Handbook of
THE PHILOSOPHY OF SCIENCE

General Editors: DOV M. GABBAY, PAUL THAGARD, AND JOHN WOODS

PHILOSOPHY *of* PSYCHOLOGY *and* COGNITIVE SCIENCE

Edited by **Paul Thagard**

GENERAL PREFACE

Dov Gabbay, Paul Thagard, and John Woods

Whenever science operates at the cutting edge of what is known, it invariably runs into philosophical issues about the nature of knowledge and reality. Scientific controversies raise such questions as the relation of theory and experiment, the nature of explanation, and the extent to which science can approximate to the truth. Within particular sciences, special concerns arise about what exists and how it can be known, for example in physics about the nature of space and time, and in psychology about the nature of consciousness. Hence the philosophy of science is an essential part of the scientific investigation of the world.

In recent decades, philosophy of science has become an increasingly central part of philosophy in general. Although there are still philosophers who think that theories of knowledge and reality can be developed by pure reflection, much current philosophical work finds it necessary and valuable to take into account relevant scientific findings. For example, the philosophy of mind is now closely tied to empirical psychology, and political theory often intersects with economics. Thus philosophy of science provides a valuable bridge between philosophical and scientific inquiry.

More and more, the philosophy of science concerns itself not just with general issues about the nature and validity of science, but especially with particular issues that arise in specific sciences. Accordingly, we have organized this Handbook into many volumes reflecting the full range of current research in the philosophy of science. We invited volume editors who are fully involved in the specific sciences, and are delighted that they have solicited contributions by scientifically-informed philosophers and (in a few cases) philosophically-informed scientists. The result is the most comprehensive review ever provided of the philosophy of science.

Here are the volumes in the Handbook:

Philosophy of Science: Focal Issues, edited by Theo Kuipers.

Philosophy of Physics, edited by Jeremy Butterfield and John Earman.

Philosophy of Biology, edited by Mohan Matthen and Christopher Stephens.

Philosophy of Mathematics, edited by Andrew Irvine.

Philosophy of Logic, edited by Dale Jacquette.

Philosophy of Chemistry and Pharmacology, edited by Andrea Woody and Robin Hendry.

Philosophy of Statistics, edited by Prasanta S. Bandyopadhyay and Malcolm Forster.

Philosophy of Information, edited by Pieter Adriaans and Johan van Benthem.

Philosophy of Technological Sciences, edited by Anthonie Meijers.

Philosophy of Complex Systems, edited by Cliff Hooker and John Collier.

Philosophy of Earth Systems Science, edited by Bryson Brown and Kent Peacock.

Philosophy of Psychology and Cognitive Science, edited by Paul Thagard.

Philosophy of Economics, edited by Uskali Mäki.

Philosophy of Linguistics, edited by Martin Stokhof and Jeroen Groenendijk.

Philosophy of Anthropology and Sociology, edited by Stephen Turner and Mark Risjord.

Philosophy of Medicine, edited by Fred Gifford.

Details about the contents and publishing schedule of the volumes can be found at http://www.johnwoods.ca/HPS/.

As general editors, we are extremely grateful to the volume editors for arranging such a distinguished array of contributors and for managing their contributions. Production of these volumes has been a huge enterprise, and our warmest thanks go to Jane Spurr and Carol Woods for putting them together. Thanks also to Andy Deelen and Arjen Sevenster at Elsevier for their support and direction.

INTRODUCTION TO THE PHILOSOPHY OF PSYCHOLOGY AND COGNITIVE SCIENCE

Paul Thagard

This *Handbook* provides informative and insightful treatments of many of the key current issues in the philosophy of psychology and cognitive science. The purpose of my introduction is to provide an overview of this branch of philosophy of science. Because I hope the handbook will be useful for scientists and students as well as philosophers, I begin with elementary expositions of the fields of psychology, cognitive science, and the philosophy of science. I then describe the fundamental epistemological, metaphysical, and ethical questions that arise in the practice of science, and sketch the forms they take in psychology and the other cognitive sciences. Finally, I preview the chapters in this Handbook.

1 PSYCHOLOGY AND COGNITIVE SCIENCE

Psychology, the investigation of human mind and behavior, goes back at least to Plato and Aristotle. However, it became an experimental science only around 1879, when Wilhelm Wundt established the first psychology laboratory. Philosophers such as Aristotle, Descartes, Locke, Hume, and Kant had developed many interesting conjectures about how minds work, but experimental tests of psychological theories began only in the latter part of the nineteenth century. Psychological theorizing was gravely limited from the 1920s to the 1950s by the dominance of behaviorism, the view that scientific psychology must restrict itself to the study of observable behavior. But since the 1960s scientific psychology has been cognitive as well as behavioral, allowing for the postulation and experimental testing of mental structures and processes that are not directly observable.

Many theoretical ideas in psychology over the last fifty years have originated from computer science, because the development of digital computers in the 1950s provided a powerful way of thinking about mechanisms by which information can be processed. Psychology is now part of cognitive science, the interdisciplinary study of mind and intelligence, which also embraces the fields of neuroscience, artificial intelligence, linguistics, anthropology, and philosophy. In the past decade, neuroscience has made rapidly increasing experimental and theoretical contributions to psychology, because of the advent of brain scanning techniques that provide ways of observing neural processes. Psychology also overlaps with artificial intelligence in the development of computational models of thinking, and with linguistics in the study of how minds understand and produce language. In addition, psychology overlaps with anthropology in the study of cultural aspects of

Handbook of the Philosophy of Science. Philosophy of Psychology and Cognitive Science
Volume editor: Paul Thagard
General editors: Dov M. Gabbay, Paul Thagard and John Woods
© 2007 Elsevier B.V. All rights reserved

social cognition, and with philosophy in a concern for fundamental issues about the nature and explanation of human minds.

The philosophy of scientific psychology must be distinguished from enterprises that have been popular in philosophy: "philosophical psychology" and armchair philosophy of mind. These enterprises assume that it is possible to learn about the mind from introspection, ordinary language, or thought experiments that generate conceptual truths about what minds must be like. In contrast, psychology is largely practiced today by means of behavioral and neural experiments that provide data used to evaluate theories about underlying mental structures and processes. Attention to introspection, everyday language and thought experiments may be useful for generating hypotheses about such structures and processes, but they are useless for evaluating hypotheses. Hence I want to sharply distinguish philosophy of psychology and cognitive science from approaches to philosophy of mind that attempt to ignore scientific developments. The point of philosophy of psychology is not to develop conceptual truths about minds, but rather to deal with philosophical issues through close attention to developments in scientific psychology and the allied areas of cognitive science. All of the essays in this handbook take a scientific rather than an armchair approach to the philosophy of psychology.

2 PHILOSOPHY AND SCIENCE

Philosophy is the investigation of fundamental questions about the nature of knowledge, reality, and morals. According to some philosophers, philosophy is inherently different from science in that it can use pure reasoning to acquire absolute certainty in answers to these questions, but no one has ever managed to generate a set of answers that seemed unassailable to anyone but their generators. The view of philosophy I prefer is naturalistic, seeing philosophy and science as strongly interconnected attempts to understand the world, including the operation of human minds. Naturalism does not proclaim that philosophy is reducible to science, because philosophical questions about the nature of knowledge, reality, and morals are more general and more normative than questions that are usually investigated in empirical science. Philosophical questions are more general in that they do not concern the particular kinds of entities and processes investigated by a science such as the plants and animals studied in biology; rather they concern the general nature of existence and our knowledge of it. Moreover, philosophical questions differ from scientific ones in being normative as well as descriptive, concerned with how the world should be as well as how it is.

Despite their generality and normativity, philosophical questions are intimately connected with descriptive, scientific ones. The connections are best seen by considering the three main branches of philosophy: epistemology, metaphysics, and ethics. Epistemology, the theory of knowledge, asks whether people know anything, what we know, and how we know it. It has often been pursued in an a priori fashion, independent of any ordinary or scientific experience, but no one has ever established any a priori truths. A more naturalistic approach to epistemology

investigates the structure and growth of human knowledge by locating it in human minds and societies that can in part be investigated by the empirical methods of psychology and the other cognitive sciences. Although science does not directly address the most general epistemological questions about whether we know anything at all and how we *ought* to go about acquiring knowledge, psychology does provide much theoretical and experimental evidence concerning the mental structure of what we know and the mental processes by which knowledge is acquired, ranging from perception to high level reasoning. Epistemology that is blind to how minds actually work is misleading and pointless.

Similarly, metaphysics, the theory of reality, is best pursued in close connection with scientific developments. Metaphysical questions concern the existence and nature of different kinds of entities, for example gods, minds, concepts, material objects, and numbers. A priori reasoning has been as fruitless in metaphysics as it has been in epistemology, but naturalistic philosophy that allies metaphysics and science provides a more promising way of pursuing questions of existence. For example, the hypothesis that God exists can be evaluated with respect to whether it adds anything to explanations from physics and biology about the origins and nature of the universe. The nature of human minds and concepts can be pursued by considering what has been learned about these entities from empirical investigations. Deliberations about the nature of material objects, numbers, causation, and space and time can also be informed by scientific theories about the nature of the universe, matter, and human minds.

For many philosophers, ethics has seemed to be the branch of philosophy most immune from empirical considerations, because it is concerned with how people ought to behave and not with how they do behave. However, the psychology and neuroscience of judgment can tell us a lot about why people make the ethical evaluations that they do and why they sometimes fall short of their own ethical standards. Ethics is indeed normative rather than descriptive, and is not reducible to or replaceable by the cognitive sciences. But ethical theories need to be coherent with the moral capacities of human beings, which requires that they pay attention to moral psychology based on empirical investigations rather than on a priori conceptual constructions.

Epistemology, metaphysics, and ethics are highly connected with each other. Views of the nature of knowledge interact with views of the nature of reality: what there is affects what can be known, and what we think can be known often affects what we think exists. Epistemology and ethics also interact in discussions of the nature of ethical knowledge. Finally, ethics and metaphysics have important interconnections, for example between views about the nature of minds as capable or incapable of free will and views about holding people morally responsible for their actions. From a naturalistic perspective, the intersections of these areas of philosophy are also intersections with areas of science. For example, moral epistemology is intimately tied to experimental and theoretical investigations of moral psychology.

If one accepts this kind of naturalistic approach to metaphysics, epistemology, and ethics, then philosophy of science becomes of central importance to philosophy. Reflections on reality, knowledge, and morals may fruitfully intersect with reflections on the development of scientific understanding. The field of philosophy of science goes back at least to Aristotle and Francis Bacon, but became richly developed only in the nineteenth century through the writings of William Whewell and others. It has flourished both with general studies of the methodology and results of science in general and with specific studies of individual sciences. I will now describe the central epistemological, metaphysical, and ethical issues that arise in general philosophy of science, and sketch the particular forms that they take in the philosophy of psychology and cognitive science.

3 EPISTEMOLOGICAL ISSUES

The premier epistemological issue in the philosophy of science might be taken to be whether science acquires any knowledge at all. However, while philosophers of science often debate the nature of scientific knowledge, few claim that it does not exist, unlike some sociologists who claim that science is merely a social construction. Similarly, few philosophers with a serious knowledge of psychology and neuroscience would claim that these fields have been utter failures in telling us anything about how the mind works.

So what is the structure of scientific knowledge? On one traditional view, scientific knowledge consists of a kind of pyramid with experimental data on the bottom, generalizations about data in the middle, and theories about the underlying causes of the data at the top. This view indeed applies roughly to psychology and neuroscience, where researchers do perform experiments that produce observable results that are statistically summarized. The summaries are usually called "effects" rather than "laws", since universal laws such as those found in physics are rare in the cognitive sciences. Theoretical cognitive science goes beyond such summaries to propose theories about the kinds of mental representations and processes that produce the various behaviors and brain activities observed in experiments.

What's missing in this attractive picture is the importance of models that provide a crucial connection between the theoretical postulation of representations and processes and the empirical phenomena that they are intended to explain. A model is a simplified representation of the entities and their interactions hypothesized by a theory. In psychology and neuroscience, models are usually computational, in keeping with the plausible contention that what brains do is a kind of computation. Computer programs have data structures and algorithms that can be taken to correspond to the mental or neural representations and processes that operate in the mind. Running the programs provides a way of determining whether the theories that postulate structures and processes can effectively produce the kinds of behaviors observed in experiments. Hence the structure of knowledge in cognitive science consists of a complex of experimental reports, effects, models, and theories.

This account of the structure of scientific knowledge should mesh with a dynamic description of the scientific methods that produce the structure. Many familiar pictures of scientific method are much too simple. According to the inductivist picture, science proceeds by collecting observations and then generalizing them into laws. Such activity does take place in psychology and neuroscience when researchers attempt to understand their data, but their aim is not usually just to collect some data and generalize about them. Rather, experiments are often done with the motivation of testing some theoretical claims about what might produce mental behavior. This motivation fits better with the hypothetico-deductivist picture, according to which scientists form hypotheses and then set out to test them. However, this picture oversimplifies the relation between theory and experiment, because what often happens is that scientists have theories that suggest experiments that provide data that suggest new theories. Thus scientific method proceeds neither from data to laws, nor from theories to experiments, but rather consists of a complex interconnected feedback process of theorizing, experimentation, and reasoning.

One of the most important issues in the epistemology of science is theory evaluation: how do scientists decide whether proposed theories should be accepted or rejected? According to inductivists, the question never arises, because scientific inference is only to generalizations of what is observed rather than to theories about underlying entities that are unobservable. More plausible is the hypothetic-deductive picture that says that a theory is used to make predictions that can either confirm or refute it. However, even in mathematically rich fields like physics, deductive relations are never enough to produce the refutation of a theory, because the failure of predictions can always be attributed to other factors such as flaws in experimental design. Refutation is even more difficult in statistical sciences such as psychology, because the experimental results do not follow deductively from theories. Rather, theories suggest models that suggest experiments that produce statistical effects, which are never definitive enough to strictly falsify a theory. This flexibility does not mean that one psychological theory is as good as another, because competing theories can be evaluated according to how well they provide explanations of a broad range of experimental results. Such competitions are a frequent and valuable part of discourse in the cognitive sciences, concerning the best explanation of such important phenomena as language learning and mental imagery. Much work in the philosophy of psychology and cognitive science has been addressed to characterizing and assessing major theoretical controversies.

I have described the relation between theories and experimental results as explanation, which presupposes an understanding of the nature of explanations. This topic is a central and controversial one in the philosophy of science, which has prominent advocates of several different accounts. Explanation has variously been construed as a deductive relation between sentences, a probabilistic relation between sentences, a relation of fit between schemas and representations of data, and an ontological relation between mechanisms and the phenomena they produce. My own view is that explanations in psychology and cognitive science are

mechanistic, in that theories postulate entities such as mental representations and neural structures that are supposed to produce observable behaviors, just as the parts of a machine produce its behaviors. On this view, the nature of explanation is a question of metaphysics as well as epistemology.

4 METAPHYSICAL ISSUES

What can we learn from science about what exists? The most general question in the metaphysics of science concerns whether we can be justified in believing in the existence of theoretical objects that are hypothesized rather than directly observed. In physics, for example, the question concerns whether we should think of entities such as atoms, electrons, and quarks as really existing, or merely as ways of talking that are convenient for predicting observed phenomena. In philosophy of psychology, this issue of scientific realism concerns whether we do or could have good grounds for believing that there really are mental representations such as rules, concepts, and connectionist networks.

The central metaphysical question in the philosophy of psychology is the mind-body problem: What is the relation between minds and brains? The commonsense theory, entrenched in many religious doctrines, is that mind is inherently different from matter. Dualism is the claim that human beings consist of two fundamentally different kinds of entities, minds and bodies. Most philosophers of mind, like most cognitive scientists, reject dualism in favor of the materialist view that mind is physical, but there is much debate about what version of materialism is most plausible. The most contentious issue concerns the nature of consciousness, particularly the kinds of qualitative experiences that people have when they perceive or feel. Whereas dualists think that conscious experience is inherently beyond the explanatory reach of natural science, reductive materialists think that it will someday be adequately explained in terms of physics, biology, and other sciences. A third position, antireductive materialism, holds that minds are inherently physical, but doubts that explanations at the biological or physical levels are likely to be achievable.

5 ETHICAL ISSUES

There are many ethical issues in the philosophy of science, but they are less frequently discussed in the philosophy of psychology than epistemological and metaphysical ones. Some issues arise in the conduct of psychological research, for example whether it is legitimate to deceive experimental subjects as often is done in research in social psychology. Informed consent is necessary if people are ethically to participate in psychological or neurological experiments.

However, the most interesting ethical questions related to psychology and cognitive science concern possible implications of research in these fields for questions of right and wrong. Commonsense morality, embedded like dualism in prevalent

religious traditions, assumes that people can legitimately be praised and blamed for actions that they perform freely. But the notion of free will is challenged by the materialist view that mental processes are brain processes constituted by biophysical mechanisms. With rapidly increasing knowledge about the proper functioning of the human brain and about the malfunctions that lead to serious mental illnesses such as schizophrenia, bipolar disorder, and psychopathy, it becomes harder to attribute immoral behavior to failures of will. More specific psychological problems arise in the context of specific ethical theories such as utilitarianism, which assumes that people are capable of taking equally into account the happiness of all people, and Kantianism, which assumes that people have autonomy in their moral decisions.

Additional ethical issues derive from the prospect of technological breakthroughs arising from the view that thinking is a kind of computation and the possibility of reverse engineering the brain to produce human-level artificial intelligence. Some writers speculate that it is only a matter of time before machines with exponentially increasing processing speed will surpass humans in computational power and intelligence. It might even be possible for people to download their mental structures from brains into more durable and easily duplicated computer hardware. These scenarios are highly fanciful, but they illustrate how future interconnection between psychology and technology may raise difficult questions about the nature of persons.

I hope that this brief sketch of the major areas of philosophy and their mutual relevance with psychology and cognitive science will provide the newcomer to philosophy with a general idea of the structure of this area of philosophy of science. For this Handbook, it was not possible to arrange contributions on all possible topics within the philosophy of psychology. Rather, my strategy was to invite philosophers who have made important contributions to philosophical discussions concerning the sciences of mind to discuss topics related to their previous work. The next section provides concise summaries of the chapters to follow.

6 OVERVIEW OF THE BOOK

Eric Dietrich discusses how minds encode information about themselves and their environments. After reviewing some philosophical answers to this fundamental problem of representation, Dietrich discusses different kinds of representations used in psychological theorizing, including traditional symbolic structures and more recent connectionist representations modeled on brain activity. He concludes that many kinds of representations are required to explain how the mind works.

Cory Wright and William Bechtel review the history of machine-based approaches to psychology since Descartes. They describe current work on mechanistic explanation that is having a large impact in current philosophy of science, and show how it applies to psychology and neuroscience, particularly in explaining motivation and reward. They show how the mechanistic approach to psychological explanation reconfigures related topics in the philosophy of psychology, especially

the question of whether psychology can be reduced to neuroscience.

Robert Wilson and Carl Craver discuss the concept of realization which has been important for understanding the relation of minds to their possible physical instantiations in brains and computers. Some philosophers have argued that we should not simply identify minds with brains, because minds could be realized in other material ways such as computers or alien physical systems. Wilson and Craver discuss several varieties of realization used by scientists and metaphysicians, and develop a useful prototype.

Robert McCauley provides a comprehensive discussion of the philosophical problem of reduction, in particular the question of whether psychology can be reduced to neuroscience. He argues against anti-reductionist views that psychology is independent of neuroscience, and also against ultra-reductionist views that see psychology as being replaced by neuroscience. His explanatory pluralism and heuristic identity theory provide rich and plausible models of current interdisciplinary developments.

Austen Clark deals with important philosophical problems that arise in the psychology of perception. He describes how psychological experiments about unconscious stages of perception reveal aspects of perception that are very different from our conscious experience. He argues that research shows the need to separate the phenomenal, qualitative character of perception from conscious awareness. This chapter is a good illustration of how scientific findings can challenge common sense conceptions of mind.

Uriah Kriegl investigates the relationship between different aspects of consciousness, challenging the distinction between phenomenal consciousness and access consciousness by showing an intimate connection between them. He elucidates the presuppositions of experimental studies of phenomena such as subliminal perception and blindsight, arguing that investigation of access consciousness is a valuable way to investigate phenomenal consciousness.

Corinne Iten, Robert Stainton, and Catherine Wearing address an important question in the scientific investigation of language. They critically examine several arguments that theories in linguistics are independent of evidence from psychology and neuroscience. They argue convincingly that theories in cognitive science, including linguistic theories, should attempt to account for a diversity of kinds of evidence. Hence linguistics and psychology are intimately related.

Jesse Prinz considers controversial issues in the psychology of emotion, including whether emotions are primarily cognitive or physiological, conscious or unconscious, and homogenous or heterogeneous. He argues that emotions all involve similar neural mechanisms, and describes several ways in which emotions can valuably contribute to reasoning.

Alvin Goldman and Kelby Mason challenge the common view that people primarily understand others by theorizing about them. They develop the alternative view that people often predict the mental states of others by simulating them. Simulation is supported by a basic neural mechanism performed by mirror neurons through which mental states in one agent can be replicated in an observer.

Valerie Gray Hardcastle gives an overview of the rapidly developing field of cognitive neuroscience, which provides neurological insights into perception, memory, and many other psychological processes. She describes the different methodologies employed in cognitive neuroscience, discussing the strengths and limitations of single cell recording and other experimental techniques.

Chris Eliasmith shows how computational models contribute to theoretical neuroscience. Psychological theories of mental representation can be deepened by showing how populations of neurons can encode and process information. He draws connections between current progress in computational neuroscience and traditional philosophical problems about the nature of meaning.

George Graham and G. Lynn Stevens provide an informative overview of the philosophy of psychopathology, the study of mental illness. They argue that conscious representational content is an important part of the origin, symptoms, and effective treatment of prototypical mental illness. Mental illnesses are not just problems in living or brain diseases.

Robert Richardson critically examines the research program of evolutionary psychology, which explains psychological mechanisms in terms of environmental demands that shaped human evolution. He holds evolutionary psychologists to the standards of adaptive explanations in evolutionary biology and argues that the evidence for adaptations offered by evolutionary psychologists is deficient in many respects.

Miriam Solomon discusses the investigation of situated cognition that has occurred in several areas of cognitive science, including psychology, anthropology, and artificial intelligence. She describes important connections between work on situated cognition and philosophical research in social epistemology and feminist philosophy of science. She concludes that normative epistemological recommendations will only apply to particular kinds of situation.

B. Jack Copeland and Diane Proudfoot provide a history of artificial intelligence and discuss philosophical issues concerning its contribution to cognitive science. These issues include the usefulness of the Turing test for whether computers can think, Searle's Chinese room argument against the possibility of strong AI, and the possibility of computations that go beyond what Turing machines can do.

I hope that readers will enjoy these investigations of important philosophical topics that arise in psychology and cognitive science.

Paul Thagard
University of Waterloo

CONTRIBUTORS

William Bechtel
Department of Philosophy, University of California, San Diego, USA.
bill@mechanism.ucsd.edu

B. Jack Copeland
Department of Philosophy, University of Canterbury, New Zealand.
jack.copeland@canterbury.ac.nz

Carl F. Craver
Department of Philosophy, Washington University, USA
ccraver@artsci.wustl.edu

Austen Clark
Department of Philosophy, University of Connecticut, USA.
austen.clark@uconn.edu

Eric Dietrich
Department of Philosophy, Binghamton University, USA.
dietrich@binghamton.edu

Chris Eliasmith
Deptartment of Philosophy, University of Waterloo, Canada.
celiasmith@uwaterloo.ca

Alvin Goldman
Department of Philosophy and Center for Cognitive Science, Rutgers University, USA.
goldman@philosophy.rutgers.edu

George Graham
Department of Philosophy, Wake Forest University, USA.
grahamg@wfu.edu

Valerie Gray Hardcastle
Department of Philosophy, Virginia Tech, USA.
valerie@vt.edu

Corinne Iten
Independent Scholar

Uriah Kriegel
Department of Philosophy, University of Arizona, USA, and Department of Philosophy, University of Sydney, Australia.

Robert N. McCauley
Department of Philosophy, Emory University, USA.
philrnm@emory.edu

Kelby Mason
Department of Philosophy, Rutgers University, USA.
mason@philosophy.rutgers.edu

Jesse J. Prinz
Department of Philosophy, University of North Carolina at Chapel Hill, USA.
jesse@subcortex.com

Diane Proudfoot
Department of Philosophy, University of Canterbury, New Zealand.
diane.proudfoot@canterbury.ac.nz

Robert C. Richardson
Department of Philosophy, University of Cincinnati , USA.
robert.richardson@uc.edu

Robert Stainton
Department of Philosophy, University of Western Ontario, Canada.
rstainto@uwo.ca

Miriam Solomon
Department of Philosophy, Temple University, USA.
msolomon@temple.edu

G. Lynn Stephens
Department of Philosophy, University of Alabama at Birmingham, USA.
stephens@uab.edu

Paul Thagard
Department of Philosophy, University of Waterloo, Canada.
pthagard@uwaterloo.ca

Catherine Wearing
Department of Philosophy, Wellesley College, USA.
cwearing@wellesley.edu

Robert A. Wilson
Department of Philosophy, University of Alberta, Canada.
rob.wilson@ualberta.ca

Cory Wright
Department of Philosophy, University of California, San Diego, USA.
cory@mechanism.ucsd.edu

REPRESENTATION

Eric Dietrich

1 THE FUNDAMENTAL PROBLEM OF REPRESENTATION

Mental representations of many sorts are used in almost all psychological explanations. An extremely important fact about psychology and cognitive science is that all of this research goes on without a solution to what is known as the *fundamental problem of representation*. Put briefly, though there is a vast quantity of on-going research dependent on representations, and though there is a vast quantity of on-going research *on* representation, no scientist knows how mental representations represent. Not deeply understanding its main theoretical tool happens in many sciences — for a while. What makes this state of affairs unusual here is that lacking an understanding of how representations represent has persisted since the inception of the cognitive sciences. There are, however, several plausible and useful theories on offer, from philosophy, psychology, and artificial intelligence. The most important of these will be presented here.

To set up the fundamental problem, we must start with the basics. Representing is a relation between two things whereby one stands in for the other. The stand-in is the representation. The two things can be objects or processes (there are four combinations). Pictures of things — people, for example — are relatively pure examples of representations. A picture of B. F. Skinner is a representation of him. Words are also important examples. The meanings of words constitute or instantiate the representation relation. The English word 'dog' means dog; i.e., 'dog' represents dogs (or a dog or that dog, etc.)

What's interesting about pictures and words is that they represent what they do because a third thing (a cognizer . . . a person, for example) interprets them that way. In the case of words, representation is achieved (in part) via convention: words mean what they do because there is a social system of rules governing their meaning. Conventions have an arbitrariness to them. The English word 'dog' could represent cats if we so chose. And a picture of B. F. Skinner is only a picture of him because it looks like him to someone else interpreting the picture that way. In this sense only, pictures are not as arbitrary as the meanings of words because the 'looks like' relation is somewhat objective. However, eventually pictures come to represent conventionally, similar to the way words do. A picture of Shakespeare or a sculpture of Socrates is conventional nowadays because there is no one now alive who knows what Shakespeare or Socrates looked like.

Handbook of the Philosophy of Science. Philosophy of Psychology and Cognitive Science
Volume editor: Paul Thagard
General editors: Dov M. Gabbay, Paul Thagard and John Woods
© 2007 Elsevier B.V. All rights reserved

Probably the most important scientific fact about the mind is that it is a representer. It constructs, processes, and stores representations of things external (and internal) to it. Minds represent things and events in their environments, for example. Thinking, which is what minds are for (like lungs are for breathing), is manipulating representations, which makes representations rather basic. The representations in minds are called *mental representations*. Mental representations are of course implemented in the brain's neural architecture (no one is quite sure how, yet), but are only sometimes best approached or understood that way (the relation between mind as a representer and brain as a representer will be discussed in the next section).

The most important fact about minds as representers is that they cannot represent in the way words and pictures do, at least not fundamentally. Specifically, minds cannot represent by interpreting mental representations as standing in for things in the world. It is useful to understand this in detail.

First, let us state the (incorrect) hypothesis explicitly: mental representation is interpretation of a representation as standing for something. We have four entities: a cognizer, M, a representation, R, a thing represented, Y, and a process of interpretation which is doing all the theoretical work. Mental representation is M's interpretation of R, as standing for Y. The question now is, 'What is interpretation?' It would be explanatorily unsatisfying in the extreme were this notion basic and foundational, for it is too abstract, metaphorical even, to be basic. It is also, most likely, a complex, made up of many parts and processes. So interpretation itself needs explaining. If the explanation of interpretation requires the notion of representation, then we've got a vicious circle: we started off using interpretation to explain representation, and now we are using representation to explain interpretation. This won't do. (One way to see this is note that vicious circles are related to infinite regresses. If a mind represented some specific dog, Fido, say, by interpreting one of its mental representations, R, as standing in for Fido, and if interpreting is representing, then the mind must have a separate mental representation of R, R', to represent R. That is, R' is interpreted as standing in for R. But if R' is interpreted as standing in for R, then the mind will need another mental representation to interpret R'. Hence, the mind will need another representation R'' to represent R'. R'' is interpreted as standing in for R'. And so on to infinity. But the mind is only a finite thing — finitely big, with only a finite amount of memory and energy.)

So, interpretation cannot be representation. What other candidates are left? No viable suggestions are known, at least none are known that sharply distinguish between representation and interpretation. Any reasonable candidate is better construed as an explanation for representation without invoking interpretation as some intermediate notion.

Representation, then, not interpretation, is the more basic notion.

Therefore, when a mind represents Fido by using mental representation R, R must somehow represent Fido straight on, and not via containing or referring to something else thus interpreted. The fundamental problem of representation

can now be stated: *How does the mind do this?* How can minds represent in an unmediated, uninterpreted way? Words, sentences, pictures, etchings, etc. all represent via interpretation. But mental representations do not, and cannot, represent that way. So how do they do it?

Currently, no one has a definitive, detailed answer to the fundamental problem (an examination of it can be found in [Bickhard and Terveen, 1995]). Indeed, for the most part, cognitive scientists do not know how minds represent (nor how brains represent). However, certainly the big picture is set. Brains represent by storing information in the patterns of activation of their neurons. Also, the way certain neural clusters activate and stay active is relevant to what and how they represent. Minds, which is what an active brain is, represent by using data structures of various sorts. The data structures are implemented as patterns of neural activation, but do not reduce to the patterns of neural activation (see below). It is widely agreed that, for both brains and their minds, a causal story is needed about how information from the environment flows into the cognizer (and vice versa). Other ingredients are also needed, since it is also widely agreed that causation or information alone will not be able to provide us with a complete theory of representation. So, for example, many cognitive scientists think that how representations relate to each other will also be a crucial aspect of the emerging theory. The details of all of this remain to be worked out, and the work is difficult — nevertheless, progress is being made.

2 THE STANDARD MODEL OF REPRESENTING AND REPRESENTATION

The discovery of the mind's ability to represent is arguably cognitive science's greatest contribution to understanding the mind. Yet, oddly, there are cognitive scientists who denounce and repudiate mental representations. It is true that there is little agreement on how representations represent, on what kinds of representations the mind uses, and how representations support cognition, but to deny that there are representations either because of this lack of agreement, or because progress in the cognitive sciences is slow (or both) is to jump to a hasty conclusion.

The term 'discovery' is appropriate in the preceding paragraph because while even the ancient Greeks wondered about the representing capacity of thoughts, it was only after the invention of the computer that a physical (or physically plausible) model of how representations work could be sustained. And it was only after this, that scientific evidence of the mind's capacity to represent was found. This is not to say that one has to be a *computationalist* to think that the mind uses representations. Computationalism is the view that thinking is a computational process, the execution of myriad complex algorithms. The view that minds represent is more fundamental than, and independent of, the view that minds are computers of some sort (decidedly *not* like any modern PC, laptop, or even a supercomputer). Computationalism entails that minds represent; but the claim that minds represent does not entail computationalism. Nothing discussed in this

entry strictly requires that minds be computational. That said, the best science we have says that minds are computational; indeed, this is the best explanation of their representational capacities.

If one is willing to be just a little generous, there is enough agreement on some issues to warrant enshrining these agreements as something worthy of the accolade 'standard model'. (In the next section, a specific, older version of the standard model, called the 'classical model', will be introduced for use in classifying representations.) On the standard model, representations are internal mental states analogous (perhaps only roughly) to data structures in computers, implemented in a brain's neural hardware, used by cognitive systems (defined next) to navigate in, think about, and manipulate objects and processes found in their environment. Cognitive systems can use their representations thusly precisely because the representations *represent* the cognitive system's environment. In this sense, representations have *semantic* or *representational content*. Any cognitive system will have many different kinds of representations as well as many different levels of representation (this is discussed in the next section). Low-level representations are built by sensory systems (e.g., the eyes), and from there, higher-level, more robust representations are built until the system is acutely perceiving its environment. Even at higher-levels, representations come in a variety of types for use in different cognitive activities, e.g., problem solving, planning, recognizing, learning, etc.

Representations accomplish all that they do by

(1) carrying information about the environment, which of course includes the body of the cognitive system, and by

(2) being related to each in other in all sorts of ways.

These ways of being related to one another include being related logically (e.g., inferentially, if Jones is speaking, Jones is alive), inductively (e.g., dogs usually have tails), phenomenologically (e.g., the smell of roses can remind one of a certain spring day), etc. To use a general term, in the mind of a given cognitive system, representations are *associated* one way or another with other representations. The technical term used here is *functional role*. Associations define a representation's functional role with respect to other representations, i.e., they define how the representation is used by other representations. (Still another technical term is *conceptual role*. It means basically the same thing as 'functional role'. 'Functional role' will be used here since it leaves open what occupies the various functional roles.) It is these associations that, together with representations' information carrying capacity, make representations so useful and robust when representing a cognitive system's environment. (1) and (2), above, are sometimes described as providing the two avenues for a representation's representational content. (1) captures the world-mind relation, and (2) captures the mind-mind relation. Both are needed to specify a representation's content. Currently, there is no agreed upon theory of representational content. But (1) and (2) are often considered to be at least necessary, if not jointly sufficient.

Though rigorously defining 'cognitive system' to the satisfaction of all or even many cognitive scientists is impossible, some sort of definition is required. Any entity that exercises some control on its environment via feedback loops must have internal states for comparing the actual states of its environment with 'desired' states. Call the desired states *goal states*, or just goals. Call the internal states which denote the actual states of the environment (within a certain degree of accuracy) *information states* because the internal states contain information. Goals have long been recognized as important components of cognitive systems (e.g., [Lewin, 1935]) and some computational modelers suggest internal states are necessary for systems to pursue goals [Agre, 1995]. Define a *system* as any entity that uses information in an attempt to satisfy its goals. Systems, on this definition, must have feedback loops (at least negative ones) because they need to determine whether or not their goals are satisfied. The goals of the system might only be implicit, i.e., not explicitly represented anywhere in the system. A thermostat controlling a heater in an enclosed building is a good example of a system. The goals of the thermostat system are not in any sense known to it. Systems are capable of making errors but not necessarily of correcting them. For example, one can hold a match under a thermostat to get it to behave as if the room temperature were quite high, though the temperature in the room might be below freezing.

All *cognitive systems* are systems. As the term is used here, all cognitive systems are *cognitive* because they use representations (e.g., structures making up information and goal states) rich in semantic content. All known cognitive systems to date are naturally occurring animals. It is unclear how simple an organism has to be not to be counted as a cognitive system. Certainly all primates are cognitive systems, as are dogs and cats. Almost certainly, all mammals are. Probably, too, are all birds and fish. Indeed, all vertebrates could reasonably be counted as cognitive systems. It is not unreasonable to include some invertebrates, notably octopi and squid, and perhaps some crustaceans (lobsters, for example, which can learn). However, the status of the rest of the arthropods is very uncertain. It is likely that one day, perhaps this century, an artificial cognitive system will be built. Some would argue that primitive, simple versions of such artificial systems already exist. This last claim is quite controversial, however, since it is a fact that we cannot yet build a robot as robust as a cockroach (even ignoring the size restriction). The controversy centers around whether this failure matters or not.

The features of the standard model that will be explicated here are the fact that representations carry information, and the nature of representation implementation. The results of the discussion of implementation will be needed to properly discuss representational typologies. Little will be said about the associations between mental representations within the mind, since that is such a large and diverse topic. However, in the next section, on typologies, models of representing that take certain types of associations as basic will be discussed.

3 REPRESENTATION AND INFORMATION

Mental representations are mediators (see [Markman and Dietrich, 2000]). They mediate between a cognitive system possessing the representation and that system's environment. Clearly, such mediation requires information flow between the system and its environment (going both ways). Theoretically, how this information flow should be understood is still a topic of debate. (The seminal work in this area is Dretske, 1981. His sort of approach is derived from the mathematical theory of information. See [Shannon, 1949].) The prime contender is couched in terms of probabilities relating the chances that something, Y, in the world has property, F, given that a certain mental representation is active (or a certain event of mentally representing is occurring). In symbols, we want to know what the value of θ should be in the formula: $Prob[F(Y)|\mathbf{R_c}(F(Y))] = \theta$ (read: 'The probability that Y has property F, given that cognitive system C represents Y as having property F). The technical details mostly revolve around where to set θ to guarantee that the mental representation is accurate, while still allowing for representational error. If the probability is set at 1, then this is equivalent to saying that representations cannot fail to represent correctly. Obviously, this is too strong. Humans, for example, routinely mistake things for something else. For example, suppose you wake up in the middle of the night and see what looks like your dog in your room, and indeed form the belief that your dog is lying there using the relevant dog representation in your head. In the morning you discover that what you took to be your dog is in fact a pile of rumpled clothes. If the conditional probability is 1 that there is a dog lying there given that you represented your dog as being there, then your dog had to have been where your clothes are. Hence, there is no room for error on this definition θ. Thus, this definition is too strong. So perhaps θ should be set at some number less than 1 (but greater than 0, obviously). Relaxing the requirement that $Prob[F(Y)|\mathbf{R_c}(F(Y))] = \theta = 1$ is an option that is available only when researchers are not interested in making information the sole basis of the semantic content of representations, or making veridical representations the sole basis of knowledge. The problem is that setting θ less than 1 generates difficulties of its own. For example, any number less than 1 will probably be arbitrary, unless experimental results can help fix it.

There are several technical solutions to these problems of where to set θ. One possibility is to give up the idea that information is important to representing. This solution is unappealing, because the idea that information is important to representing is really just a formal way of saying that there must be a *causal link* between an organism and its environment, and this seems fundamentally required if we are to make sense of how minds enable successful living in a complex world.

One interesting solution sets θ to be 1 for low-level representations, and has θ grade off from 1 as more complex, cognitively available representations are formed by higher-level perceptual systems which, among their other capacities, also produce abstractions. The idea behind this proposal is that representing that a dark, moving spot has been registered by, say, a retina is usually error-free, whereas

representing that a fly has just flown by, or even that food has just flown by, is error-prone. As higher-level representations represent not what is happening on the sensory surfaces of the cognitive system, but what is happening out in the system's environment, error is introduced. This is the cost of being able to recognize things, especially new things, out in the environment.

Consider the psychology of mistaking a pile of clothes in the dark for a dog. Presumably, in the prior history of a cognitive system that can recognize dogs, it has learned a set of perceptual features that reliably indicate the presence of a dog, including dogs it has not seen before. Enough of these features were activated by looking at the pile of clothes that the system reached the judgment that a dog was present (also, none, or not enough, of the pile-of-clothes features were active to override the activation of the dog representation). Some of the features are quite specific and are tied to the perceptual environment (e.g., the presence of particular edges, as detected in the pattern of light on the retina). Other features are more abstract, and may arise as a function of many different possible patterns of sensory stimulation (e.g., ears or snout). The more abstract features are crucial for categorization because they allow a person to recognize dogs they have never seen before. On this view, the low-level features are perfect (or close-enough to perfect) indicators that the relevant edges and patterns of light are present, whereas the abstract, higher-level features are reliable, but not perfect, indicators of the presence of a ears and snout. This is because the connection between low-level perceptual features and the presence of these more abstract features in the environment is imperfect and merely reliable (i.e., the probability of the abstract features being present given that the low-level perceptual features are is less than 1). Furthermore, most theorists in categorization believe that there are no necessary and sufficient features that determine membership in a category. Thus, the probability that an object is a member of a particular category given a set of features (even when the presence of those features in the environment is known with certainty) is less than 1. Hence, the nature of the conceptual system introduces the possibility of error in category judgments. The up-side to such a hierarchy of representations is enabling the cognitive system to recognize novel instances in the world.

The proposal then is that the conditional probability threshold, θ, should be set at 1 only for the activation of the lower-level features, and not with the higher-level representations, and certainly not with the category (e.g., 'dog') itself. The higher-level representations inherit the lower-level information (via abstraction), as well as acquire informational structure (via some sort of poorly understood construction-completing process, which might be the abstraction process itself, see below), only with some level of reliability, i.e., only with a probability less than 1.

Let $D^* = \langle DF_1, DF_2, \ldots, DF_n \rangle$ be a vector (an ordered n-tuple) of low-level representational features active in some cognitive system and sufficient for recognizing a dog. These may only be a subset of all the low-level features that can be used to recognize a dog, but let us assume that they work reliably for recognizing

dogs. Let $E^* = \langle EDF_1, EDF_2, \ldots, EDF_n \rangle$ be environmental conditions, whatever those are, that *cause* the low-level dog features to become activated. Then the following is true:

$$Prob[E^*|D^*] = 1.$$

It *is* true that there are conditions in the subject's environment that are sufficient for activating the subject's low-level dog features because, in fact, those initial features got activated. They must have been activated some way. And the only way to do it is via some environmental conditions. (Some regard this conditional probability as definitional of the relevant environmental conditions.)

Normally, or at least when things are working happily, activating those low-level features leads to the correct conclusion that a dog is present via the activation of some higher-level representations (e.g., the concepts for 'ears', 'snout', and finally 'dog'). But these latter features can be falsely activated by a rumpled pile of clothes in dim light. In this case, the high-level perceptual judgment that there is a dog present is wrong because the inheritance relationship is not perfectly reliable, but the low-level features nevertheless got activated for the right reasons, namely, there were the right low-level environmental conditions in the environment. There just didn't happen to be a dog in the environment. The low-level features get activated, not by properties of dogs *per se*, but by shapes in certain attitudes and relations.

Since the inheritance relation includes a mechanism for building abstract representations (abstractions of the lower-level ones), one problem, then, is with the abstraction mechanism. As with much of representation research, abstraction is also poorly understood. Abstraction does leave out information in the usual case, but most researchers agree that this cannot be the whole story of how abstraction works; indeed, it might not even be much of the story. Abstraction also seems to build or add in certain high-level information. Structural information is the likely candidate here: information based on relations, such as 'is-on-top-of', or 'is-at-the-front-of', etc.

Though a great many details remain to be worked out, the proposal discussed in the last few paragraphs plausibly suggests a multi-tier approach to representational error: the low-level mediating states are not in error (usually, unless a perceptual organ is malfunctioning), but higher level states introduce error in the process of completing, and creating abstractions of, the relevant low-level information. Thus, additional levels of representations may provide a road toward a solution to the problem of misrepresentation while still tying information to semantic content without making information the whole of semantic content.

The semantic content of a representation, what it means or what it represents or refers to, is a very difficult problem in cognitive science (in fact, it is not even clear if 'means', 'represents', and 'refers to' pick out the same relation). Certainly, the information content of a representation must be included in a theory of its semantic content. But what more is needed? One thing that is needed is the notion of functional role. Higher-level representations inherit the initially accurate low-level

information only with some probability less than one. What keeps the higher-level representations in sync with the cognitive system's environment and with each other is their functional role with respect to each other. Functional role is, in part, inferential role. Thus, if a representation, R, (materially) implies another representation, S, R, functions as the antecedent for S, and S as a consequent for R. R might also function as a premise in an argument for S, or in any number of other roles, including explanatory roles. The literature on this topic is large. For an entrée into it, see [Block, 1987; Field, 1977; Fodor and LePore, 1992; Harman, 1987].

4 REPRESENTATION AND IMPLEMENTATION

The relation between mind and brain frequently causes confusion and quite a bit of energetic discussion. And this confusion spills over into representation research, hampering development of theories describing the relations between low-level and high-level representations. Before turning to representations, we will consider implementation with respect to minds and brains, in general.

The crucial fact is this: Minds supervene on brains. To say that X supervenes on Y is to say that fixing all the Y properties (i.e., the properties of Y) automatically fixes the X properties. For example, temperature supervenes on molecular motion and life supervenes on biochemistry. All computer software, including the word processor used to write this entry, supervenes on the hardware of the computer the software resides in. Note, to say that X supervenes on Y, is not to say that X cannot supervene on something else, also. A given piece of software can supervene on different kinds of hardware.

Implementation is the inverse of supervenience. Minds supervene on brains; brains implement minds (or, minds are implemented in brains). Often, implementation is thought of as a human, engineering activity. Software is implemented on computers. But life is not usually thought of as being implemented in biochemistry. Certainly, 'implementation' has the human engineering meaning as one meaning, but it is clear that this technical term has a wider application: substrates implement their supervening processes. Both perspectives are necessary, and indeed they are logically tied together (X supervenes on Y if and only if Y implements X).

The confusing part of implementation is this. Though it seems otherwise, to say that Y implements X is *not* to say that X reduces to Y, or that there is nothing to X beyond Y. This is a deep, epistemic truth about minds, brains, representations, and even computer programs (and indeed, the world). Some would say it is a deep metaphysical truth, but either way it is an important truth. Implementation is not reduction because many of the crucial, identifying properties of the implemented level vanish at the implementing level. Though life supervenes on biochemistry (and though biochemistry implements life), it is not true that life reduces to biochemistry. Reduction of properties implies that the reducing properties are more fundamental than the reduced properties, more

'real', in a way. In general, this is much too strong of a claim. Biochemistry is not more real than life, nor is it more fundamental. And it is decidedly not correct when it comes to minds, brains, representations, and computers. Implemented properties are just as fundamental, just as real as the properties of the system doing the implementing. An example from computing (simplified) will make this clear.

Consider the word-processing software used to write this entry. It comes with many tools to help make a writing task easier. For example, it has a spell-checker which can find, e.g., instances of 'teh' and change them to 'the' (in fact, it has to be turned off or fooled into allowing 'teh' to be typed at all). In the high-level software implementation of the spell-checker, one can find variables bound to 'teh' which then allow it to substitute the correct spelling. These variables simply don't exist at the level of the word-processor itself, and the writer (or misspeller) has no access to them. All that happens is that 'teh' is automatically fixed in the document. Concomitantly, and more importantly, there are no misspelled words at this first implementation level. There are only bound variables and replacements of their contents with other bindings. Continuing further down, the bound variables are implemented as chunks of memory. Here, too, there are no misspelled words. But also, there are no variables. There are merely names and pointers using various techniques to get at the contents of various memory locations. Going even lower, into the hardware, engineering takes over, pointers disappear, and contents of memory locations get shuffled around via buses and the like. Below this, it is all chemistry and physics. The important properties of the hardware implementation are simply not present at the level of actually using the word-processor, and vice versa: the important properties of a word-processor, like having a spell-checker, are absent from the all the lower-level implementations. The spell-checker exists only at the level of the word-processor.

What is true of word-processors is true of all software on the computer, and indeed is true of all computational devices with more than one level. Thus, implementation is not reduction; at least not in computers. Assuming minds and brains are computational gives us the needed conclusion for minds. But, one needn't be a computationalist to reach this conclusion. One could assert that the mind is not a computer and that, hence, implementation is a different sort of thing in minds and brain than it is in computers (the plausibility of this claim won't be debated here). However, the conclusions drawn here are universal and apply to any form of implementation.

While this point about implementation is easily seen when, e.g., the word-processor is compared to its hardware implementation, it is often lost when the word-processor is compared to its high-level software implementation. Here, one can easily slide into the thought that the variable containing 'teh' and related operations *are* the spell-checker. Hence, at least some implementations are reductions, so the objection goes.

But this is incorrect. Given any two adjacent levels, $L1$ and $L2$, $L2$ implementing $L1$, if $L2$ both implements and reduces L1, then, assuming that reduction is

transitive, we have that the lowest level implements and reduces the highest level, which we have already seen is not true. The crucial property of a spell-checker is that it finds and corrects misspellings. There is no such property at the hardware level (nor at any level below the spell-checker). The properties of hardware that make it a good implementation of a word-processor's spell-checker (or not) are not more fundamental than those of the word-processor, they are merely different.

Now, with this understanding of implementation, it is easy to see that mental representations used for higher cognition (e.g., planning, and problem solving) will have properties that their implementations in lower-level representations and processes, and eventually in neural hardware, do not (and vice versa). This is a fairly deep point. It means that to understand the mind, scientists will need to use a large variety of representations, each at a different level of psychological processing. No single type of representation will work for all levels. This in turn means that psychological explanations couched at different levels will use different types of representations.

5 TYPOLOGIES OF REPRESENTATION

There are many ways to classify representations. One way is just to list the various major kinds together with the properties that distinguish them. Here, the focus will be on psychological explanation. Specifically, the different kinds of explanations and what they explain will be couched as a function of the different kinds of representations used in those explanations.

There is an older version of the standard model properly called the 'classical model'. It was committed to the view (and at times insisted) that mental representations had five extra properties that not included in the standard model (above). These are: (1) being enduring, (2) being discrete, (3) having compositional structure, (4) being abstract, and (5) being rule-governed. It is now known that representations may or may not have these properties (either singly or in various combinations) depending on where in the cognitive system they do their work: e.g., whether close to the sensorimotor boundary of the organism or deep within the reasoning, thinking part of the mind. Where in the system a type of representation operates in turn determines what kind of psychological explanation it is used in [Markman, 1999; Markman and Dietrich, 2000]. The discussion will now be run by stepping through these five properties and seeing what kinds of representations and hence explanations can be got by relaxing them in various ways.

6 VARIETIES OF REPRESENTATIONS

This section examines five properties that representations can have. These properties are: (1) being enduring, (2) being discrete, (3) having compositional structure, (4) being abstract, and (5) being rule-governed.

6.1 Enduring Versus Transient Representations

Classically, representations were thought to be enduring — they lasted through time (sometimes a considerable amount of time — decades, for example, as in long-term memory). But enduring representations, it is now known, are not sufficiently flexible to account for some of the fine details of early informational processing, notably controlling movements of the limbs and processing information on a cognitive system's sensory surface. In place of enduring representational systems, cognitive scientists posit perceptual models that involve transient, dynamic representations that capture the moment-by-moment changes in the internal states of the system. [Thelen and Smith, 1994; van Gelder and Port, 1995].

As a first example, transient representations govern a horse's gait. The gait is controlled by an elaborate dynamic interaction that involves the speed of the limbs at that moment, but is not influenced by previous gaits [Kelso, 1995]. Another example involves connectionist models. The patterns of activation on units in distributed connectionist models are transient. When a new pattern of activity arises on a set of units (perhaps on the output units as the response to a new input), the old pattern is gone. Similarly, the current state of a dynamic system is transient, and changes to some new state as the system progresses.

Although cognitive systems clearly involve transient changes in the activation of their representations, they also require states that endure over longer periods of time. Connectionist models use the weights on the connections between units as a trace of past activity. (Of course, these representations are highly distributed. No particular weight (or unit) can be identified as a symbol.) These connection weights are enduring representations; without such states, connectionist models could not learn. In general, cognitive systems must have some enduring energy landscape that determines how a new state can be derived from an existing one. This landscape determines key aspects of the behavior of the system, such as the location of attractor states.

There is an appealing insight in the view that cognitive systems do not rely exclusively on enduring representations — namely, that not all behaviors that involve representations require enduring representations. For example, the classic studies of the gill withdrawal reflex in the sea slug Aplysia have demonstrated that this reflex can be habituated with repeated stimulation. Kandel and his colleagues [Klein *et al.*, 1980] demonstrated that with repeated stimulation of the gill, the pre-synaptic motor neuron in the circuit involved in the habituation releases less neurotransmitter into the synaptic cleft than it did initially. This decrease in transmitter released appears to be mediated by a blockage of calcium channels in the pre-synaptic neuron. When stimulation of the gill is terminated, the calcium channels become unblocked and the gill withdrawal reflex returns to its original strength. In this system, the amount of transmitter released into the cleft is a representation that controls the desired strength of the gill withdrawal reflex, which is translated into an actual strength of the reflex by the post-synaptic motor neuron. This state is not enduring, however. With repeated gill stimulation, more

calcium channels become blocked, or, conversely, as the habituation stimulus is extinguished, more channels become unblocked. In either case, the reflex only has a current level of activity. It does not store past states.

So, some representations are transient. Nonetheless, not all representations can be transient. Rather, systems that learn and make use of prior behavior must have some enduring states that allow the system to react to new situations on the basis of past experience. A sea slug's gill may operate on the basis of transient states, but enduring representations are required in models of more complex cognitive behaviors, like learning cognitive science, which is definitely beyond the capacities of sea slug [Agre, 1995].

6.2 Discrete Versus Continuous Representations

Many cognitive models assume that representations are discrete (and composable, see the next section). Such representations are often called *symbols*. Discrete representations range from feature-list representations [Chomsky and Halle, 1968; Tversky, 1977] to semantic networks [Anderson, 1983; Collins and Loftus, 1975] to structured representations and schemas [Norman and Rumelhart, 1975; Schank and Abelson, 1977], and beyond. Despite the variety of proposals explaining cognition using discrete representations, there are mental processes where such representations fail to capture key aspects of cognitive processing.

Before delving into this matter, some definitions are needed. A system has discrete representations if it contains more than one representation, and the representations are bounded and uniquely identifiable. There is something right about this, of course, but there is something deeper to be said. This will be taken up at the end of this section. Continuous representations have 'no gaps'. They vary continuously with something — usually, something in the system's environment. For an example of a continuous representation, consider vibrating ear drums (vibrating ear drums are representations, they are just not *mental* representations).

It might seem that being a discrete, composable representation follows directly from being an enduring representation, but not all enduring representations make finite, localizable, and precise contributions to larger states. For example, attractor states in dynamic systems and in iterative connectionist models are enduring, but they are not discrete representations (or symbols). Attractor states are not clearly separable from each other by distinct boundaries; hence their semantic interpretations are not precise. In these systems, transitions from one state to the next occur as the result of processes like energy minimization that operate over continuous states.

The idea that not all representations in cognitive systems have to be discrete leads to natural explanations for the fact that new cognitive states are never (or almost never) exact duplicates of past ones. New states may bear some likeness to past states, but they are not identical. For example, in a distributed, connectionist model, this can be explained by the idea that new states are activation vectors that are similar to (i.e., have a high dot-product with) activation vectors that have

appeared in past states. In a dynamic systems model, a cognitive system whose behavior is characterized as a point in a continuous state space may occupy points in neighboring regions of the state space without occupying the same point more than once.

Smolensky [1988; 1991] made this point explicitly in his defense of connectionist models. He proposed that nondiscrete representations could be used to model the effects of context on cognitive processing. For example, a discrete symbol for *cup* captures very little information about cups. Rather, the information about cups that is relevant to a cognitive system changes with context. For thinking about a cup full of hot coffee, insulating properties of cups are important, and examples of cups that have handles may be highly accessible. For thinking about a ceremonial cup, its materials and design may be more important than its insulating properties. Smolensky argued that the high degree of context sensitivity in virtually all cognitive processing militates against discrete representations as the basis of cognitive states.

Clark [1993] raises a related question about where discrete representations come from. He suggests that connectionist models might actually be better suited to cognition than classical symbolic models because their sensitivity to statistical regularities in the input may help them develop robust internal states that still have most of the desirable properties discrete representations are supposed to provide.

One suggestion for how context sensitive representations might be acquired was presented by Landauer and Dumais [1997]. They describe a model of the lexicon (Latent Semantic Indexing, LSA) that stores higher-order co-occurrence relations among words in a sentence using a high-dimensional space (e.g., one with 300 dimensions). One interesting property of this model is that its performance on vocabulary tests improves both for words seen in the passages presented to it on a given training epoch and for words that were not seen during that training epoch. This improvement on words not seen is due to the general differentiation of the semantic space that occurs as new passages are presented. Despite its excellent performance on vocabulary tests (when trained on encyclopedia articles, LSA performs the TOEFL [Test of English as a Foreign Language] synonyms test at about the level of a foreign speaker of English), it contains no discrete representations corresponding to elements of word meaning.

A second line of research that suggests that continuous (or non-discrete) representations are needed in cognition focuses on the metacognitive feelings engendered by cognitive processing [Metcalfe and Shimamura, 1994; Reder, 1996]. For example, we often have a 'feeling of knowing'. When we are asked a hard question, we might not be able to access the answer to it, but we may be quite accurate at saying whether or not we would recognize the answer if we saw it. This feeling seems to be based on the overall familiarity of the retrieval cue [Reder and Ritter, 1992], as well as on partial information retrieved from memory [Koriat, 1994]. Neither of these processes seems to involve access to discrete properties of the items being processed.

Despite this solid evidence and use for continuous representations, it is known that complex cognitive systems must use some discrete representations (see [Dietrich and Markman, 2003]). For example, when people make comparisons among concepts, their commonalities and differences become available [Gentner and Markman, 1997; Markman and Gentner, 1993; Tversky, 1977]. When comparing a car and a motorcycle, people find it easy to list commonalities (e.g., both have wheels; both have engines) as well as differences (e.g., cars have four wheels, motorcycles have two wheels; cars have bigger engines than motorcycles). If a model has discrete representations, then it is possible to access those substitutional parts. In contrast, if a model does not have discrete representations, then the individual parts of a representation cannot be accessed and used by cognitive processes. For example, in a distributed connectionist model, the active representation at a given time consists of a pattern of activity across a set of units. Typically, processing involves comparing vectors using a holistic strategy like the dot product, which calculates the amount of one vector that projects on another. A scalar quantity like the dot product loses all information about what aspects of one vector are similar to another, yielding only a degree of similarity — which is useless for detailed comparing. Only when there are discrete representations can there be access to the content of the commonalities and the differences.

A similar problem arises for of Landauer and Dumais's high dimensional semantic space model described above. As discussed, this model performs well on the synonyms test from the TOEFL by finding words near to it in semantic space (in this case by having a high dot product). Its success on this test is offered as evidence of its adequacy as a model of human lexical processing that does not require discrete representations. However, this system would have difficulty with an antonyms test. Antonyms are also words that are highly related to each other, but differ along a salient dimension (e.g., 'up' and 'down' differ in direction, and 'up' and 'down' are more similar to each other than either is to 'giraffe'). Selecting the word most similar to the target would likely find synonyms, but finding the word most dissimilar to the target would find unrelated words. Determining the antonym of a word requires analyzing the parts of the relevant lexical representation, and these parts are simply not available in a purely high-dimensional semantic space.

Another reason discrete representations are crucial for cognitive processing comes from studies demonstrating that people can (depending on the circumstance) have a preference for or against exact matches along a dimension. In a study of similarity, Tversky and Gati [1982] found that people tend to give high weight to identity matches (see also [Smith, 1989]). In contrast to the predictions of mental space models of mental representation (which are continuous), Tversky and Gati found that pairs of stimuli that could each be described by values on two dimensions were considered more similar when one of the dimensions for each stimulus was an exact match than when both dimensions had similar but not identical values. Interestingly, the opposite result has been found in studies of choice [Kaplan and Medin, 1997; Simonson, 1989]. When faced with a choice, people

often select an option that is a compromise between extreme values. For example, an ideal diet meal might be one that tastes good and has very few calories. People on a diet given a choice among (1) a meal that tastes good and has many calories, (2) a meal that tastes fair and has a moderate number of calories, and (3) a meal that tastes bad, and has very few calories, are likely to select the middle option (2), because it forms a compromise between the extremes. In this case, the exact match to an ideal is foregone in favor of an option that partially satisfies multiple active goals. In these examples, the objects have a part identity rather than an overall identity. A system without discrete representations would not be able to operate on such pairs.

To summarize, it is likely that not all cognitive processes require discrete representations. Dynamic systems and connectionist models that use spatial representations are often good models of certain cognitive behaviors. These processes may often be sensitive to context. Nonetheless, the influence of context can also be modeled with discrete representations that have a small grain-size. Other processes, such as finding antonyms and making comparisons, seem to require representations that are discrete.

However, probably the deepest reason cognitive processes need discrete representations is to *categorize* inputs. That is, if a system categorizes environmental inputs then it has discrete representations (Dietrich and Markman, 2003). A system categorizes inputs if it has internal states that impose classes of sameness on those inputs. This means that the system will be in the same state for different inputs: though the inputs themselves differ, the system is unable to discern the difference.

There is a strong connection between being able to categorize and being able to discriminate. (A system discriminates inputs if it 'notices' that, and can respond to the fact that, there is more than thing in the input stream). Categorizing inputs is identifying them or classifying them. Discriminating inputs is necessary for categorizing, but not sufficient. To categorize, enduring classes of sameness are needed besides just discrete representations (see the discussion above). A system cannot discriminate between two external, environmental states with one, single continuously varying representation (think again about ear drums versus tonal representations). With a continuous representation, the system can be in different states at different times (the sets of states and times having (in theory only) the cardinality of the real numbers), but it cannot distinguish among those states. To distinguish between two external states, $S1$ and $S2$, say, the continuous infinity of intermediate states between $S1$ and $S2$ have to somehow be elided. The only way to do that is if the system has two internal representations, $R1$ and $R2$, say, that chunk all the states in some neighborhood of $S1$ in with $R1$, and all the states in some other neighborhood of $S2$ in with $R2$. This means that it will represent states of its external environment near $S1$ as *being* $S1$ (and the same with $S2$.) If a system cannot discriminate (and hence cannot categorize) inputs, but it still represents its environment, then the system's internal states must be in continuous correspondence with states in its environment. Hence the system has continuous

representations, and, accordingly, since it won't be able to discriminate inputs, it will merely respond to the continuous variation in its inputs (besides ear drums, think of the bimetal strip in old-style thermostats.)

This all suggests the following definition: A system has discrete representations if and only if it can discriminate its inputs. It follows that if a system categorizes (i.e., discriminates inputs at the level of conceptualization), then it has discrete representations.

6.3 Compositional Versus Noncompositional Representations

An important observation about cognitive processing is that concepts combine. This ability to form more complex concepts from more primitive ones is particularly evident in language, where actions are described by the juxtaposition of morphological units that represent objects (typically nouns) with other units that represent relations between those objects (typically verbs). Because we combine concepts freely and easily in this manner, it is often assumed that representations have a compositional, or role-argument, structure that facilitates combination (e.g., [Fodor and McLaughlin, 1990; Fodor and Pylyshyn, 1988]).

A central problem with cognitive processes that require a role-argument structure is that they require processes that are sensitive to the bindings between predicates and their arguments. Structure-sensitive processes are often much more complex than processes that can operate on non-compositional structures (or states). For example, when a representation is spatial, processing involves measuring distance in space (like the dot product in connectionist models). When structures are independent symbols (or sub-symbolic features), then sets of features can be compared using elementary set operations (as in Tversky's [1977] contrast model). However, when structures have bindings, a compositional procedure that is sensitive to those bindings must be created. Often, the processes proposed by cognitive scientists have been quite complex.

Consider the act of comparing two representational structures, as in the example of comparing one's representation of an atom to one's representation of the solar system. One popular and successful model of comparison, Gentner's [1983; 1989] structure-mapping theory, suggests that comparisons seek structurally consistent matches, meaning that the match must obey both parallel connectivity, and one-to-one mapping. In parallel connectivity, for each matching predicate, the arguments to those predicates must also match (e.g. the electrons corresponds to the planets, because both are revolving around something). One-to-one mapping requires that each element in one structure match at most one element in the other (e.g., mapping the electrons to the planets means they cannot also correspond to the sun). Thus, the comparison process takes into account the bindings between predicates and their arguments. A number of computational procedures for determining analogical matches have been developed [Falkenhainer *et al.*, 1989; Holyoak and Thagard, 1989; Hummel and Holyoak, 1997; Keane *et al.*, 1994].

While it may be appropriate to assume that some cognitive processes have such

complexity, it has been suggested that structure-sensitive processes are inappropriate as models of cognitive development. Indeed, a central problem that Thelen and Smith [1994] raise with the representational view of mind is that it posits representations and processes that seem far more complex than make sense on a developmental account. As one way to address this point, they discuss explanations for Baillargeon's [1987] classic studies demonstrating that infants have object permanence.

As an example, an infant is habituated to an event in which a screen is lowered, and then a car on a track rolls down a ramp to go behind the screen and re-emerges on the other side. This task is presented repeatedly, until the infant's looking time to this event subsides. Then, both possible and impossible test events are presented. In the possible event, a block sits behind the track, the screen lowers, and the car again rolls down the ramp behind the screen and emerges on the other side. In the impossible event, a block sits *on* the track, the screen lowers, and the car rolls down the ramp to behind the screen and emerges on the other side (the block having been secretly removed). Infants show greater looking time to the impossible event than to the possible one. This result is interpreted as a recognition that the block continues to exist behind the screen and should have stopped the progress of the car.

An explanation of this event involving a compositional symbol system would assume that infants store specific relationships such as that the block was *on* the track or the block was *behind* the track, as well that the car was *on* the track. It is critical to this explanation that the child can make a distinction between the consequences of the block being *on* the track and the block being *behind* the track. The process underlying this behavior might be specific to cars and tracks (or perceptual objects of particular types); or it might be general to moving objects and obstructions.

Thelen and Smith suggest that this explanation grants too much knowledge to an infant. In particular, they reason that if infants could form a representation of the scene, then it is not clear why they should require a sequence of habituation trials in order to form their representation. Moreover, if infants have such elaborate knowledge of objects, it is not clear why they should act as if hidden objects did not exist in traditional Piagetian object permanence tasks. Thus, Thelen and Smith suggest that the symbolic account of this task provides a gross description of infants' behavior, but fails to explain the details.

In place of a symbolic model, Thelen and Smith propose a dynamic systems account, which eschews classical representations in favor of continuous ones. They suggest that the infant reacts to regularities detected by the visual system. The infant visual system is assumed to have subsystems that specify *what* objects exist in the world and *where* those objects are located. These outputs form a state space (which is continuous). The impact of habituation is to form an expected trajectory through the state space. Then, during the test events, Thelen and Smith assume, the child dishabituates to the impossible event because its trajectory starts out similar to that of the habituation event but then diverges from it at some point. In

contrast, the trajectory of the possible event does not diverge sufficiently enough from that of the habituation event (though it does diverge because the habituation event used no block at all), so no dishabituation is observed.

This dynamic systems account is intriguing, but it cannot explain the infants' behavior without positing a complex compositional structure — a representation with a role-argument structural description. In the highly impoverished form of the 'what' and 'where' systems in the example, it is not clear what information is supposed to be captured in the visual array. However, even if a complex array of values were sufficient to model the output of these two systems, there is no account of why the trajectory divergence caused by having a block on the track is more surprising than the trajectory divergence caused by having the block behind the track. That is, Thelen and Smith provide no account of how an undifferentiated notion of trajectories in a state space distinguishes between trajectory differences that matter and those that do not. This suggests that infants' behavior in this case must reflect a recognition of the spatial relationships between objects, and that augmenting the dynamical systems view to account for these data will ultimately require the addition of a capacity for storing discrete, compositional spatial relations. That is, representations with a role-argument structure will be needed.

A brief examination of research in visual object recognition suggests that visual representations may be profitably characterized as having components that encode spatial relations between parts. Kosslyn [1994] marshals behavioral, computational and neuropsychological evidence in favor of the hypothesis that there are two different modes that the visual system uses to describe relationships between elements in images. The right hemisphere system describes the visual world in terms of metric aspects, and the left hemisphere system uses qualitative relations between elements to describe the world (although see [Ivry and Robertson, 1998], for an alternative explanation of these findings). Other behavioral and computational evidence that visual object recognition requires attention to relations between parts in images comes from [Biederman, 1987; Hummel and Biederman, 1992; Palmer, 1977]. For example, Biederman [1987] suggests that representations denoting objects consist of primitive shapes connected by spatial relations (see also [Marr, 1982]). As evidence, he demonstrates that the ability to recognize objects in line drawings is disrupted more by eliminating information at the junctions of line segments (which carries information about relations between parts) than by eliminating an equivalent amount of line information between the joints. This work further suggests that the visual array required by Thelen and Smith's explanation of the object permanence studies is likely to involve some relational elements. This interpretation is reinforced by the observation that spatial prepositions refer to spatial relations that are abstracted away from many specific details of objects (e.g., [Herskovits, 1986; Landau and Jackendoff, 1993; Regier, 1996]).

Finally, as discussed at the beginning of this section, compositional structure seems necessary for models of linguistic competence. Many linguists and philosophers have pointed out that people effortlessly distinguish between sentences like 'The Giants beat the Jets' and 'The Jets beat the Giants' Even connectionist mod-

els of phenomena like this make use of structured internal states (e.g., [Chalmers 1990; Elman,1990; Pollack, 1990]). We are also able to keep track of others' beliefs when they are explicitly stated. Thus, Jack may believe that the Giants beat the Jets last week, but that Jill believes the opposite. A 'propositional attitude' like belief requires not only that Jack be able to encode the elements in the original proposition itself (that the Giants beat the Jets), but that Jill believes the opposite proposition (so Jack must be able to represent the relevant meta-proposition). This sort of processing admittedly requires effort and does not develop immediately [Perner, 1991; Wellman, 1990], but it eventually becomes a significant part of human linguistic competence. It seems unlikely that these abilities could be modeled without representations that have a role-argument structure.

In sum, one insight underlying the proposal that representations do not have compositional structures is that such representations require significant effort to construct, and also significant effort to process. This complexity seems to go beyond what is required to carry out many cognitive tasks. It would be an overgeneralization, however, to conclude that compositional structures are not needed at all. Tasks as basic as those that demonstrate object permanence in infants and processes like object recognition clearly involve at least rudimentary relations between objects in a domain. A model that has no capacity for role-argument binding cannot explain the complexity of such higher-level cognitive and linguistic processing.

Noncompositional representations, then, are likely used by a cognitive system's low-level perceptual systems, such as the retina and early visual information processing systems, such as edge-detection. Also, such representations might be used in coordinating bodily movements.

This is a good place to point out that one of the important aspects of noncompositional, continuous, and transient representations is that they nicely explain the dynamical aspects of sensorimotor interaction. But it would be wrong to conclude from this that enduring, discrete, compositional representations are poor choices for explaining any of the dynamical aspects of cognition. Quite the opposite is true. Enduring, discrete, compositional representations can be used in a wide variety of explanations of dynamical, high-level, cognitive activity — activity which is far-removed from sensorimotor activity. The conceptual change during analogy making is a good example (see [Dietrich *et al.*, 2003]). For an extended discussion of cognitive dynamics, see [Dietrich and Markman, 2000].

6.4 Abstract Versus Nonabstract Representations

A common intuition is that abstract thought is central to cognitive processing. At one level, it is trivially true that representations are abstract. The world itself does not enter into our brains and affect behavior. Even sense data are the result of neural transformations of physical stimuli that reach our sense organs. Hence, the question being raised is more accurately cast as a search for the level of abstraction that characterizes representations. The classical assumption is that

the information we store is extremely abstract, and hence that it applies across domains. Indeed, when a logical statement like $P \rightarrow Q$ is written, it is assumed that any thinkable thought can play the role of P or Q. This assumption has been called into question.

One source of the attack on highly abstract, stored information comes from demonstrations that people's performance on logical reasoning tasks is often quite poor. For example, in the classic Wason selection task [Wason and Johnson-Laird, 1972], people are told to assume that they are looking at a set of four cards that all have a number on one side, and a letter on the other, and that they must select the smallest set of cards they would have to turn over in order to test the rule 'If there is a vowel on one side of the card, then there is an odd number on the other side'. The four cards show an A, 4, 7 and J, respectively. In this task, people appear sensitive to the logical schema called modus ponens ($P \rightarrow Q, P \vdash Q$), as virtually all people state that the card with the A on it must be turned over. In contrast, people generally seem insensitive to modus tollens ($P \rightarrow Q, \sim Q \vdash\sim P$), as few people suggest that the card with the 4 must be turned over. Further support for this finding comes from studies of syllogistic reasoning in which people exhibit systematic errors in their ability to identify the valid conclusions that follow from a pair of premises [Johnson-Laird, 1983].

These errors have been explained by appealing to abstract logical rules that differ in their ease of acquisition [Rips, 1994]. However, much work has focused on more content-based structures that might be used to solve logical problems. For example, Johnson-Laird and his colleagues [Johnson-Laird, 1983; Johnson-Laird *et al.*, 1989; Johnson-Laird and Byrne, 1991] have suggested that people solve logical reasoning problems by constructing mental models that contain familiar objects and use inference rules derived from familiar situations. Consistent with this claim, it has been demonstrated that people's performance on logical reasoning tasks like the Wason selection task is much better when the situation is specific and familiar than when it is abstract. Studies have demonstrated that people perform well on the Wason task when the scenario involves social rules like permission, obligation or catching cheaters [Cheng and Holyoak, 1989; Cosmides, 1989]. Although debate continues over the exact nature of people's reasoning processes, there is general agreement that context has a strong influence on how people reason.

The context-bound nature of reasoning has led some researchers to assume that the bulk of human reasoning is carried out by fairly concrete representations. The robots developed by Brooks [1991] embody this assumption. Brooks's robots do not form extensive structures to describe their environments; they only use information that is immediately available and store only transient information as they navigate the world. Modules in the robot communicate with each other only by allowing one module to inhibit the activity of another, without passing any information between them.

A related approach is taken in psychology in the study of situated action [Clancey, 1997]. For example, Hutchins [1995] performed a far-reaching study of navigators aboard naval ships. He argues that the complex task of plotting a course for a ship involves deep cognitive work by (at least some) of the participants, but it also requires extensive use of tools and of shared information processing. No individual has an abstract structure of the entire navigation task. Instead, the task itself is structured by the tools used to complete it (such as maps and protractors).

Cognitive linguists have also taken the view that mental structures are not entirely abstract. Langacker [1986] suggests that syntactic structures reflect concrete, modal information. The encoding of prepositions like 'above' and 'below' are assumed to be tied to structures that encode spatial information rather than simply reflecting abstract structures. The linguistic representations are symbolic, but they are assumed to be symbols that are closely tied to perceptual aspects of the world. This contrasts with the amodal verbal symbols often used in linguistic models. Thus, cognitive linguistics assumes a much closer connection between syntax and semantics than does classical linguistics.

Mainstream research in cognitive science has also shifted away from the use of abstract logical forms toward more context-based approaches. In the study of categorization, significant progress has been made by assuming that people store specific episodes rather than abstractions of category structure [Barsalou 1999; Brooks, 1978; Medin and Schaffer,1978; Nosofsky, 1986]. Research on problem solving has demonstrated that people solve new problems by analogy with previously encountered problems rather than on the basis of abstracted solution procedures [Bassok *et al.*, 1998; Novick, 1990; Reed and Bolstad, 1991; Ross, 1984]. For example, Bassok *et al.*, [1998] found that arithmetic word problems written by college undergraduates were affected by the content of the word problems. If arithmetic knowledge were truly abstract, then these content effects would not be expected. In AI, the field of case-based reasoning has taken as a fundamental assumption that it is easier to store, retrieve and tweak existing cases of actual events than to form abstract rules, derive procedures for recognizing when they should be used and then adapting them to be applied in a new situation [Kolodner, 1993; Schank *et al.*, 1994].

These examples demonstrate that there are unlikely to be many general-purpose, context-free representations and associated processes that are ready to be deployed in whatever domain they are needed. The fact that cognitive processing generally shows strong effects of context means only that most representations contain some information about the context in which they were formed. It does not mean that there is no highly abstract information stored in some representation somewhere. It is likely that the information within an individual may differ in its degree of abstractness. Some types of inference schemas (like modus ponens) seem so obvious and independent of the domain that we may very well store them as abstract rules (see Rips 1994 for a similar discussion). Other types of inferences seem to rely heavily on the domain. The main question to be answered by cognitive science is

how many kinds of representations are abstract and how many are concrete, and what level of abstraction is used by different cognitive processes. Currently, the balance seems to favor concreteness for many cognitive processes.

6.5 Rule-Governed Versus Non-rule-governed Representations

Many classical models of cognitive processing posit rules. For example, Piagetian stage theory was based on the idea that the end state of development is a competence with formal operations. Chomskian linguistics (the earlier versions) assumed that grammar involved rules that transformed a deep structure into a surface structure. Classical AI assumed that problems could be solved by applying operators in which the presence of a set of antecedent conditions caused a consequent that changed values of representations and controlled effectors that interact with the world.

In classical AI reasoning systems, the rules are generally inference schemas or productions [Anderson, 1983; Newell, 1990; Pollock, 1994]. The rules may also be statistical procedures that are supposed to capture crucial elements of expert reasoning behavior. In some AI research on problem solving, a problem is cast as a discrepancy between a beginning and an end state, and problem solvers are assumed to have an array of rules (or operators) that can be applied that reduce this discrepancy [Newell and Simon, 1963]. On this view, problem solving is a search through a problem space generated by the application of rules to the current state.

A rule-governed approach is also evident in developmental psychology. As Smith and Sera [1992] point out, many developmental theories begin with the adult behavior as the expected end-state and then develop a theory that leads inexorably from an inchoate beginning state to an orderly adult competence. Adult behavior is often described in terms of a system of rules and children are then monitored until they show sensitivity to the proper set of adult rules. For example, in Piagetian studies of the balance beam, children are given long blocks of various shapes and encouraged to try to balance them on a fulcrum. They are monitored for the development of the correct rule that the downward force of a weight is a function of the weight and the distance of the weight from the fulcrum. In this task, children's behavior is often described as the development of intermediate (and incorrect) rules like 'the fulcrum must always be in the center'. On this view, developmental milestones consist of the acquisition of particular rules.

In many ways, these models of development resemble linguistic models. A central tenet of modern linguistics is that syntactic structure is guided by a highly abstract and universal set of rules determining which sentences are grammatical in a given language. Linguistics is concerned primarily with linguistic competence — an accurate description of the grammar of a given language. Psychologists who have adopted this framework (and have studied linguistic performance) have assumed that there is some mental representation of these syntactic structures. On this view, the sentences of a language are constructed through the application of

grammatical rules. Many psycholinguistic models posit processes in which rules are applied to linguistic input that allow the structure of the sentence to be determined from its surface form.

Anti-classicalist arguments have often centered on rules. A central argument by Thelen and Smith [1994] is that cognitive development does not involve the acquisition of rules. They use the development of locomotor ability as an example. As they point out, very young infants exhibit a stepping motion when their feet are stimulated if their weight is supported externally. This ability later seems to disappear, only to re-emerge still later in development. Many theories of locomotion used this description of the behavior of the average child as the basis of theories of motor development. These theories often posit maturational changes that permit the observed behaviors to occur.

Thelen and Smith [1994, Thelen, 1995] argue that the rule-based view does not properly characterize children's development. Children supported in water exhibit the same stepping behavior as younger infants supported out of water, leading to the conclusion that increases in the weight of the legs may be causing the observed cessation of stepping behavior. Support for this comes from studies in which leg weights are attached to very young infants, which causes the stepping movements to stop. The fine details of locomotor behavior suggest that children's development is guided not by the acquisition of a small set of rules, but rather by the interaction of multiple physical and neural constraints. Behavior is guided in part by the maturation of brain and tissue. It is also guided by a child's interaction with the outside world. A variety of factors must come together to shape development. Finally, the rule-based view of motor development focuses on the progression of the average child, and ignores individual differences. In contrast, the dynamic view of motor development considers individual variation to be important data (see also [Kelso, 1995]).

These examples provide compelling evidence that rules are not needed in explanations of many cognitive processes. Many systems can be described by rules, but that is not the same thing as using rules to carry out a process. For example, the steam-engine governor has a representation (the speed with which the governor spins), and a mechanism that makes use of that state (a combination of arms and levers that closes the valve as the height of the arms increases). But, the system is not checking the state of memory in order to determine the appropriateness of a rule. Thus, the system is not actually using a rule to carry out its behavior.

Although many cognitive processes do not need rules, it does not follow that rules are not a part of cognitive systems. Some parts of cognitive processes seem like good candidates for being rule-based systems. For example, there have been no convincing accounts to date that the statistical structure of a child's linguistic input is sufficient to lead them to acquire a grammatical system consistent with their language. Furthermore, there have been some impressive demonstrations of rule-use. For example, Kim, Pinker, Prince, and Prasada [1991] demonstrated that verbs derived from nouns are given a regular past-tense form, even when the noun is the same as a verb that takes an irregular past-tense form. So in describing a

baseball game, an announcer will say 'Johnson flied out to center field his first time up' rather than 'Johnson 'flew' out...', because the verb 'flied' is derived from the noun fly [ball] rather than the verb 'to fly'. Furthermore, Marcus *et al.* [1995] have demonstrated that the way a verb is given its past tense form or a noun its plural form need not be a function of the frequency of that morphological ending or the similarity of the verb or noun to known verbs and nouns in the language. These findings support the view that at least some aspects of grammar are mediated by rule-governed processes.

This discussion of rules requires one technical point, and one methodological point related to it. If the computational hypothesis about the nature of cognition is correct (and it is a hypothesis, not a loose metaphor [Dietrich, 1990; 1994]), then it *must* be possible in principle to model cognition using rules, because it is a theorem in computability theory that a rule-based machine can do everything a Turing machine can do. Put another way, if cognition involves the execution of algorithms then, at least in principle, we can model all those algorithms using rule execution. Hence, arguments like Thelen's and Smith's are really founded on pragmatic constraints only.

Even if the computational hypothesis is wrong, and cognition is carried out in some non-computational way, it is still reasonable to *use* rules in cognitive models when rules provide a descriptive language that is both explanatorily adequate and easy to use. This use of rules is akin to a programmer's use of high-level programming languages like C++ or Java rather than assembly language. At present, rule-based systems should not be removed as a technique for cognitive explanation when all that has been demonstrated so far is that *some* cognitive processes are not well characterized as being rule-based and that cognitive science often uses rules that are too coarse-grained.

In conclusion, many kinds of representations are required to explain how the mind works, depending on what mental, perceptual, or cognitive process is the explanatory focus, and some representations differ quite radically from others. No one kind of representation can explain all that needs to be explained about our mental lives. But virtually nothing can be explained without them. This strongly suggests that there are many kinds of representations being *used* by the mind when it does the large number of things it does.

BIBLIOGRAPHY

[Agre, 1995] P. E. Agre. Computational research on interaction and agency. *Artificial Intelligence*, 72:1–52, 1995.

[Anderson, 1983] J. R. Anderson. A spreading activation theory of memory. *Journal of Verbal Learning and Verbal Behavior*, 22:261–295, 1983.

[Baillargeon, 1987] R. Baillargeon. Object permanence in 3.5- and 4.5-month old infants. *Developmental Psychology*, 23:655–664, 1987.

[Barsalou, 1999] L. W. Barsalou. Perceptual symbol systems. *Behavioral and Brain Sciences*, 22, 1999.

[Bassok *et al.*, 1998] M. Bassok, V. M. Chase, and S. A. Martin. Adding apples and oranges: Semantic constraints on application of formal rules. *Cognitive Psychology*, 35(2):99–134, 1998.

[Bickhard and Terveen, 1995] M. Bickhard and L. Terveen. *Foundational Issues in Artificial Intelligence and Cognitive Science*. The Netherlands: North-Holland Elsevier, 1995.

[Biederman, 1987] I. Biederman. Recognition-by-components: A theory of human image understanding. *Psychological Review*, 94(2):115–147, 1987.

[Block, 1987] N. Block. Functional role and truth conditions. *Proceedings of the Aristotelian Society*, 61:157–181, 1987.

[Brooks, 1978] L. Brooks. Non-analytic concept formation and memory for instances. In E. Rosch and B. B. Lloyd, editors, *Cognition and Categorization*, pages 169–211. Hillsdale: Lawrence Erlbaum Associates, 1978.

[Brooks, 1991] R. A. Brooks. Intelligence without representation. *Artificial Intelligence*, 47:139–159, 1991.

[Chalmers, 1990] D. J. Chalmers. Why fodor and pylyshyn were wrong: The simplest refutation. In *The Proceedings of the Twelfth Annual Conference of the Cognitive Science Society*. Cambridge, MA, 1990.

[Cheng and Holyoak, 1989] P. W. Cheng and K. J. Holyoak. On the natural selection of reasoning theories. *Cognition*, 33:285–313, 1989.

[Chomsky and Halle, 1991] N. Chomsky and M. Halle. *The sound pattern of English*. Cambridge, MA: The MIT Press, 1991.

[Clancey, 1997] W. J. Clancey. *Situated action: On human knowledge and computer representations*. New York: Cambridge University Press, 1997.

[Clark, 1993] A. Clark. *Associative Engines: Connectionism, concepts, and representational change*. Cambridge, MA: The MIT Press, 1993.

[Collins and Loftus, 1989] A. M. Collins and E. F. Loftus. A spreading-activation theory of semantic priming. *Psychological Review*, 82(6):407–428, 1989.

[Cosmides, 1989] L. Cosmides. The logic of social exchange: Has natural selection shaped how humans reason? studies with the wason selection task. *Cognition*, 31(3):187–276, 1989.

[Dietrich and Markman, 2000] E. Dietrich and A. Markman, editors. *Cognitive Dynamics: Conceptual change in humans and machines*. Mahwah, NJ: Lawrence Erlbaum, 2000.

[Dietrich and Markman, 2003] E. Dietrich and A. Markman. Discrete thoughts: Why cognition must use discrete representations. *Mind and Language*, 18(1):95–119, 2003.

[Dietrich *et al.*, 2003] E. Dietrich, A. Markman, and M. Winkley. The prepared mind: The role of representational change in chance discovery. In Yukio Ohsawa and Peter McBurney, editors, *Chance Discovery by Machines*, pages 208–230. Berlin: Springer-Verlag, 2003.

[Dietrich, 1990] E. Dietrich. Computationalism. *Social Epistemology*, 4(2):135–154, 1990.

[Dietrich, 1994] E. Dietrich. Thinking computers and the problem of intentionality. In E. Dietrich, editor, *Thinking computers and virtual persons*. San Diego: Academic Press, 1994.

[Dretske, 1981] F. Dretske. *Knowledge and the Flow of Information*. Cambridge, MA: MIT/Bradford, 1981.

[Elman, 1990] J. L. Elman. Finding structure in time. *Cognitive Science*, 14(2):179–212, 1990.

[Falkenhainer *et al.*, 1989] B. Falkenhainer, K. D. Forbus, and D. Gentner. The structure-mapping engine: Algorithm and examples. *Artificial Intelligence*, 41(1):1–63, 1989.

[Field, 1977] H. Field. Logic, meaning and conceptual role. *Journal of Philosophy*, 69:79–408, 1977.

[Fodor and LePore, 1992] J. Fodor and E. LePore. *Holism: A Shoppers' Guide*. Oxford: Blackwell, 1992.

[Fodor and McLaughlin, 1990] J. Fodor and B. McLaughlin. Connectionism and the problem of systematicity: Why smolensky's solution doesn't work. *Cognition*, 35:183–204, 1990.

[Fodor and Pylyshyn, 1988] J. A. Fodor and Z. W. Pylyshyn. Connectionism and cognitive architecture: A critical analysis. *Cognition*, 28:3–71, 1988.

[Forbus *et al.*, 1995] K. D. Forbus, D. Gentner, and K. Law. Mac/fac: A model of similarity-based retrieval. *Cognition*, 19:141–205, 1995.

[Gentner and Markman, 1997] D. Gentner and A. B. Markman. Structural alignment in analogy and similarity. *American Psychologist*, 52(1):45–56, 1997.

[Gentner, 1983] D. Gentner. Structure-mapping: A theoretical framework for analogy. *Cognitive Science*, 7:155-170, 1983.

[Gentner, 1989] D. Gentner. The mechanisms of analogical learning. In Y. S. Vosniadou and A. Ortony, editors, *Similarity and Analogical Reasoning*, pages 199–241. New York: Cambridge University Press, 1989.

[Harman, 1987] G. Harman. (non-solipsistic) conceptual role semantics. In E. Lepore, editor, *New Directions in Semantics*. London: Academic Press, 1987.

[Herskovits, 1986] A. Herskovits. *Language and spatial cognition: An interdisciplinary study of the prepositions in English*. New York: Cambridge University Press, 1986.

[Holyoak and Thagard, 1989] K. J. Holyoak and P. Thagard. Analogical mapping by constraint satisfaction. *Cognitive Science*, 13(3):295–355, 1989.

[Hummel and Biederman, 1992] J. E. Hummel and I. Biederman. Dynamic binding in a neural network for shape recognition. *Cognitive Science*, 99(3):480–517, 1992.

[Hummel and Holyoak, 1997] J. E. Hummel and K. J. Holyoak. Distributed representations of structure: A theory of analogical access and mapping. *Psychological Review*, 104(3):427–466, 1997.

[Hutchins, 1995] E. Hutchins. *Cognition in the wild*. Cambridge, MA: The MIT Press, 1995.

[Ivry and Robertson, 1998] R. B. Ivry and L. C. Robertson. *The two sides of perception*. Cambridge, MA: The MIT Press, 1998.

[Johnson-Laird and Byrne, 1991] P. N. Johnson-Laird and R. M. J. Byrne. *Deduction*. Hillsdale, NJ: Lawrence Erlbaum Associates, 1991.

[Johnson-Laird *et al.*, 1989] P. N. Johnson-Laird, R. M. Byrne, and P. Tabossi. Reasoning by model: The case of multiple quantification. *Psychological Review*, 96(4):658–673, 1989.

[Johnson-Laird, 1983] P. N. Johnson-Laird. *Mental Models*. New York: Cambridge University Press, 1983.

[Kaplan and Medin, 1997] A. S. Kaplan and D. L. Medin. The coincidence effect in similarity and choice. *Memory and Cognition*, 25(4):570–576, 1997.

[Keane *et al.*, 1994] M. T. Keane, T. Ledgeway, and S. Duff. Constraints on analogical mapping: A comparison of three models. *Cognitive Science*, 18:387–438, 1994.

[Kelso, 1995] J.A.S. Kelso. *Dynamic patterns: The self-organization of brain and behavior*. Cambridge, MA: The MIT Press, 1995.

[Kim *et al.*, 1991] J. J. Kim, S. Pinker, A. Prince, and S. Parasada. Why no mere mortal has ever flown out to center field. *Cognitive Science*, 15(2):173–218, 1991.

[Klein *et al.*, 1980] M. Klein, E. Shapiro, and E. R. Kandel. Synaptic plasticity and the modulation of the ca++ current. *Cognitive Science*, 89(2):117–157, 1980.

[Kolodner, 1993] J. Kolodner. *Case-based Reasoning*. San Mateo, CA: Morgan Kaufmann Publishers, 1993.

[Koriat, 1994] A. Koriat. Memory's knowledge of its own knowledge: The accessibility account of the feeling of knowing. In J. Metcalfe and A.P. Shimamura, editors, *Metacognition*. Cambridge, MA: The MIT Press, 1994.

[Kosslyn, 1994] S. M. Kosslyn. *Image and Brain*. Cambridge, MA: The MIT Press, 1994.

[Landau and Jackendoff, 1993] B. Landau and R. Jackendoff. 'what' and 'where' in spatial language and spatial cognition. *Behavioral and Brain Sciences*, 16(2):217–266, 1993.

[Landauer and Dumais, 1997] T. K. Landauer and S. T. Dumais. A solution to plato's problem: The latent semantic analysis theory of acquisition, induction, and representation of knowledge. *Psychological Review*, 104(2):211–240, 1997.

[Langacker, 1986] R. W. Langacker. An introduction to cognitive grammar. *Cognitive Science*, 10(1):1–40, 1986.

[Lewin, 1935] K. Lewin. *A dynamic theory of personality*. New York: McGraw-Hill, 1935.

[Marcus *et al.*, 1995] G. F. Marcus, U. Brinkmann, H. Clahsen, R. Wiese, and S. Pinker. German inflection: The exception that proves the rule. *Cognitive Psychology*, 29(1):189–256, 1995.

[Markman and Dietrich, 2000] A. Markman and E. Dietrich. In defense of representations. *Cognitive Psychology*, 40:138–171, 2000.

[Markman and Gentner, 1993] A. B. Markman and D. Gentner. Structural alignment during similarity comparisons. *Cognitive Psychology*, 25(4):431–467, 1993.

[Markman, 1999] A. B. Markman. *Knowledge representation*. Mahwah, NJ: Lawrence Erlbaum Associates, 1999.

[Marr, 1982] D. Marr. *Vision: A Computational Investigation into the Human Representational and Processing of Visual Information*, San Fancisco, W. H. Freeman, 1982.

[Medin and Schaffer, 1978] D. L. Medin and M. M. Schaffer. Context theory of classification. *Psychological Review*, 85(3):207–238, 1978.

[Metcalfe and Shimamura, 1994] J. Metcalfe and A. P. Shimamura, editors. *Metacognition: Knowing about knowing*. Cambridge, MA: The MIT Press, 1994.

[Newell and Simon, 1963] A. Newell and H. A. Simon. Gps: A program that simulates human thought. In E. A. Feigenbaum and J. Feldman, editors, *Computers and Thought*. R. Oldenbourg KG, 1963.

[Newell, 1990] A. Newell. *Unified Theories of Cognition*. Cambridge, MA: Harvard University Press, 1990.

[Norman and Rumelhart, 1975] D. A. Norman and D. E. Rumelhart. *Explorations in Cognition*. San Francisco: W.H. Freeman, 1975.

[Nosofsky, 1986] R. M. Nosofsky. Attention, similarity and the identification-categorization relationship. *Journal of Experimental Psychology: General*, 115(1):39–57, 1986.

[Novick, 1990] L. R. Novick. Representational transfer in problem solving. *Psychological Science*, 1(2):128–132, 1990.

[Palmer, 1977] S. E. Palmer. Hierarchical structure in perceptual representations. *Cognitive Psychology*, 9:441–474, 1977.

[Perner, 1991] J. F. Perner. *Understanding the representational mind*. Cambridge, MA: The MIT Press, 1991.

[Pollack, 1990] J. B. Pollack. Recursive distributed representations. *Artificial Intelligence*, 46(1-2):77-106, 1990.

[Pollock, 1994] J. L. Pollock. Justification and defeat. *Artificial Intelligence*, 67:377–407, 1994.

[Reder and Ritter, 1992] L. M. Reder and F. E. Ritter. What determines initial feeling of knowing? familiarity with question terms, not with the answer. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 18(3):435–451, 1992.

[Reder, 1996] L.M. Reder. *Implicit memory and metacognition*. Mahwah, NJ: Lawrence Erlbaum Associates, 1996.

[Reed and Bolstad, 1991] S. K. Reed and C. A. Bolstad. Use of examples and procedures in problem solving. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 17(4):753–766, 1991.

[Regier, 1996] T. Regier. *The human semantic potential*. Cambridge, MA: The MIT Press, 1996.

[Rips, 1994] L. J. Rips. *The Psychology of Proof: Deductive reasoning in human thinking*. Cambridge, MA: The MIT Press, 1994.

[Ross, 1984] B. H. Ross. Remindings and their effects in learning a cognitive skill. *Cognitive Psychology*, 16:371–416, 1984.

[Schank and Abelson, 1977] R. C. Schank and R. Abelson. *Scripts, Plans, Goals and Understanding*. Hillsdale, NJ: Lawrence Erlbaum Associates, 1977.

[Schank *et al.*, 1994] R. C. Schank, A. Kass, and C. K. Riesbeck. *Inside case-based explanation*. Hillsdale, NJ: Lawrence Erlbaum Associates, 1994.

[Shannon, 1949] C. Shannon. *The Mathematical Theory of Communication*. Urbana, IL: University of Illinois Press, 1949.

[Simonson, 1989] I. Simonson. Choice based on reasons: The case of attraction and compromise effects. *Journal of Consumer Research*, 16:158–174, 1989.

[Smith and Sera, 1992] L. B. Smith and M. D. Sera. A developmental analysis of the polar structure of dimensions. *Cognitive Psychology*, 24(1):99-142, 1992.

[Smith, 1989] L. B. Smith. From global similarities to kinds of similarities: The construction of dimensions in development. In S. Vosniadou and A. Ortony, editors, *Similarity and Analogical Reasoning*, pages 146–178. New York: Cambridge University Press, 1989.

[Smolensky, 1988] P. Smolensky. On the proper treatment of connectionism. *Behavioral and Brain Sciences*, 11(1):1-74, 1988.

[Smolensky, 1991] P. Smolensky. Connectionism, constituency, and the language of thought. In B. Loewer and G. Rey, editors, *Meaning in Mind: Fodor and his critics*. Cambridge, MA: Blackwell, 1991.

[Thelen and Smith, 1994] E. Thelen and L. B. Smith. *A Dynamic Systems Approach to the Development of Cognition and Action*. Cambridge, MA: The MIT Press, 1994.

[Thelen, 1995] E. Thelen. Time-scale dynamics and the development of an embodied cognition. In R. F. Port and T. van Gelder, editors, *Mind as motion*, page 69-100. Cambridge, MA: The MIT Press, 1995.

[Tversky and Gati, 1982] A. Tversky and I. Gati. Similarity, separability and the triangle inequality. *Psychological Review*, 89(2):123–154, 1982.

[Tversky, 1977] A. Tversky. Features of similarity. *Psychological Review*, 84(4):327–352, 1977.

[Van Gelder and Port, 1995] T. Van Gelder and R. F. Port. It's about time: An overview of the dynamical approach to cognition. In R. F. Port and T. van Gelder, editors, *Mind as motion*. Cambridge, MA: The MIT Press, 1995.

[Wason and Johnson-Laird, 1972] P. C. Wason and P. N. Johnson-Laird. *Psychology of Reasoning Structure and Content*. London: Routledge, 1972.

[Wellman, 1990] H. M. Wellman. *The Child's Theory of Mind*. Cambridge, MA: The MIT Press, 1990.

MECHANISMS AND PSYCHOLOGICAL EXPLANATION

Cory Wright and William Bechtel

1 INTRODUCTION

What is it to explain a psychological phenomenon (e.g., a person remembering a name, navigating through campus, understanding humor)? In philosophy, a traditional answer is that to explain a phenomenon is to show it to be the expected result of prior circumstances given a scientific law. Influenced by this perspective, behaviorists directed psychology toward the search for the laws of learning that explained all behavior as the consequence of particular conditioning regimens. Although discussion of laws remains commonplace in philosophical accounts of psychological practice, appeal to laws in the explanation of psychological phenomena has become increasingly peripheral in psychology proper — especially with the rise of the cognitivist tradition. Examination of the explanatory discourse of psychologists reveals a shift in emphasis from laws to *mechanisms* — mechanisms of motivation and drug addition, mechanisms of motor development, auditory recognition mechanisms, etc. This raises a substantive philosophical issue: what is a mechanism, and how does discovering and specifying one figure in an explanation?

Although largely neglected in recent philosophical discourse about psychology, the search for mechanisms is one of the principal strategies for rendering the natural world intelligible through scientific investigation. It lay at the foundation of the scientific revolution of the 17^{th} and 18^{th} centuries, and was enshrined in the mechanical philosophies advanced by Galileo, Descartes, and Boyle, among others. The dialectic between mechanistic and anti-mechanistic thinking played a decisive role in framing issues for theorizing about mental phenomena in the 19^{th} century, and further shaped the genesis of psychology as a discipline in the 1880s and beyond. Indeed, the rise of cognitive psychology around the 1960s was guided by a particular, information-processing, conception of mechanism.

In the next two sections, we will describe the development of the mechanical philosophy and its applications to mental phenomena, and will then turn toward a more analytical characterization of mechanism and mechanistic explanation. Understanding what mechanisms are and recognizing the role they play in psychology casts a number of traditional philosophical issues about psychology in a very different light than in most philosophical discussions. We will discuss some of these in subsequent sections.

Handbook of the Philosophy of Science. Philosophy of Psychology and Cognitive Science
Volume editor: Paul Thagard
General editors: Dov M. Gabbay, Paul Thagard and John Woods
© 2007 Elsevier B.V. All rights reserved

2 THE RISE OF THE MECHANICAL PHILOSOPHY AND ITS APPLICATION TO THE MIND

2.1 Cartesian roots

René Descartes is a pivotal figure in the history of the sciences, and, arguably, his most influential contribution was the heralding of a mechanistic view of the natural world.[1] Whereas classical thinkers primarily viewed machines as devices operating *against nature* that satisfy human purposes (e.g., to lift heavy weights or launch projectiles in opposition to their natural downwards motion), Descartes proposed that natural systems were mechanical. He noticed mechanisms at work throughout the natural world — including the bodies and nervous systems of human and non-human animals (indeed, the human mind was virtually the only domain where he took them to be absent).

The mechanisms familiar to Descartes (e.g., clocks, which were undergoing rapid development the 17^{th} century), typically produced their effects because of the shape, motion, and contact between their parts. So, if natural systems are mechanical, then they could likewise be rendered explicable by appealing to the shape and motion of their parts: "I have described this earth and indeed the whole universe as if it were a machine: I have considered only the various shapes and movements of its parts" [1644, IV, §188]. Two examples of physical phenomena — gravity and magnetism — will illuminate Descartes' appeal to mechanics. Explaining either phenomenon depends critically on the assumption that the physical universe is comprised of contiguous bodies such that no empty space or vacuum exists. Wherever space seems to be empty, as in the heavens, Descartes assumed that it was filled with a very fine material: the ether. Descartes maintained that, when an object moves, something else (such as the ether or another object) must immediately move into the space vacated. To explain gravity, then, Descartes appealed to the vortex created by the rapidly circulating ether, which forced objects downwards towards the center of the earth. In a similar manner, he proposed that the vortex surrounding the Sun served to hold the planets in their orbits. To explain magnetism, Descartes again invoked the model of vortex action, but also implicated the motion of screw-threaded particles circulating around the magnet. These particles would screw themselves into corresponding threaded channels in a nearby metallic object such that the magnet and object would move together. Thus, it was the shape and motion of microscopic particles that he thought determined the behavior of macroscopic objects.

Descartes faced several challenges in developing such accounts of physical phenomena. In particular, the parts that he posited as constituting physical objects — namely, minute corpuscles — were too tiny to be seen by the unaided eye. Yet, he was undeterred by the fact that the properties of these particles therefore had to be inferred:

[1] Our discussion of Descartes' mechanical philosophy follows the analysis offered by Garber [2002].

> I do not recognize any difference between artifacts and natural bodies except that the operations of artifacts are for the most part performed by mechanisms which are large enough to be easily perceivable by the senses — as indeed must be the case if they are to be capable of being manufactured by human beings. The effects produced by nature, by contrast, almost always depend on structures which are so minute that they completely elude our senses [Descartes, 1644, Part IV, §203]

Descartes proposed to infer the properties of these corpuscles by a kind of reverse engineering, remarking that:

> Men who are experienced in dealing with machinery can take a particular machine whose function they know and, by looking at some of its parts, easily form a conjecture about the design of the other parts, which they cannot see. In the same way I have attempted to consider the observable effects and parts of natural bodies and track down the imperceptible causes and particles which produce them [Descartes, 1644, Part IV, §203]

Descartes proposed several mechanistic processes to explain biological phenomena as well. He was quite impressed with William Harvey's account of the circulation of blood, although he did not follow Harvey in construing the heart as a pump. Instead, Descartes proposed that the heart serves to heat the blood, thereby causing it to expand and dilate, so that the corpuscles of the blood can move out through the arteries until they cool in the capillaries and return to the heart through the veins. Taking the circulation of blood as his starting point, Descartes offered similar mechanistic accounts of the behavior of various organs of the body.

The idea that natural phenomena — including physiological processes — are the activities of mechanisms was already a radical departure from the traditions based on Aristotelian science, which endorsed teleological explanation. Yet, Descartes made a further, controversial move in developing his mechanical philosophy. He maintained that *all* behavior exhibited by animals was generated mechanically and so did not require positing purposes or goals. Of paramount inspiration for this additional move were his encounters with the hydraulically controlled statues in the Royal Gardens at St. Germain-en-Lai outside of Paris. The opening and closing of valves in the plumbing, which resulted from visitors stepping on critical tiles, caused these statues to move in anthropomorphic ways. Consequently, Descartes proposed that a very fine fluid, which he — following a tradition harking back to Galen — called 'animal spirits', likewise ran through the nerves in animal bodies, causing them to respond differentially to various sensory stimulations.

> In proportion as these [animal] spirits enter the cavities of the brain, they pass thence into the pores of its substance, and from these pores into the nerves; where, according as they enter, or even only tend to enter, more or less, into one than into another, they have the power of altering the figure of the muscles into which the nerves are inserted,

> and by this means of causing all the limbs to move. Thus, as you may have seen in the grottoes and the fountains in royal gardens, the force with which the water issues from its reservoir is sufficient to move various machines, and even to make them play instruments, or pronounce words according to the different disposition of the pipes which lead the water [Descartes, 1664, Part VI, §130]

Moreover, Descartes did not see any reason to distinguish human and non-human animals in this respect; any human behavior that was comparable to that of non-human animals was likewise the product of mechanisms operative in the physical body. Of course, humans do perform some activities which non-human animals do not; for some of these — such as the construction and comprehension of novel sentences — Descartes could not conceive of a mechanism, and so concluded that mechanistic explanation of all human activities was not possible:

> We can easily understand a machine's being constituted so that it can utter words, and even emit some responses to action on it of a corporeal kind, which brings about a change in its organs; for instance, if it is touched in a particular part it may ask what we wish to say to it; if in another part it may exclaim that it is being hurt, and so on. But it never happens that it arranges its speech in various ways, in order to reply appropriately to everything that may be said in its presence, as even the lowest type of man can do [Descartes, 1637, Part V]

A second activity for which he thought mechanistic explanation failed is the ability to reason with regard to any given topic. Although animals might exhibit intelligent behavior in particular domains, they would fail to behave intelligently in others. Descartes maintained that this particularity revealed that their apparently intelligent behavior was therefore not due to reason, but to a cleverly designed mechanism. He compared an animal's superior performance in a given domain to "a clock which is only composed of wheels and weights," yet "is able to tell the hours and measure the time more correctly than we can do with all our wisdom" [1637, Part V] Descartes seemed to be reasoning that reason, as found in humans, is a capacity with universal applicability to any subject, and that, if animals did act from reason, their greater capacity in one domain would result in greater capacity in all domains.

Thus, for Descartes, mechanisms lacked the flexibility needed to account for the variability and context-sensitivity of language use and reason. So, instead of explaining these activities in terms of mechanisms, he attributed them to an immaterial mind, i.e., a distinct substance construed in terms of the attribute of thinking: The mind is "a thing that thinks. What is that? A thing that doubts, understands, affirms, denies, is willing, is unwilling, and also imagines and has sensory perceptions" [Descartes, 1658, Meditation 2]. Hence, unlike material bodies, which were defined in terms of being extended, a Cartesian mind was not something that occupied, or could be located in, space.

Having made a sharp distinction between material bodies and the immaterial mind, Descartes faced the potentially embarrassing problem of explaining how the one could affect the other. Famously, he located the site of the mind's interaction with the body at the pineal gland. Since Descartes viewed the nerves as conduits for animal spirits, the mind had to affect the flow of animal spirits if it were to have any impact. For Descartes it was the pineal gland's central location that made it a good candidate for the locus of interaction, for there it could alter the flow of fluids through the ventricles of the brain through slight shifts in position.

2.2 Mechanizing thought

Descartes' substance dualism was an unstable feature of his philosophy. Some of his followers, such as Julian Offray de La Mettrie [1748], argued for extending the mechanistic view to the human mind. But far more common was an anti-mechanistic attitude toward mental processes. Even some who found the problem of interaction to pose sufficiently serious problems to undercut Cartesian dualism nonetheless rejected a mechanistic conception of mental processes.

This attitude is well-illustrated at the beginning of the 19^{th} century in Jean-Pierre-Marie Flourens' opposition to Franz Joseph Gall's proposal to distinguish a number of different mental functions and localize each in a different brain area based on cranial shape. Up to a point, Flourens supported the project of localizing functions in the brain; but he based his own inferences on experimental lesions (primarily in birds), rejecting the reliance on correlations and cranial measures that subsequently inspired much of the negative commentary on Gall's phrenology. Flourens made the important discovery that coordinated movement is controlled in the cerebellum, and also found support for Gall's overall claim that mental activity is localized in the cerebral hemispheres.[2] However, Flourens attacked Gall's key claim that mental activity should be divided into isolable parts, each seated in its own area of the brain. Deploying evidence from his own lesion experiments, Flourens concluded that cognitive capacities were not differentially localized in the cerebral cortex; rather, the cortex was a unitary organ. He cited his finding that, to the extent that a lesion compromised one mental capacity, (e.g., perception), so too to the same extent would other capacities be compromised. For Flourens, this pointed to a Cartesian view that the mechanistic program of decomposing a system into parts with distinctive functions ends at the mind. It is noteworthy that Flourens dedicated his *Examen de la phrenology* to Descartes: "I frequently quote Descartes: I even go further; for I dedicate my work to his memory. I am

[2]Identifying mental abilities with the brain was one of the few features of Gall's views about which Flourens had anything positive to say (though he emphasized that Gall could not claim propriety over the view): "the proposition that the brain is the exclusive seat of the soul is not a new proposition, and hence does not originate with Gall. It belonged to science before it appeared in his Doctrine. The merit of Gall, and it is by no means a slender merit, consists in his having understood better than any of his predecessors the whole of its importance, and in having devoted himself to its demonstration. It existed in science before Gall appeared — it may be said to reign there ever since his appearance" [Flourens, 1846, 27-8].

writing in opposition to a bad philosophy, while I am endeavouring to recall a sound one" [1846, xiv].

As the 19^{th} century progressed, an alternative tradition of localizing psychological processes in the brain took root. Yet, the guiding conception of mental activities was very different than Gall's phrenology, drawing inspiration not — as Gall did — from differences between individuals (as well as between species) in mental activities, but from the associationist tradition arising from John Locke. In many respects, the associationist tradition is, at root, a mechanistic tradition, since it construes thinking as involving an assembling or associating of ideas. Although most 17^{th} and 18^{th} century associationists, such as John Locke and David Hartley, rejected any attempt to link associationist psychology to the brain, an entrée for doing so was provided by Charles Bell's and François Magendie's lesion experiments on dogs in the 1820s, which culminated in the discovery that the posterior spinal nerves are sensory while the anterior nerves are motor. The fact that the sensory inputs arrive at the brain via a different pathway than that carrying the specification of motor activity suggested that the intervening brain constituted a mechanism for making associations. This suggestion was embraced by Alexander Bain, who made his objectives clear in a letter to John Stuart Mill in 1851:

> I have been closely engaged on my Psychology, ever since I came here. I have just finished rough drafting the first division of the synthetic half of the work, that, namely, which includes the Sensations, Appetites and Instincts. All through this portion I keep up a constant reference to the material structure of the parts concerned, it being my purpose to exhaust in this division the physiological basis of mental phenomena.... And although I neither can, nor at present desire to carry Anatomical explanation into the Intellect, I think at the state of the previous part of the subject will enable Intellect and Emotion to be treated to great advantage and in a manner altogether different from anything that has hitherto appeared. (quoted in [Young, 1970, 103])

Despite Bain's stated objective, his own publications failed to advance the links between the associationist tradition and sensory and motor processing in the brain.[3] His students — especially David Ferrier and John Hughlings Jackson — however, pursued precisely that connection by using weak electrical currents to probe the brain and by analyzing deficits resulting from brain injury. Jackson, for example, commented that,

> To Prof. Bain I owe much. From him I derived the notion that the anatomical substrata of words are motor (articulatory) processes.

[3] Although Bain did not pursue the physiological component, he clearly construed the mind as a mechanism: "The science of mind, properly so called, unfolds the mechanism of our common mental constitutions. Adverting but slightly in the first instance to the differences between one man and another, it endeavours to give a full account of the internal mechanism that we all possess alike — of the sensations and emotions, intellectual faculties and volitions, of which we are every one of us conscious" [Bain, 1861, 29].

> (This, I must mention, is a much more limited view than he takes.) This hypothesis has been of very great importance to me, not only specially because it gives the best anatomico-physiological explanation of the phenomena of Aphasia *when all varieties of this affection are taken into consideration,* but because it helped me very much in endeavouring to show that the 'organ of mind' contains processes representing movements, and that, therefore, there was nothing unreasonable in supposing that excessive discharge of convolutions should produce that clotted mass of movements which we call spasm. [Jackson, 1931, 167-68]

Ferrier and Jackson were not the first to develop a link between a type of mental process and the brain. In 1861, Paul Broca established that an area in the frontal cortex was involved in speech. Broca made this connection working with a patient — Monsieur Leborgne — who lost the capacity for articulate speech. (Leborgne is better known by his pseudonym in the research literature, 'Tan' — one of the few sounds he uttered.) After Leborgne's death, Broca conducted an autopsy, and even though the brain damage was by then massive, Broca argued that it began in the frontal area that came to bear his name. Like Gall, Broca approached mental capacities from a faculty perspective, but subsequent work on language deficits by Carl Wernicke [1874] instead adopted an associationist perspective. Wernicke construed the cortex as realizing associations between sensory and motor areas, with particular types of associations realized in their own distinctive brain regions. In Wernicke's model of reading, acoustic or visual images of words were connected to motor images that controlled either speech or manual action.

In the early 20^{th} century, even the degree of localization Wernicke endorsed was challenged by researchers who adopted a very holistic conception of associations. Such an anti-localizationist view is exemplified by Karl Lashley [1948; 1950], who argued, much in the spirit of Flourens, that beyond the primary sensory and motor projection areas, cortex was non-specific and acted in a holistic manner to implement associations between sensory and motor areas. He coined the term 'association cortex', which was in common use in neuroscience through the mid-20^{th} century. This changed in the 1960s–1980s, when researchers adopting a localization perspective were able to make dramatic progress in showing that extensive brain areas anterior to primary visual cortex were involved in visual processing. By correlating activity in different areas with different kinds of stimulus characteristics, they were able to identify each area's specialization. Thus, areas that Lashley had identified as general association areas gradually became identified with processing of specific types of visual information (see [Bechtel, 2001b; van Essen and Gallant, 1994]. We will return to a detailed example of how neural processes figure in mental activity at the end of the chapter.

Associationism was rooted in more than one field (epistemology and psychology) and also influenced more than one field. In addition to its influence on neuroscience, associationism contributed to the rise of behaviorism within psychology in the United States in the early 20^{th} century. Commitment to a positivistic

philosophy of science led behaviorists to be suspicious of any appeals to psychological processes occurring in the head that could not be objectively observed. In particular, John Watson [1913] and subsequent behaviorists were skeptical of the introspectionist proposals regarding mental processing advanced by Edward Titchner [1907] and his followers. In place of appeals to mental processes, behaviorists sought laws relating behavior to objectively observable variables — stimuli in S-R psychology, reinforcers in Burrhus Frederic Skinner's accounts of operant conditioning. Behaviorists construed organisms — including humans — as learning mechanisms, and the strictest of them limited their laws to the regularities in input-output relationships that could be observed when these mechanisms functioned. They did not deny that there were internal processes occurring in the head, but rather, denied that psychology could or needed to provide an account of these processes. This was, instead, a task for physiology. Philosopher Carl Hempel's [1958] theoretician's dilemma captures the essence of the behaviorist's motivation for discounting internal mental events in explaining behavior. According to the dilemma, even if there are intervening processes caused by external variables which then cause behavior, these could be discounted in any explanation. Laws adequate to account for any behavior could be stated purely in terms of the causal external variables.

By the mid-20^{th} century, Descartes' mechanical philosophy was far more influential than his dualism, but the working conception of mechanism was still quite impoverished. In particular, most conceptions of mechanisms focused on sequential operations. Within physiology, investigators working on mechanisms within living systems had come to recognize that the component processes were often organized in cycles, not linearly; and theorists such as Claude Bernard [1865] and Walter Cannon [1929] had begun to appreciate the significance of more complex modes of organization for physiological regulation. Except for the cybernetic movement, which flourished especially in the period 1945–1955 [Wiener, 1948], however, theorists continued to think primarily in terms of relatively simply, linearly organized machines of the sort Descartes envisaged. But a new kind of machine — the digital computer with a random access internal memory — was capable of more complex patterns of behavior. It quickly supplanted the Cartesian mechanisms that had occupied traditional thinking about psychological phenomena.

3 INFORMATION-PROCESSING MECHANISMS AND THEIR APPLICATION TO PSYCHOLOGY

Theoretical ideas that contributed to the development of the digital computer were in place well before the technology needed for its construction was available. For instance, in 1777, Lord Charles Stanhope invented a device, 'the Demonstrator', which was purportedly capable of solving traditional and numerical syllogisms and elementary problems of probability [Harley, 1879]. In 1805, Joseph-Marie Jacquard introduced the idea of removable punch cards to specify the pattern a loom would weave. Charles Babbage drew upon this invention in the 1840s in his design for an

analytical engine which was to be a steam-driven computational device [Morrison and Morrison, 1961]. Although he was not able to build the analytical engine, he did engage in fruitful collaboration with Lady Lovelace (Ada Augusta Byron), who worked out ideas for programming Babbage's machine. Stanhope's machine and Jacquard's mechanical application of instructions were improved upon by William Stanley Jevons's various logic machines, such as his 'Logical Abacus', which incorporated sundry keys, levers, pegs, and pulleys into a device that anticipated modern calculators [Jevons, 1870]. Other prototypes of logic machines by Henry G. Marquand [1885], Charles P. R. Macaulay, Annibale Pastore, and Charles Peirce [1887] soon followed.

Working at a theoretical level prior to the actual construction of electronic computing machines, Alan Turing abstractly characterized a device (an automaton that came to be known as a 'Turing machine') that would perform the same operations previously performed by humans whose occupation was to perform complex calculations by hand. In advancing his characterization of a computing device, Turing [1936]; see also [Post, 1936] drew upon procedures executed by these human computers. Thus, in a Turing machine, a finite state device is coupled to a potentially infinite memory in the form of a tape on which symbols (typically just 0 and 1) are written. (The tape provides a memory that a finite state device lacks.) The finite state device has a read head that can read the symbol on one square of the tape and then may replace it by writing a different symbol or may move left or right one square. Which operation it performs depends on which of a finite number of states the device is in and which symbol it reads from the tape. The direction to perform an action also specify a new state into which the machine will enter. Turing demonstrated that — theoretically — for each computable function there exits a Turing machine that can compute it. Moreover, he established that, by encoding the description of each specific Turing machine on a tape, it is possible to devise a universal Turing machine that can simulate any given Turing machine and so compute any computable function.

Construction of actual computing devices began during World War II, although the first to be completed — ENIAC (Electronic Numerical Integrator and Calculator) — was not operational until late autumn 1945, and officially dedicated on February 15, 1946. John von Neumann designed the basic architecture that still bears his name while ENIAC was under development, but his crucial idea of a stored program was not implemented until the construction of ENIAC's successor, EDVAC (Electronic Discrete Variable Computer). By the time EDVAC itself was fully operational in 1952, the first commercially produced computer, UNIVAC I (Universal Automatic Computer) had been delivered to the Census Bureau, and the computer revolution was underway.

Although primitive and slow by contemporary standards, in their day these earliest computers were impressive in their speed of computation. But what was more theoretically significant for psychology than their speed was that they could be viewed as symbol manipulation devices. (Arithmetic computation is just one form of symbol manipulation, and the characterization of a Turing machine as

writing and reading symbols from a tape reveals that it is fundamentally a symbol manipulator that might be employed in arithmetic computation.) This inspired researchers to ask whether the same devices might perform other cognitive activities that were generally taken to require thinking and intelligence. While much of the popular attention focused on attempts to create programs that could play well-defined games such as chess, pioneers such as Alan Newell and Herbert Simon [1972] were enticed by the idea of mechanizing human problem solving. Newell and Simon's approach paralleled Turing's original work insofar as they drew upon the procedures that they believed humans explicitly follow in solving problems. Accordingly, one of their methods was to collect protocols by asking subjects to continuously describe the steps in their reasoning while solving a problem, and then to devise programs that would employ similar procedures.

Newell and Simon construed themselves as making contributions to both computer science and psychology. Within computer science, the project they pioneered was designated 'artificial intelligence'. But within psychology their work represented just one strand in the development of the tradition in cognitive psychology that construed the mind as an information-processing mechanism. The characterization of the process of symbol manipulation as information processing stems from an independent intellectual thread derived from the mathematical theory of information. In the late 1930s, Claude Shannon — in a master's thesis entitled *A Symbolic Analysis of Relay and Switching Circuits* — employed Boolean operations to analyze and optimize digital circuits that would later be used in computers. He then took a position at Bell Laboratories, focusing on the transmission of information over channels (such as a phone line). In the course of that research, Shannon [1948; Shannon and Weaver, 1949] introduced the concept of a BIT (binary unit) as the basic unit of information, and characterized the information capacity of a channel in terms of the ability of a recipient at the end of a channel to differentiate the state at the source of the channel. A particularly influential consequence of this research for psychology was Shannon's analysis of redundancy in a signal, which resulted when a given item in a signal would constrain the possibilities for another item (as the sequence of letters in 'mailbo' in an English text constrains the next letter). This result provided a basis for George Miller, who conducted his dissertation research during the war on the capacity to jam speech signals, to demonstrate that certain messages were harder to jam than others. Miller and Selfridge [1950] further developed applications of information theory in a list-learning experiment, explaining that the more closely word lists resembled English sentences (i.e., the greater their redundancy), the more words a subject could remember.

As we have noted, until this time American psychology had been dominated for forty years by behaviorist learning theory, which rejected attempts to explain behavior in terms of internal mental processes proposed on the basis of introspection. Information theory and the development of the computer provided a basis for thinking about internal processes in a much more constrained manner than introspectionism had offered. Very precise models of internal processes could be

proposed and their predictions tested against observable behavioral data. This new movement within American psychology, known as information-processing theory, changed the landscape of psychology [Bechtel *et al.*, 1998].

Miller was one of most influential researchers to develop information-processing models of cognition. For a variety of activities, such as remembering distinct items for a short period, distinguishing phonemes from one another, and making absolute distinctions amongst items, Miller [1956] showed that significant changes in processing occurred when more than a few items (7 ± 2) were involved. In addition to using behavioral data to establish limits on cognitive processing mechanisms in this and earlier work, Miller also collaborated with Eugene Galanter and Karl Pribram [1960] to provide one of the first suggestions of the structure of an information-processing mechanism. Their goal was to develop a framework that accounted for mental activities such as the execution of a plan. One challenge in executing a plan is to know when it is appropriate to initiate a behavior and when to end it. Miller *et al.* proposed a basic cognitive mechanism they called a 'TOTE unit': Test-Operate-Test-Exit. The idea is that when a test operation indicates that the conditions for an operation are met, it is performed, and continues to be performed until the conditions for it are no longer satisfied. One of the particular powerful features of the proposed scheme was that TOTE units could be embedded within other TOTE units, as diagrammed in Figure 1.

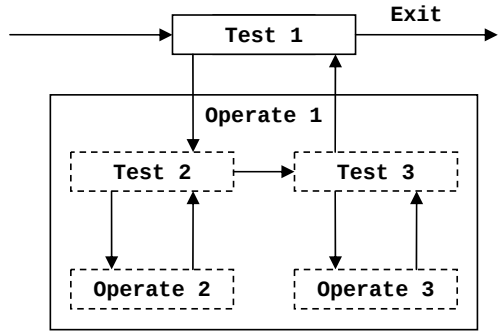

Figure 1. A TOTE unit in which two additional TOTE units are embedded. The upper level TOTE unit performs Test 1. If it is passed, the unit is exited; if not, Operation 1 is performed. This operation requires two additional tests, Test 2 and Test 3, each of which specifies an operation that is performed until that test is passed.

Artificial intelligence and psychology were not the only disciplines affected by the information-processing perspective. Linguistics was another. Starting with the efforts of Ferdinand de Saussure in the late 19^{th} century, structural linguists had developed useful proposals regarding the basic components of language. To

the extent they concerned themselves with how those components were combined, however, they found available mechanisms inadequate. In the 1950s, Noam Chomsky began to tackle this problem. He viewed a grammar for constructing possible sentences as a specialized automaton, and explored what kind would be necessary to generate all and only the well-formed sentences of a natural language such as English. One such automaton considered was a *finite state device*, in which the generation of a sentence would involve a sequential transition from state to state in the device. For example, the initial state might offer a choice among nouns with which to begin a sentence. Depending upon which noun was chosen, a specific set of choices would open up for the next word — perhaps a set of verbs. Transitions from state to state would continue until a complete sentence had been generated. Behaviorist accounts of language tended to be of this type, but Chomsky [1965] contended that finite state grammars cannot adequately characterize natural languages. An example of the problem for a finite state device is that a natural language allows a potentially unlimited number of embedded clauses to intervene between a noun and verb that need to agree in number. There are ways to build in agreement across clauses in a grammar that can be processed by a finite state device; but the more clauses, the more unwieldy becomes the grammar, and there is no way to get agreement across an indefinitely large number of such clauses.

Chomsky ultimately proposed that natural languages required grammars utilizing phrase-structure rules (which build tree structures) and transformational rules (which alter the tree structures). He also argued that a transformational grammar was equivalent in power to a Turing machine. Among the consequences of Chomsky's attempts to devise grammars and implement them in information-processing devices was that he enticed a number of psychologists to utilize his grammars in their analyses of the mental processes through which humans construct and comprehend sentences. One of these was Miller; already in his work with Galanter and Pribram he envisioned (but did not work out in detail) how a system of TOTE units could realize Chomsky's phrase structure and transformational grammars. Subsequently, he attempted to demonstrate the psychological reality of psychological transformations using reaction time data [Miller, 1962]. Chomsky's continued revisions of his grammars frustrated further attempts by psychologists to make such direct use of his grammars in understanding language processing [McCauley, 1987; Reber, 1987], although the efforts of other psychologists to reformat Chomskian grammars in more psychologically useful ways has spurred a successful lineage of research in psycholinguistics [Abrahamsen, 1987].

As noted above, information-processing psychologists primarily appeal to behavioral data to constrain their models of cognitive mechanisms. Error patterns and reaction time are two of the key behavioral measures invoked. Already in the mid-19^{th} century, Franciscus Cornelis Donders [1868] developed the idea of subtracting the time required to perform one task from the time required to perform a task requiring an additional operation to determine the time required for that additional operation. Saul Sternberg [1966] employed reaction time data to choose between candidate mechanisms for human memory retrieval. Measuring the time

subjects required to determine whether a given digit was on a just-memorized list, Sternberg found a linear relationship between the number of items on the list and how long it took to respond affirmatively or negatively to the test item. This ruled out a process of parallel access to the whole list in memory. More surprisingly, positive responses took as long as negative responses. If subjects performed a self-terminating search, stopping once they had found an item, positive responses should have taken less time. Since this was not the case, he posited a memory retrieval mechanism that incorporated exhaustive search in its design.

Most psychologists working within the information-processing approach to cognition believed that the information-processing mechanisms they were investigating were realized in the brain. However, they lacked tools for making the appropriate connections to brain processing. Although techniques for studying neurons using microelectrodes have been widely used in non-human animal studies since the 1940s (recording techniques) and even the 1870s (stimulation techniques), and have figured prominently in systems-level neuroscience research on such processes as visual perception during this same period [Bechtel, 2001b], ethical considerations limit the use of such techniques in humans. For human studies, neuropsychologists have obtained extensive behavioral data on individuals with naturally occurring lesions as a means of identifying the brain components involved in different cognitive mechanisms. However, until they began to collaborate with cognitive psychologists who employed the information-processing perspective, neuropsychologists were limited in their capacity to relate the brain regions they identified to actual mechanisms responsible for generating behavior [Feinberg and Farah, 2000]. Scalp recordings of electrical activity in the brain enabled researchers to trace some of the temporal features of information processing — especially when such recordings were time-synced to stimuli in order to measure evoked response potentials — but offered little ability to localize the brain processes spatially. Thus, it was not until the advent of functional neuroimaging with PET and fMRI that researchers interested in the mechanisms underlying psychological processes could link the component operations to brain processes [Posner and Raichle, 1994]. The introduction of neuroimaging coincided with the flowering of cognitive neuroscience as a research field involving the collaborative efforts of psychologists and neuroscientists in localizing information-processing mechanisms [Bechtel, 2001a].

The quest to explain mental activity mechanistically is well-established among practitioners of the biological, neural, and psychological sciences, who are now effectively integrating their investigations. Descartes' concerns about the inadequacy of mechanism to explain cognition have been assuaged, and the framework of information processing has provided a powerful vehicle for developing models of mechanisms responsible for cognitive behavior — one that is continuously exemplified in the statements of contemporary researchers:

> Nervous systems are information-processing machines, and in order to understand how they enable an organism to learn and remember, to see and problem-solve, to care for the young and recognize danger, it is essential to understand the machine itself, both at the level of the basic

> elements that make up the machine and at the level of organization of elements. [Churchland, 1986, 36]

With this overview of the historical route by which psychology became mechanistic, we now turn to a more analytical discussion of what a mechanism is and how a mechanical philosophy proffers a new perspective on some traditional issues in the philosophy of psychology.

4 CONTEMPORARY CONCEPTIONS OF MECHANISM

Some contemporary philosophers of science have regarded the increased sophistication of mechanistic approaches as concomitant with our best ways of understanding how reality is discovered and characterized; Wesley Salmon, for one, proclaimed, "The underlying causal mechanisms hold the key to our understanding the world" [1984, 260]. Yet, while empirical researchers frequently refer to and incessantly search for a massive array of particular mechanisms, philosophers have shown — until recently — little interest in what a mechanism is and the manner in which mechanisms might figure in explanations. Philosophers interested in the biological sciences — especially sciences such as physiology, cell and molecular biology, and neurobiology, where more traditional accounts of explanation that appeal to laws seem to yield little traction — have led the way in trying to articulate a satisfactory conception of mechanism and mechanistic explanation. William Bechtel and Robert Richardson [1993] offered one of the earliest explicit treatments, though their account was pitched in terms of *machines*.[4] They wrote,

> A machine is a composite of interrelated parts, each performing its own functions, that are combined in such a way that each contributes to producing a behavior of the system. A mechanistic explanation identifies these parts and their organization, showing how the behavior of the machine is a consequence of the parts and their organization. [1993, 17]; see also [Bechtel and Abrahamsen, 2005, 423]

Bechtel and Richardson developed and explored the consequences of this mechanistic approach by examining research in biochemical energetics, molecular genetics, and the cognitive neuroscience of memory.

In the decade since, variations on this conception have been advanced by several philosophers. Stuart Glennan, for example, has endorsed the above conception with the explicit qualification that it make room for the possibility and import of laws: "A mechanism underlying a behavior is a complex system which produces

[4]Bechtel and Richardson's focus on machines raises interesting questions pertaining to the minimal structure sufficient for something to be a machine — especially given that some entities, which are occasionally referred to as 'simple machines' (e.g., the wedge, the wheel), seem to engage in activities (as determined by their shape alone) but lack operative parts. We are inclined to distinguish machines — especially simple ones in the above sense — from mechanisms, at least of the sort that are of interest to biology and psychology.

that behavior by the interaction of a number of parts according to direct causal laws" [1996, 52]. The definition proposed by Peter Machamer, Lindley Darden, and Carl Craver — namely, that mechanisms are "entities and activities organized such that they are productive of regular changes from start or set-up conditions to finish or termination conditions" — emphasized the dual import of entities and activities, structure and function. They quip: "There are no activities without entities, and entities do not do anything without activities" [Machamer *et al.*, 2000, 3]. Additionally, Jim Woodward [2002, 374-75] characterized mechanisms as modular systems whose independent parts are subject to manipulation and control and behave according to counterfactual-supporting regularities that are invariant under interventions.

While there certainly are subtle differences between these various conceptions of mechanisms, their overall coherence reflects a growing consensus on the proper formulation. Indeed, the intersection of these and other similar conceptions is the view that many target phenomena and their associated regularities are the functioning of *composite hierarchical systems*. For example, such a system might encode nociception signals, recall an episodic memory, or plan alternate routes to a landmark. Systems may be active in isolation or in transaction with other things, many of which are themselves mechanisms, and they may be active at one time while inactive at other times.[5] As composite hierarchical systems, mechanisms are composed of *component parts* and their properties. Each component part performs some *operation* and interacts with other parts of the mechanism (often by acting on products of the operation of those parts or producing products that they will act on), such that the coordinated operations of parts is what constitutes or comprises the systemic activity of the mechanism. Stuart Kauffman [1971] proposed that component parts and their operations can be variably picked out according to the overall mechanistic activity to which they contribute causally, thereby making the characterization of the phenomenon critical to the identification of the mechanism.

Clearly though, not just any sequence of causal interactions will suffice, given that mechanisms are generally composed in such a way that they can endure for certain intervals. This raises many thorny metaphysical questions about how best to articulate the individuation and identity conditions for mechanisms — a non-trivial task for any mechanistic approach. For example, what are the necessary and sufficient conditions for a given sequence of causal interactions to be constitutive of a mechanism? How important is endurance, and when do gaps in temporal continuity cease being mere interruptions? To what extent does repair or particulate change in organization involve the creation of a new mechanism? Are there determinate limits on how much spatial or structural disconnectedness a mechanism can exhibit?

Articulating the individuation and identity conditions for mechanisms is an important task for future mechanistic approaches to tackle; in the erstwhile, we note

[5]Note that the inactivity of mechanisms does not render them explanatorily superfluous, as a mechanism's inactivity or inhibition may be just as important a factor in bringing about a given phenomenon and its associated regularities as another mechanism's activity.

that many of these questions involve another important feature of mechanisms — namely, how they are configured. The ability to endure and resist degradation importantly suggests that, in addition to the structural and temporal properties of the component parts and the functional properties of their operations, *organization* is critical to a mechanism. The relevant organizational and architectural properties — including location, orientation, polarity, cardinality and ordinality, *co*-operation, connection, feedback, frequency, duration, and so forth — enable the parts to work together effectively and perform the phenomenon Φ targeted for explanation, i.e., the presumable activity of a mechanism. The imposition of organization on components often produces more-or-less stable assemblies whose architecture then fixes the systemic activities that can be performed. Elucidating organizational properties is crucial for any mechanistic explanation, but it is especially important when the organization involves non-linearity, cyclic processes, etc. As the results of complexity theory have demonstrated, surprising behavior often results from such modes of organization.

A mechanism's spatiotemporal organization is also, in part, what makes a composite system *hierarchical* [Simon, 1969]. It has sometimes proven fruitful to conceive of organization in terms of *levels*, such that investigation and explanation of a mechanism's activity is understood as taking place at a higher level than an investigation or explanation of the constituency that composes it (see §6 below). At a higher level still, that very same mechanism may be a component part or subsystem in another, larger composite system. Its mechanistic activity would then constitute the operation of a component part).[6] Consequently, mechanisms — as composite, hierarchical systems — are multi-level; but just what these levels are has been a point of contention for mechanists (see §6.1 below).

5 MECHANISTIC EXPLANATION

5.1 Laws and the ontic/epistemic distinction

For much of the 20^{th} century, philosophers eschewed talk of both causation and mechanism, instead construing explanation nomologically in terms of laws. According to the traditional deductive-nomological (D-N) model, explanation takes the form of a deductive argument in which an event description is logically deduced from a set of statements of general laws in conjunction with a set of initial conditions [Hempel and Oppenheim, 1948]. Salmon [1984] was one of the first to dissent from the hegemony of law-based views of explanation.[7] Although Salmon

[6] Such claims need not commit mechanists to the view that mechanisms stand in mereological relations ad infinitum. Whether a given mechanism is a component of a higher-level mechanism depends upon whether it is part of an organized system at the higher level. Going the other direction, mechanists need not commit to where or whether mechanistic explanation 'bottoms out'. Whether a researcher pursues explanations that require ascending to higher levels vs. descending to lower ones, depends upon his or her explanatory goals.

[7] An earlier, but unfortunately less well-known, conception of mechanistic explanation is found in Harré's [1960] essay — some twenty years prior to Salmon's seminal work. Harré does not

spoke of 'causal/mechanical explanations', his principal focus was on causation rather than mechanisms of the sort just described. Salmon's characterization of the differences between nomological and causal/mechanical explanation continues to influence contemporary discussions, but also confuses them in an important respect.

Salmon characterized his causal/mechanical account of explanation as *ontic* and nomological accounts as *epistemic*. The motivation for this distinction is that nomological accounts have traditionally identified explanation with (deductive) *argumentation*. As such, the explanandum — a set of statements about the phenomenon Φ to be explained — is a logical consequence of the explanans — a set of statements about the relevant antecedent conditions whereupon Φ is produced together with a generalization (law) stating that, when those conditions prevail, Φ occurs. Salmon objected to rendering explanation in such terms: "An epistemic conception takes scientific explanations to be arguments..., but explanations are not the sorts of things that can be entirely explicated in semantical terms" [Salmon, 1984, 239, 273]. Instead, explaining a given explanandum "obviously involves the exhibition of causal mechanisms" leading to the occurrence of that explanandum [1984, 268]. While Salmon did note that explanations of an event take the form of fitting that event into a pattern of regularities described by causal laws, by "exhibiting it as occupying its place in the discernable patterns of the world" [1984, 17-18], he was quick to add that adequate explanations must track the mechanisms responsible for events.

> To provide an explanation of a particular event is to identify the cause and, in many cases at least, to exhibit the causal relation between this event and the event-to-be-explained.... Causal processes, causal interactions, and causal laws provide the mechanisms by which the world works; to understand *why* these things happen, we need to see *how* they are produced by these mechanisms. [1984, 121-24, 132]

So, it is because causal/mechanical explanations track mechanisms in the world that Salmon construed his account as ontic. Peter Railton [1978] articulated a similar shift in conception, suggesting that whatever lawlike generalizations range over regularities, they must be supplemented with information about the mechanisms producing those regularities:

> The goal of understanding the world is a theoretical goal, and if the world is a machine — a vast arrangement of nomic connections — then our theory ought to give us some insight into the structure and workings of the mechanism, above and beyond the capability of predicting and controlling its outcomes.... Knowing enough to subsume an event under the right kind of laws is not, therefore, tantamount to knowing the how or why of it. What is being urged is that D-N explanations

analytically define the concept of MECHANISM, but his preliminary framework partially anticipates the recasting of mechanistic explanation in terms of models as advocated in §5.2 below.

> making use of true, general, causal laws may legitimately be regarded as unsatisfactory unless we can back them up with an account of the mechanism(s) at work. [Railton, 1978, 208]

Machamer *et al.* took a more aggressive approach than either Railton or Salmon, upholding this emphasis on tracking mechanisms as the measure of explanatory adequacy while depreciating the import of lawlike generalizations with wide scope and global applicability.

> We should not be tempted to follow Hume and later logical empiricists into thinking that the intelligibility of activities (or mechanisms) is reducible to their regularity. Descriptions of mechanisms render the end stage intelligible by showing how it is produced by bottom out entities and activities. To explain is not merely to redescribe one regularity as a series of several. Rather, explanation involves revealing this productive relation. It is not the regularities that explain but the activities that sustain the regularities. There is no logical story to be told.... [Machamer *et al.*, 2000, 21–22]

The remarks of Railton, Salmon, and Machamer *et al.* point to typical motivations for this shift in conceptions in the following two respects. On one hand, the abandonment of epistemic conceptions is negatively motivated by sundry problems with the principles (e.g., subsumption of psychological phenomena under laws of nature, explanatory unification) and conceptions (e.g., D-N model, classical reduction) advanced by epistemic conceptions of explanation. Many of these problems are conceptual snags leftover from failed attempts to carry out various logical empiricist and positivist programs. The negative motivation for this abandonment can be understood, in part, as an attempt to distance mechanistic approaches from such programs. On the other hand, this abandonment is also positively motivated by the idea that an adequate conception of explanation should emphasize the mechanical production of local, individual phenomena.

Accordingly, on most ontic conceptions, an explanation counts as mechanistic for a phenomenon Φ only when it identifies a composite hierarchical system whose activities account for Φ and whose component parts are organized in certain ways and perform certain operations so as to be constitutive of the systemic activities — thereby situating Φ among a nexus of natural regularities [Salmon, 1984, 260, 268; Bechtel, 2002, 232; Bechtel and Richardson, 1993, 17-18; Glennan, 1996, 61; Railton, 1998, 752; Machamer *et al.*, 2000, 3, 22; Woodward, 2002, 373]. A common way of framing this conception is in terms of giving an answer to a how-question. Accordingly, Paul Thagard writes,

> [T]he primary explanations in biochemistry answer how-questions rather than why-questions. How questions ... are best answered by specifying one or more mechanisms understood as organized entities and activities. ... Thus answering a how-question is not a matter of assembling discrete arguments that can provide the answer to individual

> why-questions, but rather requires specification of a complex mechanism consisting of many parts and interconnections. [Thagard, 2003, 251]; see also [Cummins, 2000]

To grasp the importance of mechanists' appeal to the mechanisms themselves as explanation, consider two further examples that have figured prominently in the recent philosophical literature: electro-chemical synaptic transmission [Machamer *et al.*, 2000, 8-13; Bickle, 2003, 50-62] and stereoptic color and depth perception in state-space [Churchland and Sejnowski, 1992]. To explain the phenomenon of electro-chemical synaptic transmission, one *demonstrates* or *reveals* the operations and organization at the level of component parts, and the systemic activities performed at the level of the composite whole — e.g., the synthesis, transport, and vesicle storage of agonist/antagonist neurotransmitters and neuromodulators, their release and diffusion across the synaptic cleft, the process of binding with presynaptic autoreceptors and postsynaptic receptors, reuptake, depolarization, etc. Or again, the phenomenon of stereoptic color and depth perception is explained by demonstrating or revealing the internal structure, function, and organization of the component parts that both constitute the visual system, and comprise the production of those perceptual activities.[8]

5.2 Deconstructing and then recasting the distinction between epistemic and ontic conceptions of mechanistic explanation

The current literature on mechanistic explanation is filled with explications of the basic ontic conception wherein the explanans just is a mechanism or parts thereof and the explanandum is some phenomenon. When this view is taken literally, all references to non-ontic features (e.g., statements, inference, reference, truth-preservation) are omitted; the component parts and their operations and organization are themselves what do the explaining; and they do so in virtue of simply *being* the explanans. Accordingly, explananda are explained through or with the mechanisms that are responsible for them. Hence, in discussing long-term potentiation, Machamer *et al.* wrote, "It is through these activities of these entities that we understand how depolarization occurs" [2000, 13]. Similarly, in discussing biochemical pathways, Thagard wrote, "What explains are not regularities, but the activities that sustain regularities. Thus biochemical pathways explain by showing how changes within a cell take place as the result of the chemical activities of the molecules that constitute the cell" [2003, 238].[9]

[8] The interchangeable use of certain cognates (viz., 'demonstrate', 'reveal', 'lay bare', 'indicate', 'exhibit', 'display') seems particularly counterproductive; for they are typically left as semantic or conceptual primitives, and are not clarified, characterized, or given meaning over and above the already intuitive 'explain'.

[9] To be fair, Thagard [2003, 237, 251–52] calls for cognitive psychological research on mental representations of mechanisms, and advances 'a cognitive view of theories' whereby representations and operations thereon serve as a vehicle of mechanistic explanation. We certainly applaud the call, but go even further in jettisoning the reliance on the (naïve) ontic/epistemic distinction.

Do the component parts and their operations and organization figure in our understanding of how and why depolarization occurs? Well, yes, in a flat-footed sense: without any of these things to implicate, mechanistic explanations would be without content. But, in another sense, what our understanding *literally* proceeds 'through' is a network of linguistically- or graphically-expressed operations on representations. After all, scientists typically explain by *marshalling a narrative* — i.e., telling a story about why the explanandum is a consequence (material or otherwise) of antecedent conditions.

The problem with construing the ontic conception literally is well-illustrated in Scriven's example of the Yugoslav garage mechanic. In attempting to minimize any epistemic components of explanation, defenders of the ontic conception sometimes suggest that explanation consists in an *indicative act* — literally, a demonstrative gesturing or pointing at the component parts, operations and organization that together constitute the mechanistic activity that account for the explanandum. In critiquing Hempel's view that explanation requires argument, Scriven levels the following complaint: "Hempel's models could not accommodate the case in which one 'explains' by gestures to a Yugoslav garage mechanic what is wrong with a car" (quoted in [Salmon, 1984, 10]). But what work is a gesture or an indication doing? What the Yugoslav garage mechanic considers to be an explanation could only be what she or he finds cognitively salient about the situation. Indication, in Scriven's sense, would only be explanatorily helpful if it were made against the background of a large corpus of conceptual knowledge. And even if cognitive salience-conferring behaviors were sufficient for explanation, no such demonstrative gesturing or pageantry could alone unequivocally specify what had been indicated, since myriad structures, functions, and organizations would be consistent with any such gesture.[10]

One important consequence is that the view that mechanistic explanations are given by invoking mechanisms (qua explanans) must be understood as *metonymic* — i.e., as an emblematic stand-in for a richer, more accurate explication of what mechanistic explanation actually consists in. Presynaptic autoreceptors, sodium potassium pumps, and ligand-gated ion channels are simply inapposite candidates for any explanans; the relevant parts, operations, and organization minimally need to be captured and codified in a structural or functional *representation* of some sort. Accordingly, generating scientific understanding through mechanistic explanation necessarily involves inferences and reasoning about those representations.

[10]The same point can be made in the context of or Glennan's [1996] example of the mechanics of a toilet tank. Suppose that you are among the millions of people throughout the world who have had little-to-no experience with the recurrent flushing of a toilet and the regulation of water-level in its tank, and have no knowledge of indoor plumbing more generally. If someone attempts to explain to you how and why toilets are able to do those activities, then merely depressing the lever that initiates flushing, or taking off the lid and letting you see the internal goings-on of the tank, would be entirely deficient. Similarly, if the phenomena of salutatory conduction, memory consolidation, or creativity were cognitively abstruse, something more would be needed in addition to merely *presenting* or setting out the requisite mechanisms for observation — for even if all aspects of a mechanism are observable, there is no guarantee that one will have the "instant flash of insight" that accompanies self-explanation of how the phenomenon is produced.

(Of course, the operations on *re*presentations of mechanistic activity involved do not reduce to syntactic operations on propositions.) And so, unless the ontic conception is understood as metonymical, such accounts of mechanistic explanation will misconstrue explanatory practice by leading into absurd de re/de dicto confusions. After all, *explaining* refers to a *ratiocinative practice* governed by certain norms that cognizers engage in to make the world more intelligible; the non-cognizant world does not itself so engage. One way to appreciate this point is to recognize that mechanisms are active or inactive whether or not anyone appeals to them in an explanation. Their mere existence does not suffice for explanation; the systemic activity of a mechanism may be responsible for the presence of some psychological phenomenon Φ, but Φ is not explained until a cognizer contributes his or her explanatory labor.[11]

Now, there is clearly a contrast between the role of laws versus mechanisms in scientific understanding and explanation; but an ontic/epistemic distinction turns out to be an unfortunate way to draw that contrast because it invites trivializations and confusions. On one hand, characterizing explanation as ontic does not carry any additional commitments over and above what mechanists have already spelled out, so employing the distinction in the way Salmon intended merely reiterates the standard mechanist line on subverting the exclusive focus on laws and lawlike generalizations in the context of deductive arguments. Worse, taking ontic conceptions of mechanistic explanation at face value leads to absurdities, since mechanisms themselves are not the sorts of things that can be constituents of any explanans; and characterizing explanation as non-epistemic is clearly problematic insofar as explanation is through-and-through an epistemic practice of making the world more intelligible by providing a rational basis for inferring how some given psychological phenomena is or could be produced. On the other hand, it is entirely possible to give an account of explanation that is significantly epistemic without being overtly nomological (for related arguments, see Waskan, forthcoming). The upshot is that construing mechanistic explanation ontically and non-epistemically poorly captures what is distinctive of mechanistic explanations — i.e., the distinction is a false dichotomy.

Advocates of the ontic conception of mechanistic explanation cannot avoid appeal to epistemic conceptions; even when kicked out the front door, epistemic conceptions are quickly ushered in the back. One way in which this is manifest is in mechanists' vacillation between appealing to mechanisms themselves and identifying the explanans of mechanistic explanations with sets of descriptions of mechanisms. Hence, Machamer *et al.* aver that, "Giving a description of a mechanism for a phenomenon is to explain that phenomena and its production" [2000, 3]; see also [Craver, 2001, 68], while Glennan wrote, "A description of the

[11] There is yet a further reason why the mechanism itself is not the appropriate vehicle of explanation. Many mechanistic explanations that are offered turn out to be wrong, either fundamentally or in details. If the mechanism itself stood as a relatum in an explanation, erroneous explanation would not be possible since there would have been no mechanism for the erroneous explanation to have identified.

internal structure of the mechanism explains [its] behavior" [1996, 61]; see also [2002, 347].[12] This vacillation reflects recognition that a necessary condition on mechanistic explanation is that the structure, function, and organization of mechanisms needs to be captured and codified representationally. And this recognition shows that, while mechanists desire an account that pitches the explanation of psychological phenomena *purely* in terms of the mechanisms that produce them — thereby shedding the overbearing dependence on logical, semantic, grammatical, and inferential concerns — they cannot but simultaneously help themselves to certain resources of epistemic conceptions they abandon.

A thoroughly epistemic conception of mechanistic explanation can certainly maintain the contrast with nomological accounts by acknowledging the importance of representing mechanisms rather than laws. Such an epistemic conception can simultaneously recognize that semantic, logical, grammatical, and inferential concerns are not irrelevant while emphasizing that linguistic representations — much less those involving universal quantification — do not exhaust the range of possible representations. In many cases, graphical representations such as diagrams or figures are *much* more effective representations of mechanisms than linguistic descriptions. Because of their ability to represent objects in two or three dimensions, graphical representations are able to capture important elements of the spatial organization of a mechanism. Since one can also reserve a dimension for time, or use arrows to represent succession relations, graphical representations can also capture important aspects of the temporal organization of mechanisms. Whereas linguistic representations can only capture one component at a time, graphical representations can identify multiple components and their relations.

Just as mechanistic explanation extends representation beyond the linguistic, it also expands the way that representations are related to phenomena (see [Bechtel and Abrahamsen, 2005]). Instead of deductive argument, one must understand how the mechanism produces the target phenomenon. One strategy is to use imagination to put one's representation of the mechanism into motion so as to *visualize* how that phenomenon is generated. This is not an area in which there has yet been much work by philosophers.[13] Cognitive psychologists have begun to make some headway in characterizing the processes by which scientists and science students develop the ability to mentally simulate the behavior of simple mechanisms such as springs [Hegarty, 2002; Clement, 2003]. There is related work by philosophers [Nersessian, 1999; 2002] and psychologists [Ippolito and Tweney,

[12] Such remarks might seem to suggest that descriptions of mechanisms are not just coincident with, or derivative from, explanations — they *are* explanations. But explanations are not merely lists of descriptions of mechanisms or sets thereof; they include inferential and simulatory operations on them. (Considerations of the semantics of the explanatory connective 'because', as well as what it is that arrows in box-and-arrow diagrams represent, help in grasping this point; see [Talmy, 2000].)

[13] An exception is Waskan [forthcoming], who develops complementary arguments to those offered here against a purely ontic conception of explanation and advances an account in terms of what he calls intrinsic cognitive models that are more like scale models than lingusitic descriptions.

1995] on simulating experiments which may be useful to developing a more robust account of how cognitive agents come to understand the relation between the components of a mechanism and what it does. (The idea that what a person is doing in simulating an event points to a link with the activity of computer simulation in psychology: like mental simulations, computer simulations show that some phenomenon is what would result from the conditions specified in the simulation.)

Having resituated mechanistic explanation within an epistemic context, we need to consider briefly two traditional epistemic concepts that arise in talk of explanation — models and laws. The concept of MODEL is interspersed throughout the scientific literature, psychology included. Models stand in for the actual systems that researchers are trying to understand and are invoked in reasoning and theorizing about the actual system (see [Harré, 1960]). Beyond these commonalities, there are many disparate senses of 'model' and contexts in which they are used. In model theory, for example, a model is a set of entities, often abstract, that satisfy a set of axioms. Proponents of the structuralist and semantic views of theory [van Fraassen, 1989; Giere, 1988; 1999] construe a model as a set of abstract (non-physical) objects that conform to the theories advanced in a science. These then stand in (approximately) isomorphic relations to the actual objects in the world to which the theory supposedly applies. In the context of mechanistic explanation, the model's elements correspond to the parts of a mechanism, and their structure conforms, not to a theory, but rather to the mechanism's constituency and interactivity.

Mechanistic models may be abstract, or may be implemented physically. For example, an engineer may build a scale model to experiment with before building a device. Scientists may also build such models, but with the goal of understanding the mechanism being modeled. One of the best known examples is the physical model of DNA that James Watson and Francis Crick constructed in the course of discovering its double helix structure. A more subtle example is that a model of a cognitive activity, such as natural language understanding, may take the form of a computer program. The program itself is an abstract mechanistic model, but implementing it on a particular computer gives it a physical realization in which the consequences of its design are more readily discovered.

Yet another sense of 'model' arises from the fact that in the biological and behavioral sciences, researchers may select a particular model system (organism) on which to conduct research intended to apply to a much broader range of organisms or to humans. Mouse models of navigation and spatial memory are a prominent example in behavioral neuroscience. A model system brings out another important feature of models more generally — namely, that they typically simplify the target mechanism, abstracting from features of the system that are not taken to be essential to the generation of the phenomenon.

Lastly, reconsider laws. Mechanists often emphasize the contrast between nomological explanations that give pride of place to laws and mechanistic explanations. As such, the motivation underlying the rejection of an epistemic conception can

be understood not so much as an aversion to inferentially-based representation of the production of psychological phenomena as it is an aversion to doing so by deducing descriptions of phenomena from laws. Robert Cummins [2000, 118-22] nicely articulates one of the main reasons for rejecting the nomological conception of explanation: the explanation of psychological phenomena is not a matter of subsumption under law because psychological laws are simply 'effects', and effects are simply explananda — not explanans. He wrote,

> In psychology, such laws... are almost always conceived of, and even called, effects. We have the Garcia effect, the spacing effect, the McGurk effect, and many, many more. Each of these is a fairly well-confirmed law or regularity (or set of them). But no one thinks that the McGurk effect explains the data it subsumes. No one not in the grip of the D-N model would suppose that one could *explain* why someone hears a consonant like the speaking mouth appears to make by appeal to the McGurk effect. That just *is* the McGurk effect. [Cummins, 2000, 119]

Cummins correctly adduces that the practice of explanation by subsuming phenomena under laws is rare in psychology, and even when it is invoked, what is understood by the term 'law' tends to be a description of the target of (mechanistic) explanation.

It would, however, be a mistake to suggest that mechanists are simply opposed to the appeal to laws in explanations; on the contrary, they certainly include analyses of the significance of laws in their approaches, *where appropriate* (e.g., [Salmon, 1984; Glennan, 1996; 2002; Hardcastle, 1996]. Bechtel and Richardson wrote that, "The explanatory task begins and ends with models; we question the hegemony of laws in explanation, not their existence" [1993, 232]. Laws are sometimes needed to help characterize the regularities in the behavior of components of a mechanism and thus can play a supplementary role in mechanistic explanation. What does the major explanatory work is the modeling of the components and their operations as well as the manner in which they are organized. This work is not performed by identifying laws.

6 HIERARCHICAL MECHANISMS, LEVELS OF ORGANIZATION, AND REDUCTION

An issue that has garnered the excitement and consternation of generations of philosophers is whether theories in the psychological sciences reduce to theories in the neurosciences. One of the distinctive features of the mechanistic approach is that it demands a fundamental reorientation of this issue of reduction and reductionism. Accordingly, in one sense, a mechanistic explanation is through-and-through reductionistic: it's appeal to increasingly finer-grain component operations and parts in explaining the activity of a mechanism. But in another sense, a mechanistic explanation is non-reductionistic: explanations at a lower level do

not replace, sequester, or exclusively preside over the refinement of higher-level explanations, because mechanisms are hierarchical, multi-level structures that involve real and different functions being performed by the whole composite system and by its component parts [Wright, forthcoming]. Rather than serving to reduce one level to another, mechanisms *bridge* levels. So, while reductive and mechanistic approaches can be closely aligned, they diverge in important respects. Before developing this contrast further, we first need to clarify the basic concept of LEVEL.

6.1 Mechanistic levels of organization and analysis

Talk of 'levels' spans numerous disciplines — especially in the brain and neural sciences, cognitive science, and psychology — and it is hard to overstate the significance of the concept of LEVEL in these disciplines. Just the same, levels-talk is virtually threadbare from overuse [Craver, in press]. One reason is that the various conceptions of levels are rarely analyzed in any sustained, substantive detail despite there being a large litany of literature on the subject (for an attempt to rectify this problem, see [Wilson and Craver, this volume]. A second reason is that levels are ambiguously construed as both ontic levels of mechanistic organization and as epistemic levels of analysis [Bechtel, 1994].

A traditional conception suggests a neat hierarchical layering of entities into levels across phenomena. Here, scientific disciplines (viz., physics, chemistry, biology, psychology, sociology) are distinguished, in part, by the level of phenomena in nature that are their target of study. This conception is found in Oppenheim and Putnam's [1958] layer-cake account of disciplines, and Simon's [1969] account of metaphysically discernible levels of organization. Appealing to a variety of evolutionary considerations, Simon argued that nature would have to build complex systems all at once — an implausible conclusion — were it not for the assembly of stable, semi-autonomous, modular parts. The use of assemblies and subassemblies to facilitate increasingly organized, complex systems allows for component parts to be differentially deployed and combined. It also means that impairment of a subassembly is less likely to be disruptive to the overall system. Simon also offered an explanation of the differential stability of assemblies at successive levels of organization: the bonding energies used to create structures are greatest at the lowest levels (e.g., with the atom) and weaker at higher levels (e.g., covalent bonds in macromolecules).

This traditional construal of ontic levels of organization has been further developed by Wimsatt [1976], who also argued that the most frequent causal interactions are among entities of the same scalar magnitude. Yet, Wimsatt noted that any hope of finding neatly-delineated levels diminishes as one approaches entities the size of macroscopic objects, which interact across size scales; consequently, he introduced the concept of PERSPECTIVES for "intriguingly quasi-subjective (or at least observer, technique or technology-relative) cuts on the phenomena characteristic of a system, which needn't be bound to given levels" [Wimsatt, 1994]; see also [Wimsatt, 1974; 1976].

Several philosophers have resisted any traditional, overtly realist construal of levels. For instance, Craver [2001, 65-67] decries such accounts as typifying a problematic reification; in their stead, he recommends neutralizing some of the relevant metaphysical commitments by way of construing levels as *perspectival* in the sense of involving different views on an entity's activity in a hierarchically organized mechanism. The sense of perspectivalism involves using isolated, constitutive, and contextualized strategies for locating and differentiating components and their operations, as well as explaining how they work together to comprise mechanistic activity. The construal of levels as perspectival is also congruent with the idea that levels cannot be neatly delimited; for instance, Robert Wilson [2003] uses group and individual fitness in biological populations to propose that perspectival explanatory strategies naturally follow from the fact that levels of selection are actually fused or entwined, which precludes any determinate answers about causal efficacy or closure at a given level. Valerie Hardcastle [1996, 29-32] adopted a similar stance in arguing that the neatly divided, layer cake concept of LEVELS is nothing more than a theoretical imposition, since what counts as a level can be fairly arbitrary and relative to a variety of factors (e.g., types of questions asked, methodology); taking this point to its logical conclusion, John Heil [2002] has urged that talk of ontic levels of organization be dispensed with altogether.

Between these two extremes of realism and anti-realism about levels, an intermediate construal of levels has emerged. On a given cycle of decomposing a mechanism, this intermediate construal treats all relevant components as being at the same level [Craver and Bechtel, in press; Craver, in press; Bechtel, 2006]. Such decomposition is local to each mechanism and no attempt is made at identifying the level each component would occupy in a global portrayal of levels across all of nature. So, while this construal takes levels seriously — i.e., there are levels, and they are integrated in mechanisms — it is more restricted than the traditional conceptions advocated by Oppenheim and Putnam, Simon, and Wimsatt. For example, suppose a biological mechanism is being described in which sodium molecules cross a membrane. The traditional construal would have the membrane being at a higher level than the sodium molecules, since the membrane is itself composed of molecules. But in this more local construal of levels, the fact that molecules are constituent parts of the membrane need not entail that they be treated as existing at a lower ontic level than the membrane across which they are transported; relative to the explanatory perspective and purpose of explaining sodium transport, the membrane and the sodium molecules can be construed as existing at the same level, since each is a component part of the mechanism relevant to the activity under investigation. If an investigator pursues another cycle of decomposition, the components will themselves be analyzed into components at a lower level, but again this level is local to the analysis and bears the imprint of perspectival explanatory strategies. Multiple cycles of analysis thus give rise to a hierarchy of levels that is confined to a given mechanism, though the reality of the mechanism entails the reality of levels confined to it. An investigator who had started by focusing on a different activity may have ended up with a different pars-

ing of entities into levels (see [Kauffman, 1971]). In sum, levels on the mechanistic account are real in that they deal the particularities of actual components and their operations, but they are perspectival in that they are defined with respect to specific foci on mechanistic activities.

One of the main ways of staving off confusion has been to distinguish between the mechanistic conception of levels of components and the conception of epistemic levels of *analysis*. Perhaps the best known epistemic conception of levels in psychology and cognitive science is due to David Marr, who distinguished three "different levels at which an information device must be understood before one can be said to have understood it completely" [1982, 25]:

> 1. the abstract computational theory of the device, in which the performance of the device is characterized as a mapping from one kind of information to another, the abstract properties of this mapping are defined precisely, and its appropriateness and adequacy for the task at hand are demonstrated,
> 2. the choice of the representation of the input and output and the algorithm to be used to transform one into the other,
> 3. the details of how the algorithm and representation are realized physically — the detailed computer architecture, so to speak.

The emphasis for Marr was on the different epistemic projects an investigator can pursue and the kind of account required, not on mereological part-whole relations within a mechanism. Thus, the characterization of the algorithm (ii) and of its physical realization (iii) may be characterizations of the same thing, and thus not span ontic levels. One might be focusing on the same component parts, yet describing them in terms of their spatiotemporal properties rather than in terms of the algorithm they implement. The mechanistic account, unlike Marr's, is mereological — working down to a lower level involves decomposing something into its component parts and operations, rather than merely describing it differently.

One important consequence of the mechanistic, mereological account of levels is that it makes it much clearer how component parts operating within a mechanism can perform different functions than the composite system. Merely indicating that the properties of component parts and their operations at one level of organization are distinct from those of the overall mechanism and its activity, however, is insufficient to capture an important feature often attributed to higher levels — namely, that "composite wholes are greater than the sum of their component parts." Capturing this feature requires that one take seriously the notion of ORGANIZATION in play in phrases like 'levels of organization', yet that concept has not received sustained, detailed philosophical analysis. Perhaps the starting point par excellence comes from Wimsatt [1986; 1997], who has articulated several criteria over the years distinguishing between wholes that are *mere aggregates* of their parts, and wholes that are constituted by parts in some further way characteristic of organized systems. He suggests that in mere aggregates the parts (i) are intersubstitutable or (ii) can be reaggregated without altering the behavior, (iii) can

be added or subtracted with only qualitative changes in behavior, and (iv) exhibit no co-operative or prohibitory interactions. Composite wholes that do not satisfy one or more of these criteria possess organizational properties that give them a more complicated, systemic character.

One of the most basic types of departures from mere aggregativity arises when component parts interact *sequentially* so that at least one component performs its operations on the product of the operation of the previous components. It is often important for efficient operation that the products of one operation are immediately available to the entity performing the second operation. One way to insure this is to situate the component parts spatially and temporally adjacent to each other, fixing them in position. In human-engineered machines (e.g., cameras, cochlear implants), it is precisely the imposition of spatiotemporal order that renders parts into the sort of composite system identifiable as a machine. Biological mechanisms such as membranes often perform this function; they maintain the enzymes that catalyze a sequence of reactions in close proximity to one another in an organelle.

Particularly significant are deviations from aggregativity that result from going beyond sequential organization by allowing processes later in a sequence to feedback on those earlier in the sequence. Such feedback is often differentiated as negative or positive. In negative feedback, a product of a sequence of operations serves to inhibit one of the earlier operations. For example, the production of ATP from ADP in glycolysis and oxidative phosphorylation serves to inhibit earlier operations in these pathways, insuring that valuable foodstuffs are not metabolized until the ATP is utilized in energy demanding operations and must be regenerated. In positive feedback, a product of an operation might serve to increase the responsiveness of a component earlier in the process. Although positive feedback often results in runaway, uncontrolled behavior (e.g., when two reactions each create a catalyst that promotes the other reaction), in some cases it results in the self-organization of composite systems [Kauffman, 1993]. With positive or negative feedback, the causal interactions among component parts are a significant factor by which systems are able to exhibit the sort of integrity that allows them to form coordinated, stable subassemblies.

Organization is especially important for understanding some of the more salient characteristics of living organisms such as their ability to develop and maintain themselves over time. Not able to rely on external agents to construct or repair them, they need to perform these operations for themselves. This means that they must be organized so as to secure matter and energy from their environments and channel them into the construction and repair of their own bodies. Such systems exhibit what Ruiz-Mirazo and Moreno [2004] identify as *basic autonomy*: they maintain themselves as identifiable functioning systems by constructing and reconstructing themselves and managing the exchange of nutrients and waste products with their environment (for further discussion, see Bechtel, in press). Beyond those mechanisms required for basic autonomy, living organisms may evolve additional mechanisms. Each of these, however, must be constructed and maintained

through the organism's own processes, imposing serious organizational constraints on living systems and the mechanisms, including cognitive mechanisms, comprising them. While individual mechanisms of biological organisms may perform specialized tasks, they are necessarily highly integrated with each other, especially those responsible for fundamental energetic and generative processes.

6.2 Contrasts with philosophical accounts of reduction

Mechanistic explanations relate levels, but the relation proposed contrasts sharply with philosophical accounts of *intertheoretic reduction* that relate levels in terms of the reduction of pairwise theories. Generally, this approach characterizes each level as the locus of theories expressible as sets of axioms and postulates. The classical version of intertheoretic reduction [Oppenheim and Putnam, 1958; Nagel, 1961] held that a higher-level theory was reduced to a lower-level theory in virtue of being derived from it, together with a specification of boundary conditions and bridge principles. The boundary conditions restricted the conditions under which the higher-level theories would be applicable to specific situations, whereas the bridge principles equated vocabulary in the higher-level theory with that of the lower-level theory.

The strict derivation condition in the classical version proved difficult to satisfy, and — beginning with the work of Schaffner [1967] — a variety of post-classical accounts of intertheoretic reduction have been advanced [Hooker, 1981; Churchland, 1986; Kim, 1998; Bickle, 1998]. A key feature of these accounts is that they allow that lower-level theories might revise or refine (or, in certain circumstances, even entail replacement of or elimination of) the higher-level theories with which they are paired via a corrected image (we will overlook the nuances of these accounts for the purposes of this essay; (for an overview, see [McCauley this volume; Bechtel and Hamilton, in press]). These post-classical accounts have also been encumbered by certain difficulties [Endicott, 1998; 2001; Schouten and Looren de Jong, 1999; Richardson, 1999; Wright, 2000].

Our goal here is not the evaluation of either classical or post-classical accounts of reduction, but rather to examine the differences between such accounts and mechanist accounts of relations between levels. Perhaps the clearest difference is simply that mechanistic accounts do not start with separate theories at different levels that are subsequently related to each other with the formal apparatus of set theory or reconstructed in their idealized forms. Mechanistic explanations at each level are partial and constructed piecemeal with a focus toward actual experimental investigation, without overarching concerns that they be fit into grand, large-scale scientific theories; hence, there is no desideratum to provide a complete account of everything that happens. Further, the relation between different explanations at different levels results from the ability of a cognizer to simulate how the coordinated performance of the lower-level operations achieves the higher-level activity [Wright, forthcoming]. The result has the character of an interfield theory that identifies causal or mereological relations between phenomena described in different theories

[Darden and Maull, 1977] rather than a deductive relation between independent theories [Bechtel and Hamilton, in press].

Moreover, models of mechanisms have elements of both reduction and emergence. Mechanisms are inherently multi-level; the components and their operations occur and are investigated at one level, whereas the entire mechanism and its activity occur and are investigated at a higher level. Consequently, mechanistic explanation proceeds, in part, by modeling components of a system and how they are organized. At the level of component parts and operations, the explanation describes the operations internal to the mechanism that comprise and enable it to perform the overall activity. For example, in the case of visual processing, a lower-level account describes how cells in different processing areas extract specific information from earlier processing areas. Of course, an organism's visual processing mechanism responds to different visual arrays in order to provide information that guides higher mental activities or action. Procedurally, then, modeling components, their operations, and organization is only part of the overall explanation, since the conditions in its environment are also important parts of such explanation. A mechanism is, in one sense, an *emergent* structure that engages its environment, and explanations at the level of the whole mechanism must therefore characterize its engagement with other entities apart from itself.

Hooker [1981, 508-12] provides an instructive example (albeit one he advances in the context of providing an account of reduction). He imagines a network of electrical generators that individually exhibit fluctuations in the reliability $\langle \mathrm{r}_1, \ldots, \mathrm{r}_n \rangle$ of their net output, but, when regarded collectively, exhibit a far more stable reliability $f(\mathrm{r})$ of output. As a mechanistic system, the generators form a single, systemic-level generator — what Hooker calls a 'virtual governor' — that is able to do something which no generator on its own can. As such, the properties of the virtual governor just are those of the whole mechanistic system, and they are neither identical with the properties of any component generator, nor of sets of component generators. Each level of organization exhibits a unique integrity. The activity of the virtual governor that produces $f(\mathrm{r})$ occurs at a level of organization that is markedly different than the level of organization at which each component generator produces a specific reliability r_i.

The phenomenon Φ to be explained — virtual government — is an activity of the whole system. But it is an activity that results from the operations of the various component generators in the network. Consequently, it is important to "look down" a level — something which can be done recursively as one continues to descend to lower levels to explain how activities at a given level are comprised of component operations. Yet, the network does not meet Wimsatt's above criteria (i)–(iv) for being a mere aggregate, and so there is a sense in which the activities of government emerge from the operations and organization of components [Bechtel, 1995, 147–49]. As explanatory models of the causal mechanics involved circle back toward higher and higher levels in order to reconstruct the ultimate relation between explanans and explanandum, the significance of each component alone diminishes and one again focuses on the overall systemic activity — virtual

governance.

Accordingly, Wimsatt's [1986] reminder that issues of emergence at higher levels of organization, and of reduction at lower levels, are *inseparably entwined* is important to heed: "Aggregativity and failures of aggregativity give a clear sense in which a system property may be reducible to properties of its parts and their relations and still be spoken of as emergent" [Wimsatt, 1986, 259]; see also [Wimsatt, 1976, 206; Churchland and Sejnowski, 1992, 2-3]. In considering whether 'having a virtual governor reliability of $f(\mathrm{r})$' predicates a real property, and, if so, whether it is an irreducible one, Hooker similarly concludes that a real property *does* emerge from the structure of the system in an explanatory sense — one that cannot be reduced to the reliability outputs of any of the component generators. For his part though, Hooker waxes instrumentalistic on the question whether there is anything that exhibits virtual governance.[14] For mechanists who recognize entities at multiple levels of organization, though, a systemic-level governor itself is no less real than the component generators.

7 DISCOVERING MECHANISMS

Thus far, we have focused on conceptions of mechanism, but have only alluded to how mechanisms are discovered. Within the positivist tradition that viewed laws as the primary explanatory vehicle, philosophers following Reichenbach [1938] drew a sharp distinction between the context of discovery and the context of justification. Justification was viewed as an appropriate topic for philosophical analysis since it was construed as involving logical relations, while discovery was viewed as a non-philosophical topic relegated to psychology. The reason for this was relatively easy to appreciate: if laws are not just inductive generalizations of observations, the process of constructing new laws does not seem to be guided by principles that can be abstractly formulated. Although a number of philosophers have entered into the discussion of discovery in the last twenty-five years and have advanced constructive proposals [Nickles, 1980a; 1980b; Holland *et al.*, 1987; Darden, 1990; 1991], the task of analyzing discovery looks more manageable from the perspective of mechanism since the conception of a mechanism itself largely suggests what needs to be discovered — component parts, their operations, and the modes of organization.

[14] Hooker concludes that predication of virtual governorship is ontologically empty, extensionless. "There is no thing which is the virtual governor, so 'it' isn't anywhere, and even the property of being virtually governed cannot be localized more closely than the system as a whole" [1981, 509]. He extrapolates this conclusion about the mechanism of virtual governorship to the relationship between cognitive and neural states, suggesting that functionally-construed cognitive states recede or "disappear" at lower levels of analysis.

7.1 Explanatory strategies at multiple levels

The multi-level nature of mechanisms can be couched in terms of a trichotomy of explanatory strategies — contextual, isolated, and constitutive (e.g., [Craver 2001, 62-8]). Contextual strategies (+1 level) describe a mechanism performing operations as a component part in a higher-level, composite system, explain why that mechanism has developed or adapted, suggest how that mechanism is affected ecologically, and so forth. In many cases, contextual strategies are executed by assuming or holding constant some projected description of the organization and operations of component parts, and developing a model of the mechanism's systemic activity to test against its actual behavior [Bechtel and Richardson, 1993, 21; Hardcastle, 1996, 21-2]. Isolated strategies (0 level) identify the mechanistic activities that account for the presence of Φ without reference to ecological or evolutionary context, and without implicating lower-level structure and function. Laboratory studies, such as stimulating and recording spike trains from an isolated neuron or studying a human subject's responses to computer-generated stimuli are examples of this strategy. Setting up such studies requires making judgments about what environment conditions are and are not relevant for the activity and controlling them. Constitutive strategies (-1 level) describe the mechanism's component parts, their operations, and their organization, showing how the mechanism's constituency is responsible for its activity and so makes it more than a merely aggregative system. By themselves, constitutive strategies are a form of reductive explanation that involves "taking the mechanism apart."

Ideally, a complete understanding of a given mechanism's systemic activity would make use of each explanatory strategy in order to reflect the hierarchical nature of the composite system: situating the mechanism to get a handle on its relationship to other phenomena, identifying the specific variables that affect its ability to realize the phenomenon in question, and "looking down" a level to identify the lower-level organization of component parts and operations [Craver, 2002, 91]; for a similar perspective emphasizing the importance of adopting pluralistic perspective, not just a downward one, see [McCauley, 1996; McCauley and Bechtel, 2001]. While looking down is thus not the only explanatory strategy employed in developing an understanding of a mechanism, it is the one most distinctive of the mechanist endeavor and hence the reason why mechanistic philosophers and psychologists typically applaud work on reduction and reductive explanation, but take issue with reductionism ([Endicott, 2001, 378]).

7.2 Decomposition and localization

In explaining any mechanistically produced phenomenon, one adopts the fallible, explanatory heuristics (as opposed to algorithms) of localization and decomposition [Bechtel and Richardson, 1993; Bechtel, 1994; 1995; 2002]. *Localization* refers to mapping the component operations onto component parts. *Decomposition* refers to taking apart or *dis*integrating the mechanism into either component

parts (structural decomposition) or component operations (functional decomposition).

These two forms of decomposition typically require different experimental tools and techniques, and success in decomposing component parts and operations often occurs on different timescales in different sciences. In the brain and neural sciences, neuroanatomists such as Korbinian Brodmann [1909/1994] developed cytoarchitectural procedures for differentiating brain areas. Brodmann clearly anticipated that areas with different cytoarchitectures would perform different operations, but he had no tools for identifying these operations. When the technique of recording the electrical activity from single cells enabled researchers to determine what stimuli would drive cells in different brain regions, Brodmann's hope began to be realized. As applied to visual processing, for example, the approach allowed researchers to localize different steps in analysis of visual stimuli in different brain regions [van Essen and Gallant, 1994].

For most of its history, researchers in information-processing psychology had no access to what brain components were operative, but they developed powerful strategies for identifying the information-processing procedures that subjects were using. As we noted earlier, timing and accuracy of responses were the most common types of data to which these strategies were applied. The type of errors made in a task can often provide suggestions as to the ways subjects are performing the task. Daniel Kahneman and Amos Tversky, for example, used errors in a number of judgment tasks (e.g., do more English words have 'r' as a first letter or third letter?) to suggest the reasoning strategy subjects employed both when they produced the right answer and when they made errors [Kahneman *et al.*, 1982]. Likewise, perceptual illusions such as the Muller-Lyer illusion and the moon illusion offer clues as to the information-processing operations involved in perception. As these two examples illustrate, investigators cannot directly "read-off" the information-processing operations that give rise to false judgments: psychologists need to engage in a kind of reverse engineering to propose what kind of information processing could generate such an error. Moreover, often — as the perceptual illusions illustrate — the proposed mechanisms remain controversial long after the phenomena which prompted the search for the mechanisms are well-established.

Neuropsychological research involves yet another strategy for separating information-processing operations based on the deficits exhibited by subjects with neural damage. Until the development of neuroimaging techniques, neuropsychologists generally had little information about the locus of brain damage, but they developed a variety of tests to determine rather precisely the nature of the cognitive deficit in a patient. The discovery of patients exhibiting a distinct deficit (e.g., naming animals) provides a clue as to the specific cognitive operation that is compromised. If neuropsychologists can also find other patients who are normal in that capacity but exhibit a contrast deficit (e.g., in naming inanimate objects), the resulting *double dissociation* is taken as evidence that there are two separate operations performed in normal subjects.

Both cognitive psychology and neuropsychology provide techniques for decom-

posing cognitive function, but not for localizing these in brain regions. Starting in the 1980s with PET and in the 1990s with fMRI, neuroscientists employed tools for measuring blood flow in different brain regions, which cognitive neuroscientists treat as a proxy for neural activity. Since neural activity is always occurring in the brain, a strategy was needed to link it with cognitive operations. The strategy for doing so was adapted from Donders' original technique for determining reaction times for component mental operations: the blood flow recorded during performance of one task is subtracted from that recorded in a different task requiring additional cognitive processing (e.g., generating an appropriate verb rather than just reading a noun [Petersen *et al.*, 1989; 1988]. Like chronometric studies, these techniques require initial hypotheses about the component information-processing steps. But they also have the potential for prompting revisions in these beginning assumptions if the studies identify additional active areas. Such discovery prompts inquiry into what operations these areas are performing. Thus, neuroimaging serves both to localize functional operations onto structural components as well as to guide the search for additional functional operations.

7.3 Identifying modes of organization

The discovery of organization often lags behind the discovery of components and their operations. A common strategy when researchers start to take a mechanism apart is to identify a component within the mechanism as alone responsible for the systemic activity of the mechanism. This approach, which Bechtel and Richardson [1993] termed 'direct localization', is illustrated in both Broca's interpretation of the area to which he localized damage in Leborgne's brain as the locus of articulate speech, and in much of the first generation research in neuroimaging. Typically, however, such research results in discovering more components of the mechanism are involved in the activity, prompting researchers to investigate the operations they perform. Once multiple components are identified as being involved in the production of a given mechanistic activity, researchers attempt to relate these as operating serially. Bechtel and Richardson referred to such serial models as yielding a 'complex localization'. In many cases, however, mechanistic research gives rise to the discovery that the components are not just serially ordered, but figure in cycles and other more complex modes of organization comprising integrated systems.

When the organization being investigated remains relatively simple, it is possible to construct relatively simple explanatory models through mental simulation of the activity in the mechanism step-by-step. As more organizational complexity and *co*-operations are discovered, however, this becomes more difficult. With complex feedback loops, the mechanism can begin to behave in unexpected ways. To understand such behavior, researchers often need to offload some of the cognitive labor involved in constructing explanatory models of parallel complexity onto their (research) environment — e.g., by supplementing their own ability to mentally trace activity in a system with computer-based simulations. By supply-

ing models of the operations of the various components and the manner in which they interact with each other, researchers can discover some of the consequences of organization.

8 MECHANISTIC EXPLANATION IN PSYCHOLOGY: MOTIVATION AND REWARD

We finish our discussion by giving an example that will illustrate many of the issues discussed in earlier sections of this chapter: motivation. Surprisingly, this example has not received much discussion in philosophical accounts of psychology (a noteworthy exception might be [Schroeder, 2004]). One of the main reasons involves difficulties in imagining how the invocation of mechanistic activity could suffice to explain such an abstract — if even not simply abstruse — psychological phenomenon such as motivation. After all, motivation potentially implicates such murky issues as life-long aspirations, psychodynamic personality traits, diachronic desires for intangible psychosocial rewards, and so forth. Reticence about fully extending the mechanistic approach to psychology has been primarily motivated by the alleged inability to account for purposive, directional, or intentional behavior [Malcolm, 1968; von Eckardt, 2003].

Showing how mechanistic models could account for motivational and related states would therefore be a major boon in exorcising such reticence. In fact, significant scientific progress has been made in characterizing the mechanisms responsible for producing motivation, much of which is due to increasingly sophisticated mechanistic explanations of a psychological phenomenon initially characterized in folk idioms. These explanations are through-and-through mediated epistemically by models (an important feature of which, again, is the range of inferential operations on schematic representations of processes). Such mechanistic explanations of motivation are particularly well suited for use in and development of analyses of topics traditionally thought to fall solely within the province of philosophy (the nature of desire, weakness of will and compulsion, reasons explanations, etc.).

8.1 Early motivation research

In psychological theorizing at the turn of the 20^{th} century, motivation was generally understood in terms of a large network of "hidden forces" that subsisted below the level of conscious mental processes but manifested themselves in human conduct. In views as disparate as William James's [1890/1950] and Sigmund Freud's [1922], attempts to differentiate among the class of hidden forces focused on varieties of instincts — latent, unlearned biological causes which compel an organism to behave in particular ways in the presence of certain environmental stimuli. Whereas James proposed a taxonomy of twelve such instincts (e.g., imitation, fear, love, curiosity, cleanliness),[15] Freud confined himself to two overarching

[15]Some of James' instincts were cross-classified with the 18 'hidden forces' posited William McDougall [1908], though the latter drew heavily on Darwinian evolutionary theory and ideas

ones — the life and death instincts — but then characterized them in terms of concepts like MOTIVATION and PLEASURE and AROUSAL in his theory of the unconscious. What he called the 'Id' is simply a wellspring of urges administered by hedonistic principles.

Despite a new mandate for empirical study, early psychologists' positing of 'hidden forces', and the analysis of such forces into elaborate classification schemes of (innate) instincts, did little to illuminate the nature of the states, processes, and properties that "move" animals to behave; moreover, the paucity of consensus about both the nature and number of instincts made it clear that the taxonomies were more descriptive than explanatory. Thus, early psychologists' theorizing about motivation made few advances on centuries of philosophical rumination — from Bacon and Hobbes to Bentham, Mill, and Nietzsche.[16]

As is true of many psychological constructs, the concept of MOTIVATION entered into more recognizably scientific accounts as a variable that figured in explanations of the goal-directed behavior of both human and nonhuman animals in specific environments. In lieu of positing hidden forces, instincts, or other unobservable psychobiological states, behaviorists sought to discover the laws relating behavioral regularities to objectively observable variables. During the heyday of behaviorism, Clark Hull [1943] found it necessary to introduce intervening variables (e.g., drive, habit strength) into the laws he proposed to describe functional relations between reinforcement and patterns of behavior. But what do these variables stand for? Given the positivistic objectives of behaviorism, this was not a pressing question for Hull; it was sufficient to show how these intervening variables related to the selection, initiation, performance, and reinforcement of behavior.

Hull and other behaviorists bequeathed psychology with a concept of MOTIVATION — characterized as a dispositional variable inferred from reinforced behavior (a 'reinforcer' being defined as anything that will change the probability that immediately prior behaviors will recur) — which was well-suited for some types of experimentation. The advent of cognitivism initially only added to the clutter of motivation constructs, as post-positivist research focused on long-term planning, the influence of emotional factors, construction of prosocial skills, etc. John W. Atkinson's [1957] work, for instance, focused on cognizers' rational calculation of expectations of goal attainment and goal value; Bernard Weiner's [1986] attributional theory construed motivation as the outcome of causal ascriptions to achievement-related success or failure in the context of socially structured endeavors.

Consequently, many researchers were left wondering about the very utility of

about innateness.

[16] An important exception was Henry Murray, whose laboratory at Harvard combined an emphasis on personality factors with then-novel work on homeostatic mechanisms. Murray *et al.* [1938] catalogued dozens of viscerogenic and psychogenic "needs" and environmental "presses," and characterized them (achievement, power, affiliation, succorance, etc.) based on views of physiological regulation circulating at the time (e.g., Cannon, 1929) — namely, as electrochemical "energies" in the brain which, when activated by certain situations, create internal states of tension and insatiety that animate resolutory actions.

motivational constructs, the extent of disunification in the set of phenomena they pick out, and whether the concept of MOTIVATION is even necessary for the explanation of molar behavior. Of course, the term 'motivation' itself — derivative from the Latin verb 'movere' and the German 'Motivierung' — generally signifies that which guides movement or causes one to act; yet, due to a number of theoretical refinements and technological advances, the denotatum had clearly fragmented into a variety of related phenomena at the crossroads of affect, cognition, and action. Reflecting on its extremely wide berth, Judson Brown wrote, "The ubiquity of the concept of MOTIVATION, in one guise or another, is nevertheless surprising when we consider that its meaning is often scandalously vague" [1961]; see also [Wong, 2000, 1-2].

In a review of the motivation literature since the advent of psychology as a scientific discipline, Kleinginna and Kleinginna [1981] extracted a core group of common factors from over 100 wildly different conceptions, theories, and hypotheses. They found that numerous conceptions focused on the final common pathways of mechanisms producing goal-directed behavior. Based on this review, the authors made the following recommendation:

> It may be useful for psychologists to limit motivation, perhaps to the energizing mechanisms that are directly connected to the final common pathway for motor responses. This restriction would exclude both receptor influences and muscular/glandular reactions, as well as most analysis, storage, and retrieval mechanisms.... This view would allow for most of the energizing and some of the directing functions that psychologists traditionally have associated with motivation.
>
> These processes may not always be highly localized in the brain and may depend on cortical control as well as on the traditional subcortical motivation circuits such as the lower limbic system structures. By restricting motivation in this manner, we do not overlook the fact that psychological processes are complex and involve continuous interactions among various systems. [1981, 272]

8.2 From motivation to the mechanisms of brain reward function

Confining motivation research to better-delimited characterizations of the target phenomenon with more tractable constructs was one way in which psychologists were able to get a handle on motivation. This coincided with the move toward models of mechanistic activity in lieu of nebulous instinct-need-desire taxonomies and descriptions of lawful stimulus-response patterns — a move fostered by the serendipitous discovery of brain reward circuitry in the mid-1950s. James Olds and Peter Milner [1954] found that animals with electrodes inserted in certain areas of their brains will work extremely hard when their actions are paired with certain electric pulses. They hypothesized that the structure being stimulated — roughly, the medial forebrain bundle (MFB) — is responsible for directly mediating both

the hedonic effects of, and complex behavioral responses to, all pleasures and rewarding stimuli [Olds, 1965]. This represented an attempt to directly localize a psychological phenomenon in a brain region. Of course, since the initial research involved no decomposition of either component parts or their operations, it did not actually offer an account of a mechanism. It did, however, serve to restrict the focus of much research to the more tractable topic of reward. Moreover, localizing reward in the MFB provided a target neural structure that subsequent research could investigate. As we will see, this led to differentiation of component parts and operations involved in reward and an understanding of how they related to one another, thereby initiating a path of research culminating in a mechanistic explanation of reward.

Olds and Milner's original technique was developed into a reliable experimental test paradigm: intracranial self-stimulation [Bielajew and Harris, 1991]. As with other such test paradigms and animal models, the development of intracranial self-stimulation provided a quantitative measure for accessing a phenomenon previously thought to be impenetrably subjective; its value as a "good" animal model, however, lies not so much in being directly and comprehensively probative of pleasure and reward as in adding a new dimension for moving mechanistic research forward along with other neuropharmacological screening tests, behavioral bioassays, and simulations [Wright, 2002]. As lesion and neuroimaging techniques were eventually added to the experimental portfolio, researchers began to differentiate parts of the system, identify different operations performed by these components, and develop a schematic understanding of the workings of hedonic information processing.

A key result of this research was the identification of a cluster of highly-interconnected anatomical structures broadly constitutive of three interrelated systems: mesocortical, mesolimbic, and mesostriatal. A major component of these systems is a large number of dopamine neurons whose cell bodies are located in two proximal, evolutionarily primitive structures — the substantia nigra pars compacta (SN) and the ventral tegmentum (VTA). From these two nuclear groups, dopamine axons project through the MFB and basal ganglia to innervate a number of structures — including the ventral striatum (VS), ventral pallidum (VP), hippocampus (HC), extended amygdala (A), lateral hypothalamus (LH), and prefrontal cortex (PFC) — all of which are involved in the mediation of complex behavioral responses to reinforcing stimuli.

Together these two neuronal groups involve an interface between tightly coupled motor and reward mechanisms. Dopamine cell bodies in the SN, for instance, project throughout the mesostriatal system to the caudate (CN), putamen (P), and motor cortex (MC), and are well-known to produce an array of signals for sensorimotor control, locomotion, and the initiation of movement. Their degradation or destruction typically results in Parkinson's-like symptoms. VTA projections provide the basic structure for transmitting reward-related associative learning signals in the mesolimbic system.

With the identification of a host of interconnected brain areas apparently in-

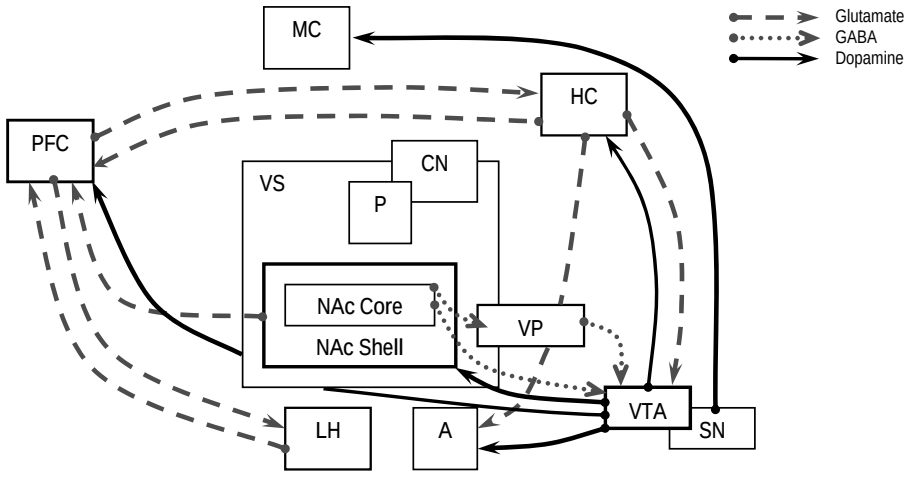

Figure 2. A schematic diagram of glutamatergic, GABAergic, and dopaminergic pathways involved in mesocorticolimbic circuitry, adapted from Berridge (2004, p. 204). Prefrontal cortex (PFC); Motor Cortex (MC); Hippocampus (HC); Ventral Striatum (VS); Putamen (P); Caudate Nucleus (CN); Nucleus Accumbens (NAc); Lateral Hypothalamus (LH); Amygdala (A); Ventral Pallidum (VP); Ventral Tegmental Area (VTA); Substantia Nigra (SN)

volved in reward, the next challenge was to figure out what each does — that is, to localize operations in the various areas. We will focus on one area that has been the locus of intense scrutiny and additional localization claims — nucleus accumbens (NAc), an area within the VS receiving major afferent dopaminergic projections from the VTA. This area has been further decomposed into a shell and a core, subcomponents which exhibit different neuronal organization and perform different operations. In particular, opioid receptors are numerously expressed at synapses in the NAc shell. These — in conjunction with mesolimbic dopamine activity originating from the VTA — provide key reward-related signals that regulate glutamatergic and GABAergic output to various cortical and striatal structures, respectively. These signals are often marked by the shift from tonic firing rates to phasic bursting.

Naturally rewarding stimuli are known to potentiate NAc shell transmission [Di Chiara, 1999; Berridge, 2003a]. But a wealth of convergent evidence about the contribution of these neurons to reward has also come from self-administration and self-stimulation studies with certain drugs of abuse whose reinforcement profile closely corresponds to those of many natural reinforcers. Generally, chronic drug use compromises the integrity of brain reward function and organization, resulting in long-term structural, functional, and organizational neuroadaptive changes in

the mesolimbic system [Franken, 2003; Koob and Le Moal, 1997; 2001]. For instance, cocaine addiction reduces the expression of endogenous κ-opioid receptors, which provide a natural inhibitory system for keeping tonic dopamine levels in the NAc in check [Chefer *et al.*, 2005]. The mesolimbic system adapts its activities to counteract the physiological changes introduced by drugs by shifting from homeostatic to allostatic mechanisms: given that the system requires continuous feedback and evaluation of its activities in response to fluctuating environmental stimuli, new set points for brain reward thresholds are constantly being generated in response to increasingly excessive environmental demands on the internal milieu of the mechanism [Koob and Le Moal, 2001]. When consumption is drastically reduced or terminated after a period of chronic administration, the drug user exhibits a range of aversive withdrawal symptoms and anhedonic deficits, given that the mechanisms for regulating normal brain reward function have been "braked" in order to accommodate the increase in the potentiation of monoaminergic neurotransmission.

Recent studies have decomposed NAc shell structure and function yet further. For instance, Ikemoto *et al.* [2005] demonstrated that the NAc shell supports heterogenous operations, as mice differentially self-administer δ-amphetamine in its ventral and medial parts; and Taha and Fields [2005] identified two distinct NAc shell populations with different encoding properties — inhibitory firing immediately before and during the initiation and maintenance of reinforcer consumption, and excitatory firing to encode palatability preferences.

The prominent role of dopamine in the mesolimbic system inspired a series of now-infamous dopamine hypotheses of reward [Wise, 2004]. These hypotheses suggest that dopaminergic transmission is the primary mediator of reward and reinforcement, and is a crucial operation in many reward-related homeostatic and allostatic mechanisms. In tandem, these hypotheses have also been used to advance claims about the pleasurable hedonic affect associated with rewarding and reinforcing stimuli: "Dopamine has often been called the 'brain's pleasure neurotransmitter', and activation of dopamine projections to accumbens and related structures has been viewed by many researchers as the neural 'common currency' for reward" [Berridge, 2003a, 32-3]. However, from increasingly sophisticated localizations and decompositions, it has become clear that this general picture of dopamine stands in need of revision. Although dopamine plays a crucial role, it does not seem to operate in isolation. Berridge and Robinson [1998] report that dopamine-selective neurotoxins such as 6-hydroxydopamine hydrobromine (OHDA-6) knock out virtually all NAc and LH dopamine operations, but not the capacity of rats to spontaneously engage in pleasure-induced behaviors. The result of these lesion studies is to functionally decompose reward-related mesolimbic dopamine operations into those governing certain 'liking' properties of hedonic affect, and those governing 'wanting' properties of incentive salience which can be dissociated from one another. This dissociation is also consistent with evidence distinguishing dopamine's role in appetitive versus consummatory behavior [Robbins and Everitt, 1996, 233].

8.3 From reward back to the complexities of motivation

Research into the production of reward-related phenomena by brain reward circuitry exemplifies how researchers were able to get a handle on a seemingly unwieldly, more complex phenomena such as motivation. By narrowing the context of research to modulation of NAc shell processes by VTA dopamine neurons, deemphasizing the utility of laws in explanation, using decomposition and localization strategies to identify parts and operations at successively lower levels, etc., researchers were able to make inroads into this perplexed area of research. Much of this progress was also born out of better tools and techniques for investigating the structure and function of hierarchical systems (mesolimbic, mesostriatal, mesocortical, etc.).

Yet the modulatory role of dopamine is much more complicated than simply providing a catch-all reward signal. Increasingly finer-grained mechanistic explanations will eventually specify just how complicated the picture actually is. But even then, dopamine is only a small piece of the puzzle: a comprehensive and complete mechanistic story of dopamine would not thereby yield a complete mechanistic explanation of reward, much less motivation. The main reason is simply that dopamine transmission is only causally efficacious in the context of larger subsystems, systems, mechanisms, and circuits working together; for example, the mesostriatal, mesolimbic, and mesocortical systems are organized in particular ways such that they comprise an extremely complex mesocorticolimbic system, in which many other factors besides dopamine turn out to be extremely important (e.g., neuropeptides, VP AMPA-Kainate expression, frontal lobe integrity). (In that sense, motivation mechanisms might share some similarities to Hooker's 'virtual governor', insofar as different mechanisms come together to produce and regulate certain phenomena that would be otherwise uncontrollable.)

Only in the interaction of these mechanistic systems does one begin to "see the forest for the trees" — a point continually emphasized across the various reflections of many of the main scientific players. For instance, in their review, Robbins and Everitt [1996] wrote, "Even leaving aside the complications of the subjective aspects of motivation and reward, it is probable that further advances in characterizing the neural mechanisms underlying these processes will depend on a better understanding of the psychological basis of goal-directed or instrumental behavior" [1996, 228]. Berridge and Robinson concur: "[F]urther advances will require equal sophistication in parsing reward into its specific psychological components" [2003, 507]. And in his review of concepts of MOTIVATION, Berridge concludes that higher-level motivation research is necessary to make sense of how the systemic interactions of neuroanatomical structures and neurochemical signals, mechanisms of protein folding, monoamine production, gene transcription, etc. realize psychological phenomena and produce behavior [2004, 205]. Others are explicit about the immediate need for systems-level functional neuroimaging results, which can place the specific components into the broader context of overall brain function [Robbins and Everitt, 1996, 233]. Consequently, although decomposition and localization

are crucial constitutive explanatory strategies, and are continuously applied in the reduction of composite systems into component parts and operations, ascending across levels is equally important as descent.

Hence, when Berridge [2003a; 2003b] and others explain motivational states in terms of attributions of the 'wanting' component of reward-related processes that transform perceptual representations into desired incentives for action, and in a way that is independent of hedonic valence, their explanations implicitly invoke models of mechanisms that exhibit much greater organizational complexity. Models of phasic bursting mechanisms of mesolimbic dopamine would fail to be pitched at the *appropriate* psychological level (i.e., in terms of transformation of representations, personal versus subpersonal systems). And indeed, a full explanation of motivation itself — especially beyond that of immediate attributions of incentive salience — must eventually involve models of mechanistic systems governing the production of planning and decision-making, the regulation of emotion and long-term memory, creativity, social role formation, and so forth [Franken, 2003; Ikemoto and Panksepp, 1999]. In sum, explaining motivation mechanistically requires illuminating the organizational collusion and interaction of these various composite systems that engage their environment at increasingly higher levels.

9 SUMMARY

The foregoing sketch of research on motivation and reward exemplifies what is now common explanatory practice in psychology — decomposing a composite, hierarchically organized system into its component parts and operations and then constructing models that abet scientific understanding of how they might be organized so as to comprise the mechanism's activity. Rather than its subsumption under sets of laws, such models feature in narratives about how a mechanism might be directly responsible for the phenomenon. As we noted at the beginning of the chapter, this explanatory practice entered science with the contributions of investigators such as Galileo, Descartes, and Boyle, and gradually became more common. Its broad acceptance in psychology was ushered in by the development of the information-processing perspective, which suggested that new understanding of how complex mechanisms work could help explain important features of psychological phenomena.

After noting this confluence between mechanistic approaches and information-processing perspectives in psychology, we turned to the task of explicating what mechanisms and mechanistic explanations are, drawing upon research from philosophers primarily focused on biology. We argued that explanation is inherently an epistemic or cognitive activity; so rather than misconstruing mechanistic explanation as ontic and nomological explanation as epistemic, what is needed is the development of the appropriate epistemic account of mechanistic explanation — one that focuses on how investigators reason with the models and representations of mechanisms. Such representation and reasoning often involves graphical representations or simulations of operations, not just linguistic representations of laws and initial conditions and deductive inferences.

Mechanistic approaches also reconfigure a number of issues in the philosophy of psychology beyond that of explanation. We have considered two: the question of reductionism, and the question of scientific discoveries. Mechanistic explanation is partially reductionistic, in the sense that it appeals to lower-level parts and their operations in explaining why a mechanism behaves as it does; but mechanistic explanation is not reductionistic in the sense of deriving higher-level theories from lower-level ones, nor in the sense of supplanting explanations of causal processes at higher levels, where the mechanism as a whole engages other entities in its environment. Causal processes at each level are different, and the ultimate result of a mechanistic account is an interfield theory that bridges levels. As to the question of scientific discoveries, mechanistic approaches are particularly apt for analyzing them, despite a tradition in philosophy of science that limits philosophy to characterizing justification of already discovered laws and disavows any prospect of contributing to the understanding of discovery. In particular, philosophers are engaged in articulating heuristics such as decomposition and localization, identifying what different investigatory techniques contribute to discovering components and operations, and understanding how scientists have discovered different modes of organization found in mechanisms, characterized their significance, and articulated relations between phenomena at different levels of organization.

BIBLIOGRAPHY

[Abrahamsen, 1987] A. A. Abrahamsen. Bridging boundaries versus breaking boundaries: Psycholinguistics in perspective. *Synthese, 72*, 355-388, 1987.

[Atkinson, 1957] J. W. Atkinson. Motivational determinants of risk-taking behavior. *Psychological Review, 64*, 359-372, 1957.

[Bain, 1861] A. Bain. *On the study of character, including an estimate of phrenology.* London: Parker, 1861.

[Bechtel, 1994] W. Bechtel. Levels of description and explanation in cognitive science. *Minds and Machines, 4*, 1-25, 1994.

[Bechtel, 1995] W. Bechtel. Biological and social constraints on cognitive processes: The need for dynamical interactions between levels of inquiry. *Canadian Journal of Philosophy, 20*, S133-S164, 1995.

[Bechtel, 2001a] W. Bechtel. Cognitive neuroscience: Relating neural mechanisms and cognition. In P. Machamer, P. McLaughlin, & R. Grush (Eds.), *Theory and method in the neurosciences*, pp. 81–111. Pittsburgh: University of Pittsburgh Press, 2001.

[Bechtel, 2001b] W. Bechtel. Decomposing and localizing vision: An exemplar for cognitive neuroscience. In W. Bechtel, P. Mandik, J. Mundale, & R. S. Stufflebeam (Eds.), *Philosophy and the neurosciences: A reader* (pp. 225-249). Oxford: Basil Blackwell, 2001.

[Bechtel, 2002] W. Bechtel. Decomposing the mind-brain: A long-term pursuit. *Brain and Mind, 3*, 229-242, 2002.

[Bechtel, 2006] W. Bechtel. *Discovering cell mechanisms: The creation of modern cell biology*. Cambridge: Cambridge University Press, 2006.

[Bechtel, in press] W. Bechtel. Biological mechanisms: Organized to maintain autonomy. In F. Boogerd *et al.*, eds *Systems Biology: Philosophical Foundations*, New York: Elsevier, in press.

[Bechtel and Abrahamsen, 2005] W. Bechtel and A. Abrahamsen. Explanation: A mechanist alternative. *Studies in History and Philosophy of Biological and Biomedical Sciences, 36*, 421-441m, 2005.

[Bechtel *et al.*, 1998] W. Bechtel, A. Abrahamsen, and G. Graham. The life of cognitive science. In W. Bechtel & G. Graham (Eds.), *A companion to cognitive science* (pp. 1-104). Oxford: Basil Blackwell, 1998.

[Bechtel and Hamilton, in press] W. Bechtel and A. Hamilton. Reduction, integration, and the unity of science: Natural, behavioral, and social sciences and the humanities. In T. Kuipers (Ed.), *Philosophy of science: Focal issues*. New York: Elsevier, in press.

[Bechtel and Richardson, 1993] W. Bechtel and R. C. Richardson. *Discovering complexity: Decomposition and localization as strategies in scientific research*. Princeton, NJ: Princeton University Press, 1993.

[Bernard, 1865] C. Bernard. *An introduction to the study of experimental medicine*. New York: Dover, 1865.

[Berridge, 2003a] K. C. Berridge. Comparing the emotional brain of humans and other animals. In R. J. Davidson, K. Scherer, & H. H. Goldsmith (Eds.), *Handbook of Affective Sciences* (pp. 25-51). New York: Oxford University Press, 2003.

[Berridge, 2003b] K. C. Berridge. Pleasures of the brain. *Brain and Cognition, 52*, 106-128, 2003.

[Berridge, 2004] K. C. Berridge. Motivation concepts in behavioral neuroscience. *Physiology and Behavior, 81*, 179-209, 2004.

[Berridge and Robinson, 1998] K. C. Berridge and T. E. Robinson. What is the role of dopamine in reward: Hedonic impact, reward learning, or incentive salience? *Brain Research Reviews, 28*, 309-369, 1998.

[Bickle, 1998] J. Bickle. *Psychoneural reduction: The new wave*. Cambridge, MA: MIT Press, 1998.

[Bickle, 2003] J. Bickle. *Philosophy and neuroscience: A ruthlessly reductive account*. Dordrecht: Kluwer, 2003.

[Bielajew and Harris, 1991] C. H. Bielajew and T. Harris. Self-stimulation: A rewarding decade. *Journal of Psychiatry and Neuroscience, 16*, 109-114, 1991.

[Broca, 1861] P. Broca. Remarques sur le siége de la faculté du langage articulé, suivies d'une observation d'aphemie (perte de la parole). *Bulletin de la Société Anatomique, 6*, 343-357, 1861.

[Brodmann, 1909/1994] K. Brodmann. *Vergleichende Lokalisationslehre der Grosshirnrinde* (L. J. Garvey, Trans.). Leipzig: J. A. Barth, 1909/1994.

[Brown, 1961] J. S. Brown. *The motivation of behavior*. New York: McGraw Hill, 1961.

[Cannon, 1929] W. B. Cannon. Organization of physiological homeostasis. *Physiological Reviews, 9*, 399-431, 1919.

[Chefer *et al.*, 2005] V. I. Chefer, T. Czyzyk, E. A. Bolan, J. Moron, J. E. Pintar and T. S. Shippenberg. Endogenous kappa-opioid receptor systems regulate mesoaccumbal dopamine dynamics and vulnerability to cocaine. *Journal of Neuroscience, 25*, 5029-5037, 2005.

[Chomsky, 1965] N. Chomsky. *Aspects of a theory of syntax*. Cambridge, MA: MIT Press, 1965.

[Churchland, 1986] P. S. Churchland. *Neurophilosophy: Toward a Unified Science of the Mind-Brain*. Cambridge, MA: MIT Press, 1986.

[Churchland and Sejnowski, 1992] P. S. Churchland and T. J. Sejnowski. *The computational brain*. Cambridge, MA: MIT Press, 1992.

[Clement, 2003] J. J. Clement. Imagistic simulation in scientific model construction, *Proceedings of the Twenty-Fifth Annual Conference of the Cognitive Science Society*, p. 25. Mahwah, NJ: Erlbaum, 2003.

[Craver, 2001] C. Craver. Role, mechanisms, and hierarchy. *Philosophy of Science, 68*, 53-74, 2001.

[Craver, in press] C. Craver. *Explaining the Brain: What a Science of the Mind-brain Could Be*. New York: Oxford University Press, in press.

[Craver and Bechtel, in press] C. Craver and W. Bechtel. Explaining top-down causation (away) *Biology and Philosophy*, in press.

[Cummins, 2000] R. Cummins. "How does it work?" versus "what are the laws?": Two conceptions of psychological explanation. In F. Keil & R. Wilson (Eds.), *Explanation and cognition* (pp. 117-144). Cambridge, MA: MIT Press, 2000.

[Darden, 1990] L. Darden. Diagnosing and fixing faults in theories. In J. Shrager & P. Langley (Eds.), *Computational Models of Scientific Discovery and Theory Formation* (pp. 319-353). San Mateo, CA: Morgan Kaufmann, 1990.

[Darden, 1991] L. Darden. *Theory change in science: Strategies from Mendelian genetics*. New York: Oxford University Press, 1991.

[Darden and Maull, 1977] L. Darden and N. Maull. Interfield theories. *Philosophy of Science, 43*, 44-64, 1977.

[Descartes, 1637] R. Descartes. *Discours de la méthode pour bien conduire sa raison & chercher la varité dans les sciences*. Leyden: I. Maire, 1637.

[Descartes, 1644] R. Descartes. *Principia philosophiae*. Amsterdam: Apud Ludovicum Elzevirium, 1644.

[Descartes, 1658] R. Descartes. *Meditationes de prima philosophia*. Amsterdam: J. Janssonium Juniorum, 1658.

[Descartes, 1664] R. Descartes. *Traite de l'Homme*. Paris: Angot, 1664.

[Di Chiara, 1999] G. Di Chiara. Drug addiction as dopamine-dependent associative learning disorder. *European Journal of Pharmacology, 375*, 13-30, 1999.

[Donders, 1868] F. C. Donders. Over de snelheid van psychische processen. Onderzoekingen gedaan in het Physiologisch Laboratorium der Utrechtsche Hoogeschool: 1868-1869. *Tweede Reeks, 2*, 92-120, 1868.

[Endicott, 1998] R. P. Endicott. Collapse of the new wave. *Journal of Philosophy, 95*, 53-72, 1998.

[Endicott, 2001] R. P. Endicott. Post-structuralist angst: A critical notice of Bickle's *Psychoneural reduction: The new wave. Philosophy of Science, 68*, 377-393, 2001.

[Feinberg and Farah, 2000] T. E. Feinberg and M. J. Farah. A historical perspective on cognitive neuroscience. In M. J. Farah & T. E. Feinberg (Eds.), *Patient-based approaches to cognitive neuroscience* (pp. 3-20). Cambridge, MA: MIT Press, 2000.

[Flourens, 1846] M. J. P. Flourens. *Phrenology examined* (C. D. L. Meigs, Trans.). Philadelphia: Hogan and Thompson, 1846.

[Franken, 2003] I. H. A. Franken. Drug craving and addiction: Integrating psychological and neuropsychopharmacological approaches. *Progress in Neuro-Psychopharmacology and Biological Psychiatry, 27*, 563-579, 2003.

[Freud, 1922] S. Freud. *Beyond the pleasure principle*. London: The International Psycho-Analytical Press, 1922.

[Garber, 2002] D. Garber. Descartes, mechanics, and the mechanical philosophy. *Midwest Studies in Philosophy, 26*, 185-204, 2002.

[Giere, 1988] R. G. Giere. *Explaining science: A cognitive approach*. Chicago: University of Chicago Press, 1988.

[Giere, 1999] R. G. Giere. *Science without laws*. Chicago: University of Chicago Press, 1999.

[Glennan, 1996] S. Glennan. Mechanisms and the nature of causation. *Erkenntnis, 44*, 50-71, 1996.

[Glennan, 2002] S. Glennan. Rethinking mechanistic explanation. *Philosophy of Science, 69*, S342-S353, 2002.

[Hardcastle, 1996] V. G. Hardcastle. *How to build a theory in cognitive science*. Albany, NY: SUNY Press, 1996.

[Harley, 1879] R. Harley. The Stanhope demonstrator. *Mind, 4*, 192-210, 1879.

[Harré, 1960] R. Harré. Metaphor, models, and mechanism. *Proceedings of the Aristotelian Society, 60*, 101-122, 1960.

[Hegarty, 2002] M. Hegarty. Mental visualization and external visualizations. In W. Gray & C. Schunn (Eds.), *Proceedings of the Twenty-Fourth Annual Conference of the Cognitive Science Society*, (p. 40). Mahwah, NJ: Erlbaum, 2002.

[Heil, 2002] J. Heil. Functionalism, realism, and levels of being. In J. Conant & U. M. Zeglen (Eds.), *Putnam: Pragmatism and realism* (pp. 128-142). New York: Routledge, 2002.

[Hempel, 1958] C. G. Hempel. The theoretician's dilemma. In H. Feigl, M. Scriven, & G. Maxwell (Eds.), *Minnesota studies in the philosophy of science, vol. 2*, pp. 37-98). Minneapolis, MN: University of Minnesota Press, 1958.

[Hempel, 1948] C. G. Hempel and P. Oppenheim. Studies in the logic of explanation. *Philosophy of Science, 15*, 137-175, 1948.

[Holland *et al.*, 1987] J. H. Holland, K. J. Holyoak, R. E. Nisbett and P. R. Thagard. *Induction: Processes of Inference, Learning and Discovery*. Cambridge, MA: MIT Press, 1987.

[Hooker, 1981] C. A. Hooker. Towards a general theory of reduction. *Dialogue, 20*, 38-59; 201-236; 496-529, 1981.

[Hull, 1943] C. L. Hull. *Principles of behavior*. New York: Appleton-Century-Crofts, 1943.

[Ikemoto and Panksepp, 1999] S. Ikemoto and J. Panksepp. The role of nucleus accumbens dopamine in motivated behavior: A unifying interpretation with special reference to reward-seeking. *Brain Research Reviews, 31*, 6-41, 1999.

[Ikemoto *et al.*, 2005] S. Ikemoto, M. Qin, Z.-H. Liu. The functional divide for primary reinforcement of D-amphetamine lies between the medial and lateral ventral striatum: Is the division of the accumbens core, shell, and olfactory tubercle valid? *Journal of Neuroscience, 25*, 5061-5065, 2005.

[Ippolito and Tweney, 1995] M. F. Ippolito and R. D. Tweney. The inception of insight. In R. J. Sternberg & J. E. Davidson (Eds.), *The nature of insight* (pp. 433-462). Cambridge, MA: MIT Press, 1995.

[Jackson, 1931] J. H. Jackson. *Selected writings of John Hughlings Jackson, vol. 1*. London: Hodder and Stoughton, 1931.

[James, 1890/1950] W. James. *Principles of psychology*. New York: Dover, 1890/1950.

[Jevons, 1982] W. S. Jevons. On the mechanical performance of logical inference. *Philosophical Transactions, 160*, 497-518, 1982.

[Kahneman *et al.*, 1982] D. Kahneman, P. Slovic and A. Tversky, eds. *Judgment under uncertainty: Heuristics and biases*. New York: Cambridge University Press, 1982.

[Kauffman, 1971] S. A. Kauffman. Articulation of parts explanations in biology and the rational serach for them. In R. C. Bluck & R. S. Cohen (Eds.), *PSA 1970* (pp. 257-272). Dordrecht: Reidel, 1971.

[Kauffman, 1993] S. A. Kauffman. *The origins of order: Self-organization and selection in evolution*. Oxford: Oxford University Press, 1993.

[Kim, 1998] J. Kim. *Mind in a physical world*. Cambridge, MA: MIT Press, 1998.

[Kleinginna and Kleinginna, 1981] P. R. Kleinginna and A. M. Kleinginna. A categorized list of motivation definitions with a suggestion for a consensual definition. *Motivation and emotion, 5*, 263-291, 1981.

[Koob and Le Moal, 1997] G. F. Koob and M. Le Moal. Drug abuse: Hedonic homeostatic dysregulation. *Science, 278*, 52-58, 1997.

[Koob and Le Moal, 2001] G. F. Koob and M. Le Moal. Drug addiction, dysregulation of reward, and allostasis. *Neuropsychopharmacology, 24*, 97-129, 2001.

[La Mettrie, 1748] J. O. d. La Mettrie. *L'homme machine*. Leyde: E. Luzac, 1748.

[Lashley, 1948] K. S. Lashley. The mechanism of vision: XVIII. Effects of destroying the visual "associative areas" of the monkey. *Genetic Psychology Monographs, 37*, 107-166, 1948.

[Lashley, 1950] K. S. Lashley. In search of the engram, *Physiological mechanisms in animal behavior, vol. iv*, pp. 454-482). New York: Academic, 1950.

[Machamer *et al.*, 2000] P. Machamer, L. Darden, and C. Craver. Thinking about mechanisms. *Philosophy of Science, 67*, 1-25, 2000.

[Malcolm, 1968] N. Malcolm. The conceivability of mechanism. *Philosophical Review, 77*, 45-72, 1968.

[Marquand, 1885] H. G. Marquand. A new logic machine. *Proceedings of the American Academy of Arts and Sciences, 21*, 303, 1885.

[Marr, 1982] D. C. Marr. *Vision: A computation investigation into the human representational system and processing of visual information*. San Francisco: Freeman, 1982.

[McCauley, 1987] R. N. McCauley. The not so happy story of the marriage of linguistics and psychology: or why linguistics has discouraged psychology's recent advances. *Synthese, 72*, 341-353, 1987.

[McCauley, 1996] R. N. McCauley. Explanatory pluralism and the coevolution of theories in science. In R. N. McCauley (Ed.), *The Churchlands and their critics* (pp. 17-47). Oxford: Blackwell, 1996.

[McCauley, 2001] R. N. McCauley and W. Bechtel. Explanatory pluralism and the heuristic identity theory. *Theory and Psychology, 11*, 736-760, 2001.

[McDougall, 1908] W. McDougall. *An introduction to social psychology*. London: Methuen, 1908.

[Miller, 1962] G. A. Miller. Some psychological studies of grammar. *American Psychologist, 17*, 748-762, 1962.

[Miller *et al.*, 1960] G. A. Miller, E. Galanter, and K. Pribram. *Plans and the structure of behavior*. New York: Holt, 1960.

[Miller and Selfridge, 1950] G. A. Miller and J. A. Selfridge. Verbal context and the recall of meaningful material. *American Journal of Psychology, 63*, 176-185, 1950.

[Morrison and Morrison, 1961] P. Morrison and E. Morrison, eds. *Charles Babbage and his calculating engines: Selected writings by Charles Babbage and others*. New York: Dover, 1961.

[Murray *et al.*, 1938] H. A. Murray, W. G. Barrett, E. Homburger, and others. *Explorations in personality: A clinical and experimental study of fifty men of college age, by the workers at the Harvard psychological clinic*. New York: Oxford, 1938.

[Nagel, 1961] E. Nagel. *The structure of science*. New York: Harcourt, Brace, 1961.

[Nersessian, 1999] N. Nersessian. Model-based reasoning in conceptual change. In L. Magnani, N. Nersessian, & P. Thagard (Eds.), *Model-based reasoning in scientific discovery* (pp. 5-22). New York: Kluwer, 1999.

[Nersessian, 2002] N. Nersessian. The cognitive basis of model-based reasoning in science. In P. Carruthers, S. Stich, & M. Siegal (Eds.), *The cognitive basis of science* (pp. 133-153). Cambridge: Cambridge University Press, 2002.

[Newell and Simon, 1972] A. Newell and H. A. Simon. *Human problem solving*. Englewood Cliffs, NJ: Prentice-Hall, 1972.

[Nickles, 1980a] T. Nickles, ed. *Scientific discovery: Case studies*. Dordrecht: Reidel, 1980.

[Nickles, 1980b] T. Nickles, ed. *Scientific discovery: Logic and rationality*. Dordrecht: D. Reidel Publishing Company, 1980.

[Olds and Milner, 1954] J. Olds and P. Milner. Positive reinforcement produced by electrical stimulation of septal area and other regions of rat brain. *Journal of Comparative and Physiological Psychology, 47*, 419-429, 1954.

[Oppenheim and Putnam, 1958] P. Oppenheim and H. Putnam. The unity of science as a working hypothesis. In H. Feigl & G. Maxwell (Eds.), *Concepts, theories, and the mind-body problem* (pp. 3-36). Minneapolis: University of Minnesota Press, 1958.

[Peirce, 1887] C. S. Peirce. Logical machines. *American Journal of Psychology, 1*, 165, 1887.

[Petersen *et al.*, 1989] S. E. Petersen, P. T. Fox, M. I. Posner, M. Mintun, and M. E. Raichle. Positron emission tomographic studies of the processing single words. *Journal of Cognitive Neuroscience, 1*, 153-170, 1989.

[Petersen *et al.*, 1988] S. E. Petersen, P. T. Fox, M. I. Posner, M. Mintun, and M. E. Raichle. Positron emission tomographic studies of the cortical anatomy of single-word processing. *Nature, 331*, 585-588, 1988.

[Posner and Raichle, 1994] M. I. Posner and M. E. Raichle. *Images of Mind*. San Francisco: Freeman, 1994.

[Post, 1936] E. L. Post. Finite combinatorial processes — Formulation I. *Journal of Symbolic Logic, 1*, 103-105, 1936.

[Railton, 1978] P. Railton. A deductive-nomological model of probabilistic explanation. *Philosophy of Science, 45*, 206-226, 1978. Reprinted in J. A. Cover (Ed.), *Philosophy of science: The central issues* (pp. 746-765). New York: W. W. Norton and Company, 1998.

[Reber, 1987] A. S. Reber. The rise and (surprisingly rapid) fall of psycholinguistics. *Synthese, 72*, 325-339, 1987.

[Reichenbach, 1938] H. Reichenbach. *Experience and prediction*. Chicago: University of Chicago Press, 1938.

[Richardson, 1999] R. C. Richardson. Cognitive science and neuroscience: New wave reductionism. *Philosophical Psychology, 12*, 297-307, 1999.

[Robbins and Everitt, 1996] T. W. Robbins and B. J. Everitt. Neurobehavioural mechanisms of reward and motivation. *Current Opinion in Neurobiology, 6*, 228-236, 1996.

[Ruiz-Mirazo and Moreno, 2004] K. Ruiz-Mirazo and A. Moreno. Basic autonomy as a fundamental step in the synthesis of life. *Artificial Life, 10*, 235-259, 2004.
[Salmon, 1984] W. C. Salmon. *Scientific explanation and the causal structure of the world*. Princeton: Princeton University Press, 1984.
[Schaffner, 1967] K. Schaffner. Approaches to reduction. *Philosophy of Science, 34*, 137-147, 1967.
[Schouten and de Jong, 1999] M. Schouten and H. Looren de Jong. Reduction, elimination, and levels: The case of the LTP-learning link. *Philosophical Psychology, 12*, 237-262, 1999.
[Schroeder, 2004] T. Schroeder. *The three faces of desire*. Oxford: Oxford University Press, 2004.
[Shannon, 1948] C. Shannon. A mathematical theory of communication. *Bell System Technical Journal, 27*, 379-423, 623-656, 1948.
[Shannon and Weaver, 1949] C. Shannon and W. Weaver. *The mathematical theory of communication*. Urbana, IL: University of Illinois Press, 1949.
[Simon, 1969] H. A. Simon. *The sciences of the artificial*. Cambridge, MA: MIT Press, 1969.
[Sternberg, 1966] S. Sternberg. High-speed scanning in human memory. *Science, 153*, 652-654, 1966.
[Taha and Fields, 2005] S. A. Taha and H. L. Fields. Encoding of palatability and appetitive behaviors by distinct neuronal populations in the nucleus accumbens. *Journal of Neuroscience, 25*, 1193-1202, 2005.
[Talmy, 2000] L. Talmy. The semantics of causation. In his *Toward a Cognitive Semantics, vol. 1*, pp. 471-549. Cambridge, MA: MIT Press, 2000.
[Thagard, 2003] P. Thagard. Pathways to biomedical discovery. *Philosophy of Science, 70*, 235-254, 2003.
[Titchner, 1907] E. Titchner. *An outline of psychology* (Revised and enlarged ed.). New York: MacMillan, 1907.
[Turing, 1936] A. Turing. On computable numbers, with an application to the Entscheidungsproblem. *Proceedings of the London Mathematical Society, 42*, 230-265, 1936.
[van Essen and Gallant, 1994] D. C. van Essen and J. L. Gallant. Neural mechanisms of form and motion processing in the primate visual system. *Neuron, 13*, 1-10, 1994.
[van Fraassen, 1989] B. van Fraassen. *Laws and symmetries*. Oxford: Oxford University Press, 1989.
[von Eckardt and Poland, 2004] B. von Eckardt and J. S. Poland. Mechanistic and explanation in cognitive neuroscience. *Philosophy of Science, 71*, 972-984, 2004.
[Waskan, forthcoming] J. Waskan. *Models and cognition*, forthcoming.
[Watson, 1913] J. B. Watson. Psychology as the behaviorist views it. *Psychological Review, 20*, 158-177, 1913.
[Weiner, 1986] B. Weiner. *An attributional theory of motivation and emotion*. New York: Springer Verlag, 1986.
[Wernicke, 1874] C. Wernicke. *Der aphasische Symptomenkomplex: eine psychologische Studie auf anatomischer Basis*. Breslau: Cohn and Weigert, 1874.
[Wiener, 1948] N. Wiener. *Cybernetics: Or, control and communication in the animal machine*. New York: Wiley, 1948.
[Wilson, 2003] R. A. Wilson. Pluralism, entwinement, and the levels of selection. *Philosophy of Science, 70*, 531-552, 2003.
[Wimsatt, 1974] W. C. Wimsatt. Complexity and organization. In K. F. Schaffner & R. S. Cohen (Eds.), *PSA 1972* (pp. 67-86). Dordrecth: Reidel, 1974.
[Wimsatt, 1976] W. C. Wimsatt. Reductionism, levels of organization, and the mind-body problem. In G. Globus, G. Maxwell, & I. Savodnik (Eds.), *Consciousness and the Brain: A Scientific and Philosophical Inquiry* (pp. 202-267). New York: Plenum Press, 1976.
[Wimsatt, 1986] W. C. Wimsatt. Forms of aggregativity. In A. Donagan, N. Perovich, & M. Wedin (Eds.), *Human nature and natural knowledge* (pp. 259-293). Dordrecht: Reidel, 1986.
[Wimsatt, 1994] W. C. Wimsatt. The ontology of complex systems: Levels, perspectives, and causal thickets. *Canadian Journal of Philosophy, 20*, S207-S274, 1994.
[Wimsatt, 1997] W. C. Wimsatt. Aggregativity: Reductive heuristics for finding emergence. *Philosophy of Science, 64*, S372-S384, 1997.
[Wise, 2004] R. A. Wise. Dopamine, learning and motivation. *Nature Reviews Neuroscience, 5*, 483-494, 2004.

[Wong, 2000] R. Wong. *Motivation: A biobehavioral approach*. Cambridge: Cambridge University Press, 2000.

[Woodward, 2002] J. Woodward. What is a mechanism? A counterfactual account. *Philosophy of Science, 69*, S366-S377, 2002.

[Wright, 2000] C. D. Wright. Eliminativist undercurrents in the new wave model of psychoneural reduction. *Journal of Mind and Behavior, 21*, 413-436, 2000.

[Wright, 2002] C. D. Wright. Animal models of depression in neuropsychopharmacology qua Feyerabendian philosophy of science. In S. P. Shohov (Ed.), *Advances in Psychology Research, vol. 13*, pp. 129-148). New York: NovaScience Publishers, 2002.

[Wright, forthcoming] C. D. Wright. Is psychological explanation going extinct? In M. Schouten and H. Loren de Jong (Eds.), *The Matter of the Mind: Philsophical Essays on Psychology, Neuroscience, and Reduction*. Oxford: Blackwell, forthcoming.

[Young, 1970] Young, R. M. (1970). *Mind, brain, and adaptation in the 19th century: Cerebral localization and its biological context from Gall to Ferrier*. Oxford: Clarendon Press.

REALIZATION: METAPHYSICAL AND SCIENTIFIC PERSPECTIVES

Robert A. Wilson and Carl F. Craver

1 A SERVANT OF TWO MASTERS

For the greater part of the last 50 years, it has been common for philosophers of mind and cognitive scientists to invoke the notion of realization in discussing the relationship between the mind and the brain. In traditional philosophy of mind, mental states are said to be realized, instantiated, or implemented in brain states. Artificial intelligence is sometimes described as the attempt either to model or to actually construct systems that realize some of the same psychological abilities that we and other living creatures possess. The claim that specific psychological capacities, such as the capacity to understand spoken language, might be realized by different individuals in different ways, has been presupposed by psychologists and cognitive neuroscientists, who design clever experiments in which measures of dependent variables — reaction times, error rates, or localized metabolic activity, for example — provide evidence about what the neurological realizers for specific psychological abilities are.

As common as it is to speak of realization and realizability, these notions have only recently been scrutinized under the philosophical microscope. Much of this work has been critical. Some have identified putative problems with standard views of realization [Wilson, 2001; 2004; Gillett, 2002]; others have challenged the widespread commitment to the thesis that the mental is multiply realized in the physical or biological [Bechtel and Mundale, 1999; Shapiro, 2000; 2004]. But along with the critiques, some positive views have emerged, and we now have a better understanding of several of the desiderata that any view of realization must satisfy.

One thing that has already become apparent is that the concept of realization serves two different masters. A caricature of the two will help to distinguish two different sets of desiderata on accounts of realization.

On the one hand, there is the *Metaphysician of Mind*, concerned primarily with, to use C.S. Lewis's phrase, the place of mind in the world order. With the exception of the recently burgeoning area of consciousness studies, analytic philosophy of mind has been predominantly physicalist (in one sense or the other) since the classic statements of the mind-brain identity theory by J.J.C. Smart and U.T. Place in the 1950s. But this broad commitment to physicalism does not take the Metaphysician of Mind as far down the path to understanding the mind as one

Handbook of the Philosophy of Science. Philosophy of Psychology and Cognitive Science
Volume editor: Paul Thagard
General editors: Dov M. Gabbay, Paul Thagard and John Woods
© 2007 Elsevier B.V. All rights reserved

might hope, and has given rise to its own problems. Consider several of these. Is the mind strictly identical to the brain, or should physicalists endorse some other kind of relationship between the two? Wherever one stands on this issue, what room does this leave for the reality of the mind, for genuine *mental causation*, for an understanding of consciousness and content, or for distinctive aspects of our grasp of the mental, such as first-person knowledge?

In serving this master, the concept of realization has been slotted into a particular network of technical concepts, such as supervenience, metaphysical sufficiency, and nomic necessity. That network of concepts is also partly constituted by a range of by now familiar "isms"–physicalism, functionalism, computationalism, and reductionism being the four most common. The Metaphysician of Mind uses conceptual analysis as a way of exploring entailments and tensions between positions that one might adopt using those concepts. This approach is exemplified by the work of Jaegwon Kim. Best known for his sustained work on supervenience and mind over the past 30 years, Kim has more recently turned his attention to realization in extending his critique of non-reductionist views of the mind.

But the Metaphysician of Mind is not the only master that the concept of realization must serve. It must also serve the *Cognitive Scientist*. For the Cognitive Scientist, it is not the place of the concept of realization amongst any network of concepts that needs to be understood, but how specific psychological functions and capacities are or can be realized by particular psychological and neurological structures and mechanisms. These realizations are best explored through the construction of models or schemata of the corresponding processes, models or schemata that can be specified at various degrees of abstraction. These schemata range from the very abstract functional decompositions ("boxology", to its detractors) that one finds in much of cognitive psychology and artificial intelligence, to biochemically detailed accounts of specific neural pathways. For the Cognitive Scientist, the realization of the mental is to be investigated through such familiar strategies as localization of function and physical decomposition. Such strategies start with the cognitive capacities or behaviors of an organism (plus or minus a bit) and then proceed to explain them in terms of the capacities or behaviors of its parts (e.g., parts of the central nervous system). This is how, for example, cognitive neuroscientists and psychologists investigate short-term memory or visual shape recognition.

In saying that the concept of realization must serve the Cognitive Scientist, we are helping ourselves to an assumption that is widespread among Metaphysicians of Mind and philosophers of cognitive science. This assumption is that although cognitive scientists seldom use the term "realization", much of what they say and do can be reconstructed with the help of this term. Psychologists and cognitive neuroscientists are much more likely to talk of the *neural correlates of*, the *neural substrates of*, the *neural mechanisms for*, or the *implementation of* psychological capacities than to talk of how those capacities are "realized" by specific mechanisms and processes. As Jaegwon Kim says, acknowledging that the meaning of "realize" has not been fully explained in the philosophical literature in which it

features, "you will not go far astray if you read 'P realizes M' as 'P is a neural substrate, or correlate, of M" [1996, 102, fn.4]. So the idea is that these ways of talking about the relationship between the psychological and the physical, prevalent in the discourse of scientists themselves, can be viewed as less metaphysically committed invocations of the concept of realization.

The slack between talk of neural correlates and mechanisms within the cognitive sciences and philosophical talk of realization deserves at least a brief comment. The short answer to the question "How do cognitive scientists conceptualize realization?" is: "They don't". Or, to put it more accurately: they do, but almost exclusively when they are attempting to sketch a broader location for their particular views, or when interacting directly with their philosophical interlocutors. Given that, the assumption that we can gloss the scientific talk of the physical correlates of, the underlying mechanisms for, or localized structures causally implicated in, the operation of some particular psychological capacity in terms of the notion of realization may not be cost-free. It may require that we adjust our view of what appeals to "realization" can do, or leave us with a philosophically emaciated view of realization (cf. [Polger, 2004]).

It is our view that realization is and will continue to be, a servant to these two masters, and that this is not altogether a bad thing. But two masters they are, and each make distinct demands on the concepts that serve them. Consider the following list of desiderata that would surely make any Metaphysician of Mind happy. They would like a view of realization that:

- elucidates the relationship between functionalism, physicalism, and reductionism
- enables us to at least clarify what mental causation and multiple realizability involve (if not tell us whether they occur in the domain of cognition)
- points the way to an explanation of how organized networks of neurons give rise to the full range of mental phenomena

They might, of course, want more, but this would be at least a start.

Contrast these desiderata with those that a Cognitive Scientist could not just live with but live for. They would like a view of the realization of psychological capacity, P, in neural structures N, that:

- identifies the N that are relevant to the operation of P
- reveals what variation exists in the relationship between P and N across different populations
- provides a step-by-step account of just how P is realized by N in any particular instance

Again, as a Cognitive Scientist, one might expect more — one might like an account of realization to guide the design of experiments, or to allow a clear distinction

between adequate and inadequate realizing explanations. But it seems that this would be a minimal list of desiderata for any Cognitive Scientist.

The trivial point about these two sets of desiderata is that they are different. Less trivially and more speculatively, projects that draw on a concept satisfying one set of desiderata may do little to satisfy the other. For example, the experiments that show that the hippocampus is involved in episodic memory will likely tell us little about how physicalism, functionalism, and reductionism are related, and will leave us none the wiser about mental causation or multiple realization in general. Conversely, suppose that we had a "deep explanation" of just why neural networks give rise to particular mental phenomena. Even such an apparent breakthrough for the Metaphysician of Mind may rely on or lead to a concept of realization that leaves us with a blank stare when it comes to providing a step by step account of just how auditory hallucinations are produced in cases of delusional schizophrenia.

More pessimistically (and even more speculatively), it is possible that the concept of realization satisfying one of these sets of desiderata must be different from the concept of realization that satisfies the other. While a certain kind of peaceful coexistence could persist were the preceding scenario to eventuate, to find that there was some kind of deep incompatibility between the desiderata of the Metaphysician of Mind and of the Cognitive Scientist would be a sort of intellectual disaster. So much so, we think, that one of the desiderata on the list of each should be that their view of realization should be at least consistent with (ideally, well-integrated with) that of their counterpart.

2 REALIZATION AND THE METAPHYSICS OF MIND

To get some sense of the role that appeals to realization have played in the philosophy of mind and cognitive science, and of how realization has come to serve two masters, let's look at how and why talk of realization was initially introduced.

Hilary Putnam brought the concept of realization into contemporary philosophy of mind in 1960 in his classic paper "Minds and Machines". In that paper Putnam described the relationship between the mental and the physical as one of realization. In doing so, Putnam drew an analogy between minds and machines. In particular, he argued that the relationship that holds between minds and brains is the same as that holding between abstract Turing machines and the physical arrangements of matter in which they are instantiated or implemented: realization.

Part of Putnam's point was to dissolve the mind-body problem. He claimed that the relationship between mind and body, or mind and brain, should be no more puzzling — indeed, no more interesting — than that between the abstract states of a given Turing machine and the structural states of the device realizing it.

Putnam's original introduction of "realization" also formed a part of the cognitive revolution that spawned the cognitive sciences as we know (and love) them. In doing so, it also provided a way of thinking about the relationship between the mental and the physical taken up within psychology, linguistics, and computer science themselves. At least that is part of philosophical lore. The idea was that

even if we did not hear the term "realization" in the mouths of those working within these disciplines, their exploration of the mechanisms underlying psychological functions, capacities, and abilities could be adequately glossed in terms of more metaphysical-sounding notions, such as realization. As we said in the first section, talk of the neural correlates of, or of the neural mechanisms for, a given psychological capacity have been viewed as loose science-speak for something like the relation of realization.

Within a dozen years or so of its introduction, realization came to be seen as useful in articulating three of the closely related "isms" that had, by that time, gained currency: physicalism, functionalism, and computationalism.

First, as a metaphysical relation weaker than identity, realization was thought suitable for developing a brand of physicalism that made room for the autonomy of psychology, for genuine mental causation, and for psychological laws. And, unlike identity, as an asymmetrical relation, realization was also thought well suited to capturing the dependence of the mental on the physical, or the *determinative* nature of underlying, physical states.

Second, as functionalism came to replace type-type identity views as the theory of choice for materialists in the philosophy of mind, the physical was seen as realizing the mental, rather than being strictly identical to it. Functionalism is the view that psychological states are to be identified with the causal or functional role they play in the overall causal network of psychological states. We identify psychological states in terms of their functional roles much as one might identify a part of a machine in terms of the causal role that it plays in the operation of the machine, or as one might identify a job in an organization, such as a manager, in terms of what the person in that position contributes to the overall set of tasks that the organization undertakes. In all of these cases, there is a distinction to be drawn between the role itself and the occupant of that role. In the psychological case, brain states and mechanisms are taken as the occupants of the functional roles defining psychological states. The concept of realization was drawn on to tighten this intuitive notion of an occupant.

If we extend these analogies, then it is seemingly easy to infer that the functionalist view of the mind is committed to the idea that mental states are *multiply realizable* in physical states. For just as different persons can fill one and the same role within an organization, and parts of a machine can be replaced by other parts, sometimes parts quite different in many of their properties, so too in principle can psychological states be realized by very different physical states. We have multiple realization just when we have the *same kind* of psychological state realized by physical states that are *different in kind* from one another.

Reflecting the proclivity of philosophers for cutting edge technology, a favorite example used to illustrate this idea was the mousetrap. What mousetraps *are*, what they are to be strictly identified as, is a kind of device that is designed to catch mice. It has that function, and a mousetrap can be thought of, perhaps arcanely, as an input-output device that accepts live or unconfined mice as inputs and delivers dead or confined mice as its output. But as every cat knows, there

are many ways to catch a mouse. (Correlatively, as every mouse knows, there are many ways to skin a cat.) Mousetraps are realizable by different kinds of physical devices: traditional neck-snappers, balance cages, and pit traps, for example.

Third, functionalism and the appeal to (multiple) realizability sat well with the rise of the computational metaphor and the nascent cognitive sciences to which that metaphor is central. As Putnam's original appeal to Turing machines suggested, minds are to brains as programs are to the hardware that runs them. Computationalism moves beyond this sort of appeal to the computational metaphor in providing a more precise way of characterizing the functional roles that define psychological states. Computationalism takes these roles to be specifiable as *algorithms*, i.e., as effective procedures that can be realized in principle, and often in practice, by machines.

Turing machine functionalism, the earliest and in many respects the most problematic form of computationalism, held that these algorithms were computable by a universal Turing machine — not an actual machine but a "theoretical machine" conceptualized by the mathematician Alan Turing in his well-known work on the foundations of logic. A Turing machine is a simply structured device that can perform several mechanical functions: it can scan a symbol, write a symbol to a specific location, or erase a symbol from that location, and move its scanner from one location to the next. Although it is common in the philosophy of mind and cognitive science to talk of mental states as being realizable by particular Turing machines, note how quickly this leads to confusion, given that Turing machines are themselves characterized purely functionally.

Computationalism has been central to classic, early work in artificial intelligence, to more recent work in the computational modeling developed by connectionists, and to those using Bayesian and statistical techniques to understand psychological functions and capacities. *Strong* artificial intelligence can be characterized as the view that appropriately designed computational programs would not simply be models of but realizations of certain psychological capacities, such as the capacity to understand (fragments of) natural language. Given that the computer hardware that realizes such programs are of a distinct physical kind from the brainware that realizes these very same programs in us, strong AI entails the multiple realizability of the mental in the physical.

In this section we have concentrated on the development of the concept of realization from the point of view of the Metaphysician of Mind, albeit one who also likes to think of himself as allied with the Cognitive Scientist. We now turn our attention to the view of realization implicit in the work of the Cognitive Scientist.

3 REALIZATION AND THE COGNITIVE SCIENCES

The Cognitive Scientist, as noted above, rarely makes the notion of realization explicit. Instead, the idea of realization is implicit in the practice of explaining by decomposing and in a variety of experimental techniques used to test claims to

the effect that some N realizes some P. We can learn something about this implicit notion of realization by looking at an exemplar case: the voltage-gated sodium channel. We can then turn our attention to a higher-level cognitive case.

Our first example suggests that the explanatory practices of science embody several different varieties of realization. The primary division is between the material realization of an entity and the explanatory realization of a property. A secondary division, holding among explanatory realization relations, turns on whether what is getting realized is a property or an activity, and on the relevant kind of organization in the realizer. The example will help to make these distinctions transparent.

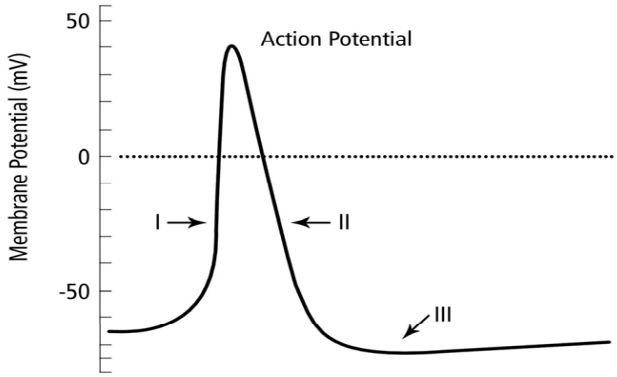

Figure 1.

Suppose that we want to understand how neurons generate action potentials. Action potentials are electrical waves that propagate along the axons of neurons and are generated near the cell body (in a structure known as the axon hillock) in response to summing excitation from neighboring cells. In their resting state, neurons have a voltage gradient across their membranes, and an action potential (as shown in Figure 1) is a fluctuation in that gradient that propagates from the hillock down the length of the axon. The explanation for the action potential involves breaking a neuron and its membrane into parts and showing how their activities are organized to produce the action potential.

The action potential is realized by the cooperative activities of a variety of ion channels in the membrane. One of these channels, the one responsible for the initial rising phase (I) of Figure 1, is the voltage-gated sodium channel. (Detailed description of the sodium channel can be found in most neuroscience texts (e.g., [Kandel and Schwartz, 2001]). One hoping to understand how the sodium channel is responsible for the rising phase of the action potential will similarly appeal to its parts and their cooperative activities. But as we look more closely at this exemplar

of explanation in neuroscience, it quickly becomes clear that there are several varieties of realization in play. These differences reflect the fact that different kinds of things are getting realized and that different kinds of things are doing the realizing.

One part of understanding the action potential will involve understanding the material composition of the sodium channel. This membrane-spanning channel is a protein, and so it is made of amino acids. If we ask what realizes the sodium channel, we might appeal to different parts — subunits of the channel protein, or stretches of the sequence within those subunits, or the atoms bound into amino acids — but in each case we are appealing to the material parts composing the sodium channel. And this is the case whenever we look for the realizer of an entity or entity kind.

Sometimes entity kinds have been the Metaphysician's primary examples. They have appealed to jade, water, or corkscrews as realized items (see [Shapiro, 2000; 2004]). And when we talk of entity kinds being realized, we tend to point to the material realizers — the units of material (however those are to be individuated) located within the entity's spatial boundaries. Jade (or more specifically, nephrite) is realized by a silicate of calcium, magnesium, and iron. Corkscrews are realized by steel. Computers are realized by silicon. When we say such things, we are talking about the realization of an entity as an entity (a spatially bounded object) rather than as some particular kind of entity (i.e., under some description). For example, we neglect the organization of components that gives the sodium channel its shape and unique activities, and pay attention instead only to the matter in its boundaries.

Learning how entities are materially realized is often an important descriptive stage of science. Learning the primary sequence of the sodium channel, for example, allows one to begin to sort its various orders of structure, from the sequence (order) of amino acids, beginning to end, to how these give rise to local spatial forms (e.g., helices and sheets) and subsequently to more complicated folding arrangements that constitute the whole protein. Cell fractionation and centrifugation were essential to the development of biochemistry and molecular biology precisely because they allowed investigators to discover the constituents of cells. Learning the material realizers of the sodium channel has been similarly important to learning how it works and why it has the properties that it does.

The task of specifying a material realizer for an entity is complete once one has listed (or specified) its constituents exhaustively. Such completeness is a descriptive success. When attention shifts from describing to explaining, however, entities and their material realizers are no longer the primary focus. Instead, attention shifts to properties and activities as realized kinds and, correlatively, to component parts, their properties, their activities and their organization. Explanatory forms of realization are not just exhaustive lists of material constituents, but selective descriptions of the relevant parts for some explanatory purpose. The explanation neglects some properties of the material realizers and accentuates others. For explanatory realization (as opposed to material realization), it frequently does matter

which parts one attends to: arbitrary parts will not be explanatorily relevant. We can distinguish three varieties of this explanatory form of realization: realization of an aggregate, realization of a structural property, and realization of an activity. Consider these in turn.

The wave form of the action potential is not the product of a single sodium channel but of thousands of channels in the cell membrane. The axon hillock is especially dense in such channels, and when large numbers of channels activate at once, they give rise to the wave form for the cell as a whole. This shift from single channels to populations of channels is often simply presumed in an explanation (note that this step is typically left out of standard textbooks) because the total current produced by opening the sodium channels in the hillock is approximately the sum of the currents flowing through individual sodium channels. The total current is an example of what Wimsatt [1997] has called an *aggregative property*. The mass of a pile of sand is realized aggregatively by the masses of the individual grains. Other subunits, such as half-grains of sand, would work just as well. Likewise, the total sodium current is a sum of the currents through individual sodium channels. In contrast to material realizers, aggregate realizers are selective in that some properties of the realizers are important for the explanation and some are not. It is the mass of the individual grains (and not their color, melting point, or texture) that fully constitutes the mass of the whole pile. In purely aggregative cases of realization, it makes sense to speak of realized properties as being composed of realizing properties: the realized property and the realizing properties are different values of the same variable, and the value of the realized variable is exhausted by the sum of the values for the same variable for the individual realizers. Unlike material realization, however, cases of aggregate realization are selective: the realizers include only some of the properties of the constituents (just those involved in the sum). Explaining the mass of the pile of sand appeals only to the masses of the components, and not to their color, for example. And understanding the current through the population of sodium channels involves calculating only the values of the currents through the individual channels (or at least so it is often presumed in explanations).

Truly aggregative properties are difficult to come by. Indeed, the opening of different sodium channels is not, properly speaking, aggregative, since the influx of sodium gradually changes the charge of the cell and the relative concentrations of sodium inside and outside of the cell, each of which is relevant to the rate of ionic flow through the channel. Nonetheless, some explanations in the sciences of the mind do depend on properties that are closer to the aggregate end on a spectrum of organization, and they represent an important type of realization in these sciences. For example, population effects of neurotransmitters at the synapse and sums of excitation to a node in a connectionist network each approximate this ideal.

Sometimes it is a structural property that calls for explanation. One may want to understand, for example, how the sodium receptor has its characteristic sequence of amino acids (an explanation fleshed out by appeal to translation and

transcription), how it forms a channel through the membrane, or how the channel selectively allows (primarily) for the flow of sodium ions. Of particular explanatory interest in the effort to understand the action potential is the sequence of amino acids in each of the repeating subregions of the channel known as the S4 region. The amino acids in this region are ordered such that every third amino acid residue is an arginine or a lysine. Given the charges on these residues, this linear order of components produces a helical structure with evenly spaced positive charges. Leaving aside for the moment the significance of this arrangement, one might want to understand why that type of amino acid sequence produces an alpha helix. For this, one will appeal not merely to the material constituents of the helix, but to the sequence of amino acids and their polar and nonpolar (hydrophilic and hydrophobic) properties. It is these that determine how the molecule folds and coils in water. Structural forms of realization thus lay out not only the (relevant) material components, but also the features that determine their overall shape and configuration: not just the matter but the spatial forms of organization.

The details of molecular folding have been worked out for relatively simple structures such as the alpha helix, but are only partially worked out for even moderately complex forms of tertiary structures. As in the case of aggregative realization, there is a straightforward sense in which the spatial properties of the parts are part of the structure of the whole. Unlike cases of aggregative realization, however, the lower-level properties are not summed, but may involve interaction and organization of the components (as the evenly spaced charges participate in S4's helical structure).

Now suppose that the realized kind is not merely an aggregate or structural property of the sodium channel but an *activity*: something that the sodium channel does. Suppose, for example, that we want to understand how the sodium channel is activated by the summing excitation of cell body. As with the above discussion of the alpha helix, the explanatory task selectively directs attention preferentially to some parts of the receptor and away from others. The helix is significant because of its putative role as an activating gate for the sodium channel. In cases of *mechanistic realization*, the burden of realization is borne by some constituents (working parts or components) more than others, and organization among the components figures increasingly in our description of the realizer. Let us continue with our example of the sodium channel.

Under the resting electrical conditions, a positive extracellular potential holds the alpha helix — lined with positive charges — in a stable position in the membrane. Weakening that positive potential, as occurs when the cell is depolarized from its resting state, allows the helix to rotate out toward the extracellular side. This rotation is thought to occur in each of the sodium channel's subunits which destabilizes the balance of forces holding the channel in its closed state and allows it to move to a new equilibrium state, one that has an open channel through the membrane. There is, as of yet, no complete story to tell about how the displacement of the helix, coupled with the attraction and repulsion among component polar and nonpolar amino acids, alters the conformation of the entire protein (the

mathematics required to predict how such a change would ramify through the molecule is daunting). But we nonetheless know the component processes and can see how those properties and activities could, if organized appropriately, change the conformation of the channels. Mathematical models are often used to simulate the folding in such complex arrangements, thereby showing how complex structures can be realized by such simple component forces and sequential organization and how they can be altered by such factors as voltage changes in the cell.

Mechanistic realizers are composed of the working parts of the mechanism, such as amino acids, membranes and ions and their activities, the things that these entities do: their repelling, rotating, opening. The components are organized such that they exhibit the behavior of the mechanism as a whole [Machamer *et al.*, 2000]. There are many varieties of organization in mechanisms. Spatial properties of the parts (e.g., their size, shape, and orientation) or spatial relationships among the parts (e.g., their positions, compartmentalization, fit, motion) are as important for mechanistic realization as they are for structural realization. But mechanistic realizers also have a temporal component: the activities in the mechanism have characteristic orders, rates and durations. And, finally, the parts of this mechanism act and interact with one another such that they exhibit the behavior of the mechanism as a whole. Mechanisms, in other words, also have a causal component to their organization.

Descriptions of mechanistic realizers are selective. There are no mechanisms simpliciter; all mechanisms are mechanisms *of* something. It is by reference to this behavior of the mechanism as a whole that the relevance of components is established. In describing the activation of the sodium channel as we just have, we selectively attend to the alpha helix and selectively neglect those aspects of the receptor responsible for the ion selectivity, or the channel's inactivation, or the binding cites for neurotoxins. Material realization includes all of these parts; mechanistic realization includes only the relevant ones. The description of a mechanistic realizer is considered complete when all of the relevant components have been included such that it is possible to describe its working in terms of intelligible activities from beginning to end, without gaps or promissory notes. (See [Wilson and Keil, 1998] on the surprising shallowness of our grasp of everyday mechanistic explanations.)

The search for mechanistic realization is embodied in many of the techniques of cognitive neuroscience. Lesion and stimulation experiments, reaction time studies, PET and fMRI experiments are all designed to tease apart the mechanistic realizer for cognitive phenomena. In "bottom up" experiments, one intervenes to remove or stimulate a component and monitors for changes in the behavior of the mechanism as a whole. In "top down" experiments, one manipulates the behavior of the mechanism as a whole and monitors changes in the states or activities of putative components. The goal is to show that an item has a function within the mechanism and to characterize what, precisely, that function is.

Perhaps the most plausible way to develop the link between the search for mechanisms in neuroscience and the notion of realization at play in the metaphysics

of mind is through attention to the notion of a *function*. We have already briefly discussed functionalism in the philosophy of mind, where "functional role" and "causal role" are often used interchangeably. But there is a more constrained notion of function that has been used by biologists, one that links functions to mechanisms, and it is this notion that is particularly relevant to linking mechanisms to realization. Roughly for now, the function of X is what X does or is supposed to do in the mechanism in which X operates. On this view, the function of X is not the total causal role that X plays but some more selective subset of the causes and effects of X: those relevant for the explanatory purposes at hand. One describes X's causal role when it is possible to show how X is organized (spatially, temporally and actively) into a higher level mechanism. Only a subset of its total causal role will be directly relevant to that higher level mechanism.

Moreover, functions are frequently *hierarchically realized* in that higher level functions are achieved through the performance of sub-functions. To take a stock example, one of the functions that the circulatory system performs is to move blood around the body, and one of the functions necessary for this is for there to be a source of force for moving the blood. The parts of the heart, their activities, and their organization together constitute the mechanism for this sub-function, even though there are many other inputs to and outputs from the heart (e.g., it makes noises, produces heat, etc.).

This appeal to biological systems, their functions, and the mechanisms that realize them has a natural application to cognition. Cognitive systems are composed of mechanisms that perform specific functions, and those mechanisms in turn are composed of further mechanisms with even more specific functions, together constituting a hierarchy of mechanisms. For example, the mammalian visual system is composed of a series of mechanisms — the retina, the superior colliculi, the lateral geniculate nuclei, primary visual cortex (V_1), and the extrastriate areas (such as V_{2-4} and MT). Each of these has its particular components that perform specific functions, and in turn they are chunked together to form larger functional units.

As we noted at the start of this chapter, scientific notions of realization are implicit in the practice of explaining by decomposition. But that practice involves different kinds of realization. This difference in kind tends to track different criteria for assessing the success of the description of a realizer. Descriptions of material realizers are complete when they include all of the material constituents of the realized entity. For aggregative properties, it should be demonstrable that the realized property is a sum of lower level properties. For structural properties, determinants of shape and position are of primary importance. And finally in cases of mechanistic realization for activities, the goal is to provide a complete description, without gaps, from the beginning of the mechanism to the end, exhibiting the relevant entities, properties and activities, and showing how they are organized together within the realized activity.

4 A WORKING ACCOUNT OF REALIZATION

One way to proceed in articulating the concept of realization further would be to provide a traditional philosophical analysis of that concept, a set of individually necessary conditions that are collectively sufficient for anything to be a realization. Because we are deeply skeptical of the likelihood of success of such an approach (here, as elsewhere), we offer instead the following working account of realization. This working account serves both to articulate some commonalities that are shared by Metaphysicians of Mind and Cognitive Scientists and to sharpen some of the differences that separate them.

Let's begin with a simple canonical statement of the realization relationship:

(R_0) An object O's having property or activity A is realized by O's having property or activity B.

We have stated R_0 as a relationship between properties, since this is a common way of proceeding in the literature. For material realization, R_0 would be a relationship between objects, and for mechanistic realization, the relata would be an activity and its mechanism. R_0 describes an ontological relationship: it is a relation between O's having A and O's having B rather than between statements about O, P and Q, or between theories or explanations involving O, A, and B. O's having B is the realizer, and O's having A is what's getting realized. O is an object, in the minimal sense that it is a bearer of properties, an agent of activities, or a relatum in relations that constitute A and B.

R_0 is typically taken to be *asymmetric*: O's having B realizes O's having A and not the other way around. This restriction helps to distinguish realization from identity. This asymmetry can be understood as a dependency relation: O's having B is *sufficient but not necessary* for O's having A. To accommodate the possibility of multiple realization, both the metaphysician and scientist will want to allow that O's having B is unnecessary for its having A; A could be realized by different objects and different properties. On this much, it seems to us that the Scientist and the Metaphysician can agree. Let's consider some candidates for strengthening the realization relation, some of the motivations for so strengthening it, and some of the limitations of doing so.

4.1 Must Realization Be Decompositional?

The most common scientific conception of realization — discussed above in association with explanation by analysis into parts — tends to be *decompositional*. In cases of decompositional realization, a property of the whole is realized by the properties of its parts:

(R_1) O's having P is realized by parts of O having Q

For example, the mousetrap's behavior of catching mice (or its capacity to catch mice, P) is realized by the parts of the trap (a trigger, a latch, an impact bar) and

their organized activities (Q). The realized property is a property of the whole, and the realizing properties are the intrinsic and relational properties of the parts. By way of contrast, if one were to say that the mouse trap's being red is realized by its being some determinate shade of red (e.g., scarlet), this would not be a decompositional realization relation. Similarly, David Marr's [1982] now-famous computational, algorithmic and implementation levels are not decompositional. Algorithms are not parts of computations but instead are different descriptions of one and the same thing. In these last two examples, it is not the parts of O that do the realizing, but rather a different property of O as a whole.

Metaphysicians, who associate realization with discussions of physicalism and functionalism, may not be concerned with understanding the relationship between a whole and its parts *simpliciter* but rather with the relationship between the whole and the parts *plus their organization*. They are asking: is there anything more to O and its properties than the parts of O, their properties, and their organization? Spooky varieties of emergence and vitalism are spooky *not* because they insist that the whole is greater than the sum of its parts — even the most ardent physicalist will grant that the organization of the parts realizes properties of the whole that would not be realized if the parts were organized differently — but because they insist that the whole is greater than the parts plus their organization together. Because many metaphysicians who are interested in mind-body relations are interested in whether the mind is anything "over and above" the organized activities of neurons (and molecules, brain regions, etc.), they need a notion of realization that frames that question.

Along these lines, Kim has suggested that realization is a relationship not between "micro-macro" levels but between orders. Kim thinks of levels as sorted by a mereological relation: things at lower levels are parts of the things at higher levels. He calls this hierarchy of levels the "micro-macro hierarchy." Orders, however, are sorted by a dependency relationship: things have their higher-order properties in virtue of having their lower-order properties. Kim stresses that levels and orders should not be confused:

> Notice the following important fact about this [realization] hierarchy: *this hierarchy does not parallel the micro-macro hierarchy* — to put it another way, *the realization relation does not track the micro-macro relation*. The reason is simple: *both second-order properties and their first-order realizers are properties of the same entities and systems.* [1999, 82]

Realization, on this view, is like Marr's levels or like the relationship between O's being red and O's being scarlet. This view can also be extended to the mousetrap. The mousetrap has the ability to catch mice because it has the property of being an organized collection of triggers, springs, and levers (and perhaps a mouse). Having the ability to trap mice and being an organized collection of triggers, springs and levers, on this view, are not properties at different levels but properties of different orders: the trap has the ability to catch mice in virtue of the organized relations

among its parts. The description of the parts plus organization and the description of the mousetrap are descriptions of one and the same thing, and so the realization relation, on this view, is not decompositional. It is not an interlevel relationship.

Scientists, unlike metaphysicians, spend little time worrying about whether wholes are greater than their parts plus their organization. It is a working assumption that the wholes just are the parts organized together. When they say that "the whole is greater than the sum of the parts," they should be taken as literally talking about sums of parts. When the whole is greater than the sum, this is because the parts are organized in a particular way, they interact with one another, and they do things together that they could not do alone. But these assertions are entirely consistent with the idea that wholes just are the parts and their organization. Perhaps empirical investigation will shake the scientist from this working assumption, as the periodic revival of forms of holism and vitalism in the history of biology and neuroscience might attest. But such revivals are not taken especially seriously by those engaged in the task of seeking lower-level explanations. The lack of any lasting success for vitalism or emergent properties in the past and the corresponding successes of sciences driven by the search for lower-level explanations are for many convincing reasons to ignore holism and vitalism and to continue to seek such explanations. Scientists spend more time worrying about whether they have identified the right parts and the right aspects of their organization to fully explain O's having P in terms of the parts of O having Q. Failure to discover such an explanation is only rarely accompanied by calls to abandon this physicalist working hypothesis and much more commonly accompanied by calls to do more experiments.

Be that as it may, this issue points to an important pragmatic difference in the uses to which scientists and metaphysicians are putting the concept of realization. The metaphysician is concerned with evaluating the status of physicalism while scientists largely presume some form of physicalism. The latter are concerned with questions like: "What sort of realizer could realize a property like that?;" "What parts are included in this realizer, and which parts are irrelevant?;" "How are these parts organized together such that the realize that property?" These latter questions are more likely to be addressed by attention to varieties of realization and criteria of adequacy for describing and evaluating hypothesized realizers (see section 3) than they are by attention to global metaphysical principles. On the issue of decompositionality, then, it may turn out that the metaphysician and the scientist do not (and perhaps need not) agree.

4.2 Must Realization be Constitutive or Intrinsic?

A second possible restriction on the realization relation would be to make it constitutive — that is, to require that the realizer of O's having P be wholly contained within O's spatial boundaries. This is the view that Wilson [2001; 2004] calls the *constitutivity thesis* about realization, a thesis that is widely shared amongst Metaphysicians of Mind. The problem with this constraint is that many of the

activities in which objects engage cannot be realized by the parts of those objects alone in that these parts — along with their properties and how they are organized — are not metaphysically sufficient for many of the activities in which those entities engage. In such cases, it is plausible to argue that if realizers are to be metaphysically sufficient for the properties and activities they realize, then they must extend beyond the boundary of O and beyond the relationships that exist among the things inside O. The constitutivity thesis would rule out such appeals to contextual realizers, and might be formulated in the following putative constraint:

(R_2) In C, O's having P is realized by the intrinsic physical properties, Q, of O and its parts.

To see why $R_2$is controversial, consider some of the functions that might be explored in mechanistic realization. We might describe the trap as firing, or as catching mice, or as protecting the cookie jar. We might describe the heart as pumping blood, as circulating blood, or as delivering the poison. If we say that the trap is catching mice or protecting the cookie jar, then we are referring to the world beyond the trap. Its function is then world-involving in a way that suggests that the realization of that function (as sufficient for "catching mice" or "protecting the cookie jar") must also extend beyond the trap or the cookie jar. If the heart is taken to be pumping the blood, circulating the blood or delivering the poison, then explanations of these activities will necessarily include different aspects of the heart's context.

So if we are to embrace a view of realization that provides an account of the realization of these sorts of relational properties, activities, and functions, then it seems that we should reject the idea that realizations are constitutive or intrinsic. But perhaps we should go further. Once we attend to the role of context in realization in cases like these, even the relatively weak claim that there are *any* instances of truly intrinsic realization seems difficult to maintain. Suppose that we consider not the heart's pumping of blood but its beating, an activity that we might think of as being purely intrinsic to it. Yet even the beating of the heart, truth be told, cannot be explained by reference only to intrinsic properties of hearts. Sympathetic and parasympathetic inputs to the heart, both from neural input and from circulating hormones, all contribute to different aspects of the heart's beating (its rate, regularity, etc.). It seems likely to us that mechanistic realization, as a species of explanatory realization, will typically (if not always) involve contextual in addition to constitutive factors. And so there seem to be good reasons for rejecting R_2 as a constraint on realization.

4.3 Must Realization be Synchronic?

It might be suggested, and most seem to assume, that realization is a synchronic relationship. We might build this into our canonical realization statement as follows:

(R_3) In C, O's being P over the interval $[t_o, t_f]$is realized by O's being Q over the interval $[t_o, t_f]$.

Features of O at $t_n < t_o$ or at $t_m > t_f$ cannot be part of P's realizer. One reason that the metaphysician might want to make realization synchronic is that P's realizer is supposed to explain the "causal powers" individuative of O's being P — that is, the set of physically possible causes and effects of O's being P. Since it seems reasonable to suppose that O's causal powers depend exclusively on the properties of O (along with its parts and its context) at a given time — and not to features of O's past — one might wish to restrict the realization relation to the occurent properties of O.

The problem with this putative constraint is that there are well-known properties that are individuated by their histories. Consider first objects of which this is true. The Mona Lisa would not be the Mona Lisa if Brad Pitt painted it instead of DaVinci. Dollar bills would not be dollar bills if Mugsy made them rather than the treasury. And a mouse would not be a mouse if it did not descend from parents that were mice. For those who have called attention to such cases, this is not an empirical fact: we could not discover that classic paintings and dollar bills could be made by anyone or that mice could in fact be made in test-tubes (perhaps they can, but that is very much beside the point). At best we could discover that things very much like the Mona Lisa, dollar bills and mice could be fabricated down to the last detail — but they would still not be the Mona Lisa, dollar bills or mice. In order to accommodate such cases, we may wish to recognize *historical realization* as a sub-species of realization. Historical realizers, in all of the cases with which we are familiar have contextual components. They include relationships to objects outside of O, such as DaVinci, the U.S. Treasury, and parents. But in addition, they include non-occurent, historical properties of O.

A number of philosophers have appealed to historical realization to, as Elliott Sober has said, "put the function back in functionalism" [Sober, 1985]; see also [Neander, 1991; Millikan, 1993]. According to one standard account of biological functions, some item S has the function of T-ing in O (as opposed to merely functioning as a T-er in O) if and only if O has S because S Ts. "Because S Ts" is then fleshed out with an historical causal story to the effect that Os with Ss T-ed better than those that did not and so were preserved in the population of Os. Many have thought that appeal to such a biologically-inspired notion of proper functions can distinguish how an item ought to behave from the way that it in fact behaves, that it can vouchsafe the existence of higher level kinds in spite of the fact that there is no common occurent realizer that all of them share, and that it can provide a criterion for individuating biological kinds (what makes a heart a heart is that it was selected for its pumping).

Again, this is an important starting point for many scientists and philosophers, and so we see no reason to build into the very notion of realization a commitment that metaphysically excludes historical realizers. Thus, we do not think that R_3 should be viewed as a constraint on realization. It seems best to treat the acceptability of historical realization (in discussions of function, or information, or

computing) on a case by case basis. For understanding higher level causation, this might be a problem, since it certainly seems that only occurent properties contribute to causal powers: the same causal powers could be achieved by any number of histories. But issues of this sort require more discussion than space affords at this time. To the extent that different scientific and philosophical discussions turn on allowing the possibility of historical realization, we think it would be imprudent to rule out historical realization from the start.

5 AUTONOMY, REDUCTION, AND MULTIPLE REALIZATION

We have already noted that, from the get-go, philosophers of mind have happily endorsed the idea that mental states are multiply realizable in different physical states. Fictitious Martians, fancied artificial intelligences, and presumed cross-species physical variation have all been invoked in support of this idea. To many, the multiple realizability of the mental provided the basis for thinking that strict mind-brain identity theories were mistaken, and that non-reductionist, functionalist views of the mind constituted a more compelling account of the relationship between the mental and the physical. Along with this metaphysics came the methodological moral that psychology and the cognitive sciences more generally are *autonomous* of the nitty-gritty details of basic neuroscience, such that we could do much (if not all) of our psychological theory construction without being hostage to developments within neuroscience.

This sort of view, which predominated throughout much of the 1970s and 1980s, came under attack toward the end of that period. There were two prongs to the critique, which together have suggested to many that the complacency about multiple realization, and its putative implications, cannot be justified.

The first was delivered by Jaegwon Kim in his attack on "the myth of non-reductive materialism". In a series of influential papers [1989; 1992; 1993], Kim argued that there were metaphysical reasons for thinking that non-reductive forms of materialism, including functionalism, were unstable hybrids. By appealing to causal powers, the nature of scientific explanation, natural kinds, and laws, Kim argued for the reductive materialism that he had begun his career defending, and against the autonomy of psychology.

To give the flavor of Kim's discussion and to see how it manifests the metaphysical strand to views of realization, consider one principle that he appeals to, what he calls the *principle of causal inheritance*: "if M is instantiated on a given occasion by being realized by P, then the causal powers of *this instance of* M are identical with (perhaps, a subset of) the causal powers of P" [Kim, 1993, 355]. The intuitive idea behind the principle is that instances of higher-order kinds, such as psychological kinds, inherit the causal powers of instances of the lower-order kinds that realize them. The plausibility of this intuition turns on the instances of each of these kinds being the very same glob of matter in the world, or close enough. And if the causal powers of instances are inherited from the physical to the mental, then it seems that all of the real causal action is captured by the lower

order description in terms of physical properties and powers. Thus, it is confused to think of mental causation as somehow distinct from or autonomous of physical causation.

The second prong to the attack on appeals to multiple realization derived from advances in the neurosciences, particularly from the emergence of *cognitive neuroscience* as a distinct subfield within the cognitive sciences. Developing as an area bridging cognitive psychology and neuroscience over the last twenty years, cognitive neuroscience has come to occupy a prominent place within the cognitive sciences. Perhaps this place will become more prominent still. In concluding their introduction to the hundred or so articles on neuroscience in *The MIT Encyclopedia of the Cognitive Sciences* [1999], Tom Albright and Helen Neville make a telling (even if tongue-in-cheek) prediction. They say that "if cognitive neuroscience fulfills its grand promise, later editions of this volume may contain a section on history, into which all of the nonneuro cognitive sciences discussion will be swept" [1999, lxix].

Hyperbole and rhetorical flourish aside, the rise of cognitive neuroscience has provided reason to rethink the phenomenon of multiple realization and what has been claimed about cognition in its name. As Bechtel and Mundale [1999] have argued, appeals to function are an intrinsic part of neurotaxonomy, and such appeals are often justified by cross-species investigations. Although the taxonomy of the brain into 47 regions by Korbinian Brodmann early in the 20^{th}-century had an anatomical basis, as Brodmann says, its "ultimate goal was the advancement of a theory of function and its pathological deviations" [1909/1994, 243], as quoted by [Bechtel and Mundale, 1999, 180].

Likewise, deficit studies in both human and non-human animals, a major source of information about psychological functioning, presupposes that the same part of the brain plays much the same physical role across different individuals. Hence the contrast between "functional", psychological levels and "physical", neural levels, and the idea that cross-species investigations and generalizations provide the basis for viewing the neural level as too fine-grained to capture genuinely psychological commonalities, both presuppose a view of the neurosciences that is sadly out of touch with the taxonomic and explanatory practices in those sciences.

Moreover, a closer look at neurotaxonomic practices raises doubts about the very terms in which the multiple realizability thesis has been stated. Philosophers typically talk of mental states as being realized in brain states. As Bechtel and Mundale also note, "the notion of a brain state is a philosopher's fiction." [1999, 177]. Cognitive scientists themselves are often interested in activities, functions, and mechanisms, not states *per se*. While brain states, whatever those are, may be multiply realizable, this will not establish that the functions and activities of the brain are. What is needed is a notion of realization, such as that discussed at the end of Section 3, that will accommodate cases of mechanistic realization.

These points — the mischaracterization of neurotaxonomy, the false dichotomy between functional and physical levels, and departure that talk of states makes from empirically-grounded practice — together suggest several ways in which tra-

ditional debates and positions which invoke the notion of multiple realization need to be rethought in light of attention to work in the neurosciences.

Larry Shapiro [2000; 2004] has also argued that both traditional and more recent putative examples of multiple realization fail because they do not constitute examples where *the same* cognitive function is realized by *relevantly different* physical mechanisms. For example, consider common appeals to the idea that silicon chips could replace neurons, one by one, to produce a silicon brain, without changing the psychological functions realized. Such examples provide support for the multiple realizability of cognition only if the *ways* in which neurons and silicon chips realize psychological functions are relevantly different. As Shapiro says, if "each neuron's contribution to psychological capacities is solely its transmission of an electrical signal, and if silicon chips contribute to psychological capacities in precisely the same way, then the silicon brain and the neural brain are not distinct realizations of the mind" [2000, 645]. At best, this would be an example of trivial multiple realization.

Shapiro's general point can be applied to more detailed examples that draw on recent work in cognitive neuroscience on neural plasticity (see [Shapiro, 2004, ch.2]). Ned Block and Jerry Fodor had appealed early on to the lability of the brain, as evidenced in the differential realization of the capacity for language in the left and right hemispheres in some cases, to defend the multiple realizability of psychological functions [Block and Fodor, 1972]. The flexibility and adaptability of both organism-level psychological functions, such as vision or memory and the determinate forms they take, have long been a subject of investigation within psychology, and much of this work has been viewed more recently under the rubric of "neural plasticity" [Gilbert, 1999]. There has been a surge of work on varieties of neural plasticity, including experience-dependent neural reprogramming [Elbert *et al.*, 1995], neural adaptation following cortical lesions [Kaas, 2000], and cross-modal neural rewiring [von Melchner *et al.*, 2000]. Consider this last bit of work in more detail.

Mriganka Sur and his colleagues have explored lesion-induced neural rewiring in the auditory system of young ferrets. In this paradigm, a lesion was induced in a part of the auditory system upstream from the auditory cortex. This involved severing the normal connection between the inferior and superior colliculus in the midbrain and the medial geniculate nucleus (MGN) in the left hemisphere of the ferrets. The effect of the procedure was that, over time, the MGN then comes to accept projections from the retina, passing on this information in turn to the auditory cortex. This allows the rewired ferrets to "see with their auditory systems". Could this be a case where cognitive neuroscience has shown that psychological functions are multiply realizable?

Like the imagined cases that are the stock and trade of traditional philosophical analysis, such cases must also jointly satisfy two conditions: they must describe the same psychological capacity and it must be realized in a relevantly different manner. Sur's work on cross-modal neural rewiring, interesting as it is, appears to satisfy neither of these conditions. Although rewired ferrets have some visual

capacity, they clearly do not have the same visual capacity as normal ferrets. But perhaps we need to focus not on the coarse-grained capacity of vision but on specific visual abilities, such as the capacity to detect light and to distinguish grates of light and dark. While even here there remain differences between normal and rewired ferrets, they are interestingly similar in their abilities, and perhaps we can grant that they do share some basic visual capacities.

This brings us to the second condition: does the auditory cortex of the rewired ferrets, for example, process visual information in a relevantly different manner from the visual cortex of normal ferrets? We have a case of multiple realization only if the answer to this question is "yes". As Sur and his colleagues point out, however, although this question remains somewhat open, the organization of the auditory cortex in rewired ferrets is strikingly similar to that of the visual cortex in normal ferrets, and strikingly unlike that of the auditory cortex in normal ferrets. Sameness of structure does not, of course, imply sameness of function, but this at least casts some doubt on whether this example would satisfy Shapiro's second condition even if it were to satisfy the first. As Shapiro says, it "seems not unreasonable to suppose that a rewired ferret sees as well as it does only to the extent that its auditory cortex resembles a visual cortex" [2004, 64]. If this supposition is not just reasonable but correct, then there is little support for the idea of multiple realizability to be garnered from at least this example of cross-modal neural rewiring.

Our view is that further consideration of the varieties of neural plasticity, and of the details of psychological and neural taxonomies, is very unlikely to support the actual multiple realizability of the psychological in the neural. Both Shapiro and Bechtel and Mundale have provided reasons that we think are the basis for skepticism here.

The two basic constraints that Shapiro identifies — same psychological function but (relevantly) different neural realization — pull in opposite directions. This is in large part because "the psychological" and "the neural" should not be viewed simply as two levels. Instead, as Bill Lycan [1987, ch. 4)]and Carl Craver [2001] have suggested, there are many interstitial levels between lowest and highest. Further, the relationship between them is complex enough that we should be wary of untutored level-talk and might even question the very utility of the levels metaphor.

Bechtel and Mundale [1999, section 5] have argued that one reason why multiple realizability seems so obviously to hold of the relation between the mental and the physical is the tacit adoption of a sort of double standard in thinking about the two. In thinking about psychological capacities, it is common to describe them coarsely — as the capacity for vision, or for short-term memory. By contrast, realizing neural structures and their "immediate" functions receive comparatively fine-grained description. For example, it is common to consider cases in which pain is the psychological state under consideration, while the physical realizers are described in terms of the different kinds of brains that, say, human beings and octopi have. Given that a coarse grain of description facilitates the view that

two cases are instances of the same (psychological) kind, and that a fine grain of description does the same for the view that two cases are instances of different (neurological) kinds, the bias that this double-standard introduces is one that creates the impression that such cases satisfy both of Shapiro's constraints.

Finally, our discussion of neural plasticity, brief as it is, raises the issue of what has been called *emergent realization* [Wilson, 2001], i.e., of a physical realizer for a given psychological capacity that could realize some other capacity were the world different in various ways. The idea of emergent realization is sometimes introduced in contrast to that of multiple realization, as a case where rather than there being a one-many relation between the psychological and the physical, there is a many-one relation. Metaphysicians of mind, in at least tacitly accepting the idea that realizers are metaphysically sufficient for the properties, states, and capacities that they realize, typically reject the possibility of emergent realization as incompatible with versions of physicalism that hold that the physical facts must, in some sense, determine all the facts there are.

The examples of neural plasticity that we have mentioned might be thought to give us pause here. For these appear to be cases in which once the world is changed in various ways — certain lesions are made, neural rewiring develops over time — one part of the brain comes to realize psychological capacities that it previously did not realize. But our preceding discussion should make us wary of any quick inferences here, and remind us to consider the empirical details more fully. What happens in these cases, after all, is that parts of the brain come to be not only differently configured in terms of how they connect to other parts of the brain, but also differently *structured* themselves. In the ferret brain, for example, the neural rewiring introduced by severing the auditory input channels to the MGN makes the physical structure of the auditory cortex, downstream, much more like the visual cortex of normal ferrets, bearing a columnar structure that is orientation-sensitive as well as a two-dimensional map of the surrounding visual space.

6 CONCLUSION: GETTING OUR HEADS TOGETHER

We began with the idea that the concept of realization was the servant of two masters. The idea of realization is at the heart of a number of debates in contemporary metaphysics and the philosophy of mind. It is also (at least implicitly) central to the explanatory and investigative practices of cognitive scientists. Here we have tried to review possible lines of rapprochement. We think that a notion of realization will continue to have an important role to play in systematically thinking about the nature of cognition, and we have tried to accommodate the demands placed on the concept by both philosophers and cognitive scientists. While we have tended to adopt an ecumenical view of how we should think about realization, we have also tried to identify a series of issues that remain contentious and that invite further discussion. Some of these concern what properties the relation of realization has, or what constraints it should be subject to — is it decompositional, intrinsic, or synchronic in nature? (section 4) — while others focus more

directly on the interplay between abstract philosophy concerns (say, about the autonomy of psychology) and ongoing explanatory practice in cognitive neuroscience (section 5). Both, we reckon, require that metaphysicians and scientists continue to get their heads together.

ACKNOWLEDGEMENTS

The authors would like to acknowledge the Society for Philosophy and Psychology for providing a venue (their 2001 meeting in Cincinnati) for the start of their collaboration. Rob Wilson's work on this paper has been supported by SSHRC Individual Grant 410-2001-0061 from the Social Sciences and Humanities Research Council of Canada.

BIBLIOGRAPHY

[Albright and Neville, 1999] T. Albright and H. Neville. Neurosciences, in R.A. Wilson and F.C. Keil (editors), *The MIT Encyclopedia of the Cognitive Sciences*. Cambridge, MA: MIT Press, pp.li-lxxii, 1999.

[Bechtel and Mundale, 1999] W. Bechtel and J. Mundale. Multiple Realizability Revisited: Linking Cognitive and Neural States, *Philosophy of Science* 66: 175-207, 1999.

[Block and Fodor, 1972] N. Block and J.A. Fodor. What Psychological States are Not, 1972. Reprinted in J.A. Fodor, *Representations*. Sussex, England: Harvester, 1981.

[Brodmann, 1909] K. Brodmann. *Vergleichende Lokalisationslehre der Grosshirnrinde*. Leipzig: J.A. Barth, 1909. English translation by L.J. Garey, *Brodmann's 'Localisation in the Cerebral Cortex'*. Cambridge, MA: MIT Press, 1994.

[Craver, 2001] C. F. Craver. Role Functions, Mechanisms and Hierarchy, *Philosophy of Science* 68: 31-55, 2001.

[Elbert *et al.*, 1995] T. Elbert, C. Pantev, C. Wienbruch, B. Rockstroh, and E. Taub. Increased use of the left hand in string players associated with increased cortical representation of the fingers *Science* 270: 305-307, 1995.

[Gilbert, 1999] C. Gilbert. Neural Plasticity, in R.A. Wilson and F.C. Keil (editors), *The MIT Encyclopedia of the Cognitive Sciences*. Cambridge, MA: MIT Press, pp.598-601, 1999.

[Gillett, 2002] C. Gillett. The dimensions view of realization: a critique of the Standard view *Analysis* 64: 316-23, 2002.

[Kaas, 2000] J. Kaas. The Reorganization of Sensory and Motor Maps after Injury in Adult Mammals, in M. Gazzaniga (editor), *The New Cognitive Neurosciences*, 2^{nd} edition. Cambridge, MA: MIT Press, pp.223-236, 2000.

[Kandel and Schwartz, 2000] E. R. Kandel and J. H. Schwartz, eds. *Principles of Neural Science*. 4^{th} Edition. New York: McGraw Hill, 2000.

[Kim, 1989] J. Kim. The Myth of Nonreductive Materialism, 1989. Reprinted in his *Supervenience and Mind*. New York: Cambridge University Press, 1993.

[Kim, 1992] J. Kim. Multiple Realization and the Metaphysics of Reduction, 1992. Reprinted in his *Supervenience and Mind*. New York: Cambridge University Press, 1993.

[Kim, 1993] J. Kim. *Supervenience and Mind*. New York: Cambridge University Press, 1993.

[Kim, 1996] J. Kim. *Philosophy of Mind*. Cambridge, MA: MIT Press, 1996.

[Kim, 1998] J. Kim. *Mind in a Physical World*. Cambridge, MA: MIT Press, 1998.

[Lycan, 1990] W. G. Lycan. The Continuity of Levels of Nature, in W. G. Lycan (editor), *Mind and Cognition: A Reader*. Oxford: Bradford, 1990.

[Machamer *et al.*, 2000] P. K. Machamer, L. Darden and C. F. Craver. Thinking about Mechanisms" *Philosophy of Science* 67: 1-25, 2000.

[Marr, 1982] D. Marr. *Vision: A Computational Investigation into the Human Representation and Processing of Visual Information*. San Francisco, CA: Freeman, 1982.

[Millikan, 1993] R. G. Millikan. *White Queen Psychology and Other Essays for Alice*. Cambridge, MA: MIT Press, 1993.
[Neander, 1991] K. Neander. Functions as Selected Effects, *Philosophy of Science* 58: 168-184, 1991.
[Place, 1956] U. T. Place. Is Consciousness a Brain Process?, *The British Journal of Psychology* 47: 44-50, 1956.
[Polger, 2004] T. Polger. Neural Machinery and Realization, *Philosophy of Science* 71: 997-1006, 2004.
[Putnam, 1960] H. Putnam. Minds and Machines, 1960. Reprinted in his *Mind, Language, and Reality: Philosophical Papers, Volume 2*. New York: Cambridge University Press, 1975.
[Shapiro, 2000] L. Shapiro. Multiple Realizations, *Journal of Philosophy* 97: 635-654, 2000.
[Shapiro, 2004] L. Shapiro. *The Mind Incarnate*. Cambridge, MA: MIT Press, 2004.
[Smart, 1959] J. J. C. Smart. Sensations and Brain Processes, 1959. Reprinted in D. Rosenthal (editor), *The Nature of Mind*. New York, Oxford University Press, 1991.
[Sober, 1985] E. Sober. Putting the Function Back into Functionalism, 1985. Reprinted in W.G. Lycan (editor), *Mind and Cognition: An Anthology*. Oxford: Blackwell, 1998, 2^{nd} edition.
[von Melchner *et al.*, 2000] L. von Melchner, S. Pallas, and M. Sur. Visual Behaviour Mediated by Auditory Cortex Directed to the Auditory Pathway, *Nature* 404: 871-876, 2000.
[Wilson, 2001] R. A. Wilson. Two Views of Realization, *Philosophical Studies* 104: 1-31, 2001.
[Wilson, 2004] R. A. Wilson. *Boundaries of the Mind: The Individual in the Fragile Sciences: Cognition*. New York: Cambridge University Press, 2004.
[Wilson and Keil, 1998] R. A. Wilson and F. C. Keil. The Shadows and Shallows of Explanation, 1998. Modified version reprinted in F.C. Keil and R.A. Wilson (editors), *Explanation and Cognition*. Cambridge, MA: MIT Press, 2000.
[Wilson and Keil, 1999] R. A. Wilson and F. C. Keil, eds. *The MIT Encyclopedia of the Cognitive Sciences*. Cambridge, MA: MIT Press, 1999.
[Wimsatt, 1997] W. Wimsatt. Aggregativity: Reductive Heuristics for Finding Emergence *Philosophy of Science* 64: S372-S384, 1997.

REDUCTION: MODELS OF CROSS-SCIENTIFIC RELATIONS AND THEIR IMPLICATIONS FOR THE PSYCHOLOGY-NEUROSCIENCE INTERFACE

Robert N. McCauley

1 INTRODUCTION

With the rise of functionalism and of new proposals about putative intrinsic properties of the physical necessary to account for conscious experience, enthusiasm among philosophers about the possibilities of reducing the theories of psychology to those of neuroscience probably reached its nadir in the 1990s. Ned Block [1997] noted that recent philosophy of psychology has witnessed (very nearly) an "antireductionist consensus," and Jaegwon Kim suggested that this consensus extended well beyond the confines of professional philosophy:

> ... "reduction," "reductionism," "reductionist theory," and "reductionist explanation" have become pejoratives not only in philosophy, on both sides of the Atlantic, but also in the general intellectual culture of today. They have become common epithets thrown at one's critical targets to tarnish them with intellectual naivete and backwardness. ... If you want to be politically correct in philosophical matters, you would not dare come anywhere near reductionism, nor a reductionist [1998, p. 89].

The widespread presumption is that proposed reductions of the psychological to the neural are so obviously hopeless that when they address the deepest philosophical problems of mind philosophers need not trouble themselves much with the details of either scientific investigations pertaining to their connections or philosophical accounts of those investigations.

On three prominent fronts these developments seem (at the very least) unexpected from the standpoint of the philosophy of science. First and, perhaps, most important, the history of science provides no grounds for such pessimism (let alone, such a dismissive view) about the prospects for successful reductions in these environs. The explanatory triumphs of the resulting theoretical integrations have richly rewarded the eagerness with which scientists have pursued reductive

Handbook of the Philosophy of Science. Philosophy of Psychology and Cognitive Science
Volume editor: Paul Thagard
General editors: Dov M. Gabbay, Paul Thagard and John Woods
© 2007 Elsevier B.V. All rights reserved

projects over the past one hundred fifty years. Reduction has probably been the single most effective research strategy in the history of modern science, engendering more precise accounts of the mechanisms (and their operations) underlying everything from magnetic forces to organisms' inheritance of traits to the visual perception of moving objects — to note but three examples from three different levels of analysis in science and three different collections of decades in the two centuries in question. Exploring reductive possibilities opens new avenues for sharing methodological, theoretical, and evidential resources. Successful reductions reliably generate productive programs of research at the analytical levels from which the candidate theories hail, squaring the lower level, mechanical details with the upper level phenomenal patterns and refining our understanding of both in the bargain.

Second, the opportunities for such theoretical integrations between the cognitive and the neural sciences have, if anything, only increased during the time period when this anti-reductionist consensus has prevailed among philosophers. The emergence in the 1960s and 1970s of the multi-disciplinary project that has become cognitive science provided example after example of researchers attempting to mine evidence from work in related disciplines and to develop cross-scientific resonances of theory. The growing use of brain-imaging technologies, especially PET and fMRI, over the next two decades has only accelerated the pace at which the multi-level theoretical proposals that have marked recent cognitive neuroscience have appeared. These proposals regularly conjoin insights from psychology, network modeling, clinical neurology, cellular neuroscience, and more. The past forty years have witnessed not only a rapid multiplication in the sheer number of findings about the mind-brain available for reductive analysis but dozens of interdisciplinary projects in the relevant sciences that have made reductive headway.

Finally, philosophers of science have diligently sought at least since the 1950s to provide general models of reductive accomplishments in science ([Kemeny and Oppenheim, 1956]; [Oppenheim and Putnam, 1958]). In the final decade of logical empiricism's reign, Ernest Nagel's [1961/1979] account of reductive relations in science (referred to hereafter as the "standard model") emerged as the touchstone for subsequent discussions among philosophers of mind. There Nagel speaks about both the reduction of scientific theories and the reduction of whole sciences. He construes both in terms of the logical derivation of a reduced theory's laws from the laws of a reducing theory, supplemented by bridge principles that specify systematic connections between the two theories' predicates and the boundary conditions within which those connections hold. The anti-reductionist consensus among both functionalists and the friends of consciousness has mostly turned on arguments that the considerations they raise establish (different) barriers to reduction — as characterized by Nagel [Bickle, 1998, p. 5].

It probably comes as no surprise, though, that philosophers of science, including those interested in the psychological and neural sciences, have not ceased to advance new models of scientific reduction and of cross-scientific relations, more generally. These models differ variously and, sometimes, considerably from Nagel's

account. Virtually all of them include provisions that disarm the principal arguments on which the anti-reductionist consensus rests. Although they differ both in detail and in the strength of the reductionism they defend, they broadly concur that reports of the death of positions that explain mentality on the basis of neural operations and that identify features of minds and brains have been greatly exaggerated and that the character of conscious experience does not constitute an insuperable obstacle to proposing such hypothetical identities.

Section 2 briefly sketches Nagel's standard model of reduction and then discusses how the machinery of the "New Wave" model of reduction has transformed one of the standard model's principal problems into a virtue. Section 3 explicates the New Wave continuum of comparative goodness of intertheoretic mapping. Section 4 situates within the New Wave framework the two major arguments informing the anti-reductionist consensus among recent philosophers of mind. These arguments concern the multiple realizability of mental states and the irreducibility of conscious experience. Sections 5 and 6 review criticisms of the New Wave model suggesting that its proximity to the logical empiricist model on two fronts renders it, first, insufficiently sensitive to the wide range of cross-scientific relations that arise and, second, capable of engendering misleading conclusions about the status and fate of the cognitive and psychological sciences relative to neuroscience. Section 7 presents a more fine-grained model of intertheoretic relations that distinguishes between two major sorts of cases that the New Wave models lump together. Coincident with work on mechanistic explanation in science (discussed in section 9), this alternative analysis contrasts two sorts of cases that exhibit diverging profiles and considers New Wave counter-arguments against distinguishing them. It elucidates the explanatory pluralism that dominates in cross-scientific settings. Section 8 suggests that, although disagreements about general models of scientific reduction persist, confluences of opinion have emerged over the last few years in the works of philosophers interested in exploring fruitful cross-scientific relations at the borders between the cognitive and neural sciences. The first concerns the distance of functionalists' multiple realizability argument from the practices and discoveries of working scientists in these fields. Section 9 takes up a second sort of confluence concerning the crucial role that mechanistic analyses play in those practices and discoveries. Recent mechanistic analyses apply the morals of explanatory pluralism to models for the detailed study of particular patterns and mechanisms in nature. These positions rule out the most ambitious aims of New Wave reductionists in interlevel settings. Finally, section 10 examines how defusing the multiple realizability argument and taking a closer look at the practices and discoveries of scientists working in these areas suggests that arguments for the unique character of conscious experience are largely irrelevant to the sorts of considerations that lead scientists to hypothesize intertheoretic identities in reductive contexts. The Heuristic Identity Theory incorporates these insights about cross-scientific research, advancing a new, more scientifically informed, version of the psycho-physical identity theory.

2 THE STANDARD MODEL OF REDUCTION AND NEW WAVE REVISIONISM

Scientific reduction, according to Nagel and the logical empiricists, is a deduction of the laws of one scientific theory (the reduced theory) from those of another (the reducing theory). This inference requires supplementing the laws of the reducing theory with a set of ancillary statements that lay out systematic connections between the two theories' predicates while incorporating the boundary conditions within which those connections are realized. (See Wilson and Craver's "Realization" in this volume.) On this account reductions are a type of explanation in which, unlike most cases of scientific explanation, the explanandum is *not* some phenomenon but rather some law or other of the theory that is being reduced. A successful reduction on this standard model demonstrates how the reducing theory's explanatory resources encompass those of the reduced theory that is to be mapped on to it. Thus, in effect, the reduced theory constitutes an application of the reducing theory in one of its sub-domains that the boundary conditions specify.

Philosophers have used a variety of phrases ("bridge principles," "reduction functions," "coordinating definitions," etc.) to refer to the ancillary statements that supplement the laws of the reducing theory, and they have offered various proposals about the connections those statements should establish between the two theories' predicates. The significant point for now is that on the standard model those connections must enable the reduction to meet two important constraints. The first constraint is logical; the second is material.

The first constraint is concerned with assuring the "derivability" of the reduced theory from the reducing theory. In order for the reducing theory to explain the reduced theory, the latter must follow from the former (supplemented by the bridge principles) as a deductive consequence. That is because explanation for the logical empiricists conforms to the deductive-nomological (D-N) model. On the D-N model explanation involves the *derivation* of statements about what is to be explained from scientific *laws*. Consequently, the reduction functions have to articulate connections between the two theories' predicates of sufficient logical strength to support the derivation.

The second constraint is the "connectability" condition. In order for the reduction to help justify a metaphysical unity as well as a theoretical unity to science, the reduction functions also have to certify substantial connections between the entities and their properties that the two theories discuss. Establishing such connections between scientific theories motivates *programs* for unifying science via "microreductions," in which the entities the reducing theory discusses constitute the components of the entities that the reduced theory endorses ([Oppenheim and Putnam, 1958]; [Causey, 1977]). Such programs not only aim to fashion a compelling case based on mereological relations for a materialist metaphysics but also envision the reduction of entire sciences. They foresee, at least in principle, the possibility of scientists eventually abandoning research at higher levels of analysis

in deference to explanatory theories at lower levels that are simultaneously more comprehensive and more detailed.

Robert Causey's [1977] proposal for a theoretical unification of the sciences is, perhaps, the most thoroughgoing. Paul M. Churchland [1979] and Patricia S. Churchland [1986] have jointly initiated one of the best known programs exploring the possibilities for such deflationary consequences, while John Bickle [1998; 2003] offers some of the more spirited contributions along these lines recently. Bickle, for example, is concerned with "the reduction of ... psychology to neuroscience," and he holds that "reduction is a proof of displacement (in principle), showing that a typically more comprehensive theory contains explanatory and predictive resources that parallel those of the reduced theory" [1998, 214 and 28].

Proposals differ about the logical and material strength of such intertheoretic connections. The major options are (1) that the various predicates of the reducing theory constitute sufficient conditions for predicates in the reduced theory, (2) that they constitute necessary and sufficient conditions, or (3) that they not only constitute necessary and sufficient conditions but that they also involve intertheoretic identities. (See [Nagel, 1961/1979, 354–355]; [Causey, 1977], respectively.) The comprehensive mapping of the predicates applicable to the entities that the reduced theory countenances on to the predicates applicable to the entities that the reducing theory countenances vindicates assumptions about correspondences between the two theories' ontologies.

All three of these options possess sufficient logical muscle to underwrite the derivation that the first constraint demands, so it is primarily the problem of ascertaining what is required for adequately linking the theories' ontologies, i.e., what is required for meeting the second constraint, that has occupied subsequent commentators. How strong of an intertheoretic "mapping" relation is required to variously achieve (a) the explanation of the reduced theory? (b) the displacement of the reduced theory? and (c) the theoretical and ontological unification of science? On all three of options (1) - (3) above, the reduction functions can be regarded, in effect, as hypotheses that call for empirical support, and, correspondingly, discussions have regularly taken up the character of that support, how advocates of any particular reductive explanation might gather it, and what it would entail concerning questions (a) - (c) above.

The appeal of the standard model's formality, clarity, and precision is uncontested. Philosophers, however, began to realize that its idealized account of intertheoretic relations came at the price of its ability to capture a large range of actual cases of intertheoretic relations that did not meet its exacting requirements [Wimsatt, 1978]. Faithful portrayals of many relationships between scientific theories often resulted in candidate reduction functions that were weaker logically than the options listed above and that supported partial mappings only. Therefore, the resulting connections frequently seemed capable of sustaining neither the derivation of the theory to be reduced nor the comprehensive mapping of its ontology on to that of the putative reducing theory. Compare, for example, Patricia Churchland's diverging assessments in the 1980s of the prospects for the reduction of

various aspects of consciousness [P. S. Churchland, 1983; 1988]. Without a doubt, the most celebrated analysis of such failures of mapping was the Churchlands' profound pessimism concerning the possible connections between our everyday folk psychology and theories in neuroscience ([P. S. Churchland, 1986]; [P. M. Churchland, 1989]).

Although this diagnosis does not tally well with the logical empiricists' conceptions of reduction and the unity of science, it is consonant with the persisting impression in many cases that the reducing theory's resources do not merely encompass those of the reduced theory. If, for no other reason, on the basis of its added precision alone, the reducing theory usually appears to *improve* upon the reduced theory's account of things. Not infrequently, it corrects it. Even the familiar case of the reduction of the classical gas laws yields corrections to their predictions at extreme temperatures and pressures. Or within cognitive neuroscience itself, David van Essen and Jack Gallant's [1994] more articulated picture of the numerous connections permitting the sharing of information in the processing streams of the primate visual systems' "what" and "where" pathways, arguably, constitutes a correction of the initial proposal of Leslie Ungerleider and Mortimer Mishkin, which construed these sub-systems' operations as basically isolated from one another ([Ungerleider and Mishkin, 1982]; [Mishkin, Ungerleider, and Macko, 1983]).

If reducing theories often *correct* reduced theories in intertheoretic reductions, then, on the standard model of reduction, the reduced theories' laws should *not* follow with deductive validity from premises about the laws of the reducing theories in conjunction with the bridge principles. In dealing with some of the most impressive reductions in the history of science, advocates of the standard model find themselves faced with the embarrassing dilemma of having to repudiate the D-N model of explanation unless they will accept reduction functions that leave enough semantic slack to render the putative derivation guilty of equivocation. After all, false reduced theories cannot be validly deduced from true reducing theories, and they cannot even appear to do so unless the argument involves an equivocation. (See [Wimsatt, 1976, 218]; [P. M. Churchland, 1989, 48].)

The next generation of philosophers interested in modeling scientific reductions came to regard our *inability* to sustain reduction functions without semantic slack, i.e., our inability to formulate defensible reduction functions capable of underwriting the derivation of the reduced theory's regularities, as a *virtue* of any putative reduction that improves upon those regularities. Instead of standing by a formally perspicuous, idealized model of intertheoretic reduction that fails to describe many cases accurately, the successors of the standard model allow for the relaxation of its requirements. For example, Kenneth Schaffner [1967] argued that strictly speaking, what can be deduced from the reducing theory is not the reduced theory itself but only an *analogue* of that theory.

This proposal inspired what Bickle [1998] has dubbed the "New Wave" model of intertheoretic reduction. (Although the accounts of scientific reduction and cross-scientific relations that Bickle, the Churchlands, and Clifford Hooker propound do not coincide in every last detail, they are sufficiently similar on the fronts that

matter here that for ease of exposition I will use Bickle's "New Wave" label.) On this "New Wave" model the reducing theory does not explain the reduced theory, so the dilemma disappears. Instead, it explains an analogue of the reduced theory constructed within the conceptual framework of the reducing theory. Thus, on the New Wave view the analogical relationship between the reduced theory and its reconstruction in terms of the reducing theory's conceptual resources enables the reducing theory both to correct the reduced theory and to explain at least something very much like it at the same time. Moreover, relying on analogy, the New Wave model of reduction, apparently[1], accomplishes all of this without needing to specify bridge principles, and, therefore, without needing to explicate either their logical or ontological status. Hooker [1981] was the first to explore this proposal at length and to provide a formal explication. He notes in the course of that exposition that the strength of the analogy can vary considerably from one case to another, resulting in a spectrum of analogical strength that ranges from retentive reduction at one end to outright theory replacement at the other.

The difficulties surrounding appeals to analogical reasoning, however, are familiar. Just how close does the analogy need to be in order to justify reductive claims and how is the "closeness" of an analogy to be measured in the first place? How well must the reduced theory map on to the reducing theory in order to establish their explanatory and ontological continuity? New Wave reductionists have offered various proposals that conform to Schaffner and Hooker's general approach. For example, Paul Churchland [1989, 49] suggests that the reducing theory should provide an "equipotent image" of the reduced theory. The *equal potency* concerns its explanatory and predictive capabilities. The equipotent analogue, formulated in terms of the reducing theory's conceptual resources, should explain and predict the phenomena that the reduced theory addresses. To constitute an *image* of the reduced theory, the analogue may not have to map it comprehensively, but it should preserve that theory's principal contours from the standpoint of the causal relations it systematizes. Similarly, Bickle, who prefers to explicate reductive relations within the framework of a non-sentential, structuralist account of theories, aims to show how the analogues "mimic the structure" of the reduced theories [1998, 65].

Although talk of either images or mimicry is unlikely to meet the derivability constraint of the standard model or to point toward conceptions of explanation that square with the D-N model, they do undergird a picture of *approximate reduction* that embraces the familiar cases and, at least, offers an initial, if not an especially precise, step toward ascertaining just how much slack is tolerable. Bickle is sensitive to the fact that the cost of the New Wave models' broader

[1]Ronald Endicott [1998] argues persuasively that New Wave reduction does *not* avoid the problem of formulating satisfactory reduction functions but only relocates it. Endicott maintains that neither Paul Churchland's [1989] nor Bickle's [1998] (different) non-sentential analyses of theories, finally, enable them to avoid the fact that scientific theories always involve at least some "public-language sentences" [1998, 71] and, thus, any account of theory reduction (including these New Wave accounts) that aims to draw ontological conclusions must face the problem of specifying "a set of intertheoretic bridge laws" [1998, 72].

applicability is their vagueness. So, he [1998] employs the formal machinery of the structuralist program to provide a means for calibrating the degrees with which the theory-analogues approximate the commitments of the reduced theories.

Bickle's structuralist account characterizes theories in terms of their models and their intended empirical applications. A theory's models are the real world and mathematical systems that possess the structures it describes. Its intended empirical applications are all of the actual systems in the world to which it applies, as specified by the relevant scientific community. Models consist of (1) base sets, whose elements are classified according to the theory's categories, (2) auxiliary sets, which are abstract spaces which the theory's explanations presume, and (3) fundamental relations and functions, i.e., operations on the elements of (1) and (2). The collection of some theory's fundamental relations and functions constitutes the structure of its models. On this specific version of the New Wave account, reduction is, in effect, the mapping of particular models and their intended applications across two theories. That mapping must satisfy a variety of conditions, but the significant point is that the itemized accounting that this model theoretic approach affords permits *a measure* of the goodness of intertheoretic mapping.

On the New Wave account, the standard model's ideal designates an end point on the continuum of the comparative levels of isomorphism between reduced theories and their analogues. The continuum orders the relative goodness-of-mapping relations possible between reduced theories and their images constructed within the frameworks of their corresponding reducing theories. None of the New Wave reductionists, though, offer any precise criteria for when the amount of slack is no longer tolerable, i.e., when the theory-analogue's approximation of the reduced theory becomes too loose to make sense of reductive talk. Bickle [1998, 100–101] readily notes this limitation. At some point on that continuum the goodness-of-mapping becomes sufficiently weak that the case for intertheoretic continuity collapses.

Ironically, both friends [Fodor, 1975] and foes [P. M. Churchland, 1989] of folk psychology agree that this is the character of its relationship with the theories of neuroscience. New Wave reductionists hold that such situations make not for theory reduction but for the "historical theory succession" that marks scientific revolutions [Bickle, 1998, 101]. The superior theory simply displaces its inferior counterpart. If their intertheoretic mappings are as tenuous as those in uncontroversial historical cases such as, say, those between Stahl's account of combustion and Lavoisier's or those between Gall's phrenological hypotheses and modern cognitive neuroscience, we are, presumably, justified in speaking of the complete *elimination* of the inferior theory. Of course, it appears that the theories that risk elimination in the case at hand are, at the very least, those of folk psychology and, presumably, those in other areas of psychology that appeal to similar notions.[2]

[2]Although section 8 below will maintain that Bickle's [2003] most recent extended discussion of these matters offers some grounds for situating his position alongside philosophical treatments that do not always foresee the elimination of psychology, it is still worth noting that he characterizes his "ruthless" reductionism as one in which all such failures to map basic entities of the

3 THE NEW WAVE CONTINUUM OF THE COMPARATIVE GOODNESS OF INTERTHEORETIC MAPPING

The New Wave model situates different cases of intertheoretic relations at various points on a continuum of comparative goodness of intertheoretic mapping. (See figure 1.)

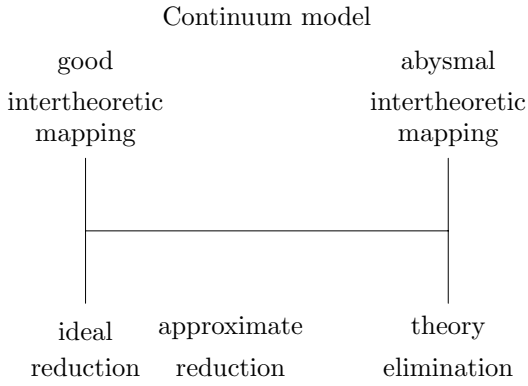

Figure 1.

These include cases situated quite near the end point of the continuum defined by the standard model's ideal, i.e., at the left end of the continuum that figure 1 portrays. Advocates of the standard model cited examples from basic physical science (such as the reduction of the wave theory of light to electromagnetic theory). Uncontroversial examples concerning the psychological and neural sciences may not exist, but, of a piece with William Wimsatt's [1978] observation, noted above, few, if any, actual cases of intertheoretic relations fully meet these exacting standards in *any* science.

New Wave theorists agree with Wimsatt's judgment. Hooker, for example, suspects that "the retention extreme of the retention/replacement continuum goes unoccupied" [1981, 45]. If even the standard model's parade cases from the physical sciences, in fact, fall at some distance on this continuum from the anchor point that designates that ideal, then that would only underscore the significance of New Wave analyses' abilities to make sense of these many familiar cases of approximate reduction. On the New Wave account the standard model's parade cases *are* only approximate reductions, since they reliably require minor counterfactual assumptions. (See [Bickle, 1998, especially p. 38 and 2003, p. 11].) Examples include the approximate reduction of the classical gas laws to principles of the kinetic theory and statistical mechanics and of Kepler's laws concerning planetary orbits

reduced theory result in their being "related *in a domain eliminating way*" to the machinery of the reducing theory [2003, 98]. What, apparently, has changed is not his analysis of contexts marked by mapping failures but, rather, Bickle's assessment of how much of the psychology-neuroscience interface that analysis captures.

to those of Newtonian mechanics. Clearly, if this is where the most thoroughgoing reductions from the physical sciences fall on the continuum in figure 1, then the consensus of philosophical opinion would locate *all* reductive proposals linking theories from the psychological and neural sciences even further away from the standard model's ideal. Even the best of the psychology-to-neuroscience cases would be comparatively *less* approximate reductions. Candidates from these domains would include the reduction of psychological proposals about the "switch" responsible for the consolidation of declarative memories in terms of the molecular mechanisms (both subcellular and extracellular) underlying both the transition from early phase (E-LTP) to late phase (L-LTP) long term potentiation and the preservation of the latter [Bickle, 2003, ch. 2].

As bases for constructing an analogue of the reduced theory dwindle, cases are arrayed further and further to the right on the continuum in figure 1. On the New Wave account the prospects for retaining either the principles or the ontology of the theory to be reduced decrease as cases exhibit fewer and fewer correspondences. Many of the classic revolutions in the history of science fall here. These include the elimination of the Aristotelian-Ptolemaic cosmology and the impetus theory with the rise, respectively, of the Copernican system and Galileo's investigations of terrestrial mechanics. New Wave reductionists, especially the Churchlands, are famous for their arguments that many cases of intertheoretic relations at the interface of psychology and neuroscience should be located at this end of the continuum. (See [P. M. Churchland, 1989, 1–22]; [P. S. Churchland, 1986, 373] as well as [Bickle, 1998, 30, figure 2.1].)

In the first sort of case (near the left end of the continuum in figure 1) the intertheoretic mapping is delightfully smooth, and the explanatory power of the reduction is transparent. In the second sort of case (falling, say, in the left half of the continuum), the analogies are close enough and the mappings remain substantial. Any improvements or corrections at the reduced theory's edges are a function of the heightened precision the reducing theory affords. Increasingly problematic cases make up the third category as they fall at greater and greater distances from the standard model's ideal (i.e., increasingly close to the right end of the continuum). Correspondingly, the intertheoretic mappings become ever more "bumpy" until, as they near the continuum's opposite end, they become prohibitively so. In this half of the continuum the outlook for reconciling the two theories moves from dim to dismal. New Wave reductionists insist that the failure of intertheoretic mapping in the dismal cases is so thoroughgoing that the success of the reducing theory impugns the integrity of the "reduced" theory and motivates its outright rejection. Ronald Endicott [1998, 57, footnote 13] says that referring to such cases as "bumpy reductions" is like referring to a divorce as a "bumpy marriage."

In these cases that fall at the displacement end of the continuum, the New Wave reductionists' presumptions in favor of neuroscientific over psychological theories has not gone uncontested. Not only have philosophers who are critical of reductionist programs objected, so have many philosophers sympathetic to reductionists' projects — though they have objected *on quite different grounds.* The next

section outlines the two most influential arguments that have arisen from the anti-reductionists. Sections 5, 6 and 7 explore the reservations of other philosophers who are less averse to reductionism, and Sections 8, 9, and 10 trace a confluence of both their and New Wave responses to the two principal arguments of the anti-reductionists that are sketched in section 4.

4 TWO ARGUMENTS FOR NON-REDUCTIVE MATERIALISM

Many philosophers, including many who profess a naturalistic orientation, subscribe to versions of non-reductive materialism that, ultimately, aim to absolve them of much need to scrutinize seriously either reductionist proposals within the psychological and neural sciences or philosophical discussions thereof (whether New Wave or other). These philosophers adopt these positions not because they reject all of the assumptions of New Wave reductionists but rather because they so heartily concur with one of them. Specifically, they agree with the New Wave reductionists' surmise that psychological theories will often show little promise for intertheoretic mappings on to the theories of neuroscience. These non-reductive positions marshal considerations that suggest that reductionists will not be able to map readily either some features of psychological theories or some features of their objects of study (i.e., minds) in to theories about brains in a fashion that will sustain any sort of displacement of the psychological. Their partisans regard one or both of those considerations, viz., the multiple realizability of psychological states or the peculiar character of conscious experience, as establishing barriers to reduction, certainly as classically construed. Instead of employing that premise (in the way that New Wave reductionists do) as promising grounds for *displacing* the psychological, they view it as reason to reject the assumption that such failures to find analogues must automatically impugn the psychological. They hold that, on the contrary, what these failures show is the *indispensability* of the psychological.

Non-reductive materialists come in various stripes but, finally, many take inspiration from what is, by now, a familiar argument concerning the multiple realizability of psychological states.[3] Hilary Putnam [1967/1975] first advanced this argument against psycho-physical identity theories. Putnam argued that the same psychological state can be realized by many different physical states. He appealed to the fact that many organisms other than humans experience pain, yet they have brains that differ considerably from the brains of *homo sapiens*. A physical

[3]Although the brief discussion that follows will address neither his anomalous monism nor his concern with the normative in interpretation, a careful reading of Donald Davidson's "Mental Events" [1970] will disclose assumptions about the range of possible relations between the psychological and the neuroscientific that press an extreme version of the functionalist argument concerning multiple realizability, which is to say that although Davidson avows that every mental event is a physical event, there are no systematic relations to be found among such pairs of event descriptions. Davidson holds that there are no psychological laws; hence there are no *theoretically substantiated* psychological kinds (thus, "multiple realizability" is a bit of a misnomer); hence there are no psycho-physical laws. Of course, Davidson published this paper just as psychology was beginning to free itself from the grip of behaviorism.

description of the arrangements that constitute the hunger of an octopus will almost certainly look quite different from the physical description of the functionally equivalent state in human beings. Therefore, pain, hunger, and, presumably, any of our other psychological states that we attribute to other organisms, cannot be identified with states of the human brain, in particular. It is easy to envision extending this conclusion. Even if we might settle on some state to identify with pain in all of the terrestrial creatures who experience it, that hypothesis would face the further challenge of having to serve as the state of affairs underlying the pains of *extra*terrestrial creatures too. The suggestion is that a wide variety of possible physical states in a wide variety of creatures might all be pain, i.e., that the psychological state of pain almost certainly has multiple physical realizations, and if that is so, then it precludes any simple mapping of this and other qualified psychological states on to the neural states of humans in the way that both psycho-physical identity theories and proposed reductions of psychological theories to theories in neuroscience would require.

The rise of cognitive science and especially of artificial intelligence inspired even more ambitious versions of the multiple realizability argument. As computers proved capable of a growing list of accomplishments from proving theorems to playing chess — accomplishments that we unequivocally regard as intellectual when *we* perform them — it appeared that the chauvinism of neural reductionism and identity theories ran even deeper. Alternative realizations of *bona fide* psychological states were not confined to other critters (including critters from other planets). Completely different forms of hardware could, perhaps, instantiate those psychological states too.

Enter functionalism. Functionalists in the philosophy of mind proposed that the best way to make sense of such a diversity of physical circumstances, all of which could be psychological states, was not to worry about describing these systems physically (since their diversity seemed to guarantee that what was of interest in common about them could not be captured by laws concerning their physical constitutions). They proposed, instead, to characterize psychological states *functionally*. Like most interesting philosophical positions, functionalism, too, comes in many flavors, but it gains traction in debates about reduction when it elevates this thesis to the level of a metaphysical claim [Polger, 2004]. Metaphysical functionalism maintains that as functional states, mental states should be delineated in terms of their causal interactions with one another, with input from the senses, and with motor outputs. On this view, a mental state is the nexus of such causal relations. It is the functionally described operations of a system, rather than any part of the system. It must be characterized at that level of abstractness in order to capture the diverse range of its possible physical instantiations. Employing such abstract accounts of mental states, functionalism, arguably, even allows for instantiations of mental states that are not physical [Fodor, 1981]. Such versions of the position would be non-reductive but examples of neither physicalism nor materialism.

Jerry Fodor [1974 and 1975] appealed to multiple realizability in his criticism of reductionist proposals in psychology. Multiple realizations of psychological states quickly yield reduction functions that are unwieldy and impractical at best. Instead of reduction functions establishing systematic connections between one type of psychological state and one type of neural state within some well-specified set of boundary conditions, the multiple realization of psychological states opens the door to reduction functions that might include immense disjunctions of neural possibilities, since any one of those states of affairs would suffice to instantiate the psychological item in question. Without mappings of psychological types on to neuroscientific types, the multiple realization of psychological states requires that the philosopher of psychology adopt no more than a "token physicalism" that affirms only that each token psychological state is identical with some token brain state. Each token of some psychological type is a token of some physical type, however, every token of that psychological type is not a token of *one* particular physical type. As the possible arrangements that might realize the psychological state increase, so will the size of the corresponding disjunctions.

Fodor contends that disjunctions of possible neural instantiations of some psychological state do not need to be very large before the bridge principles and the reductions that they are taken to inform become not merely unhelpful as guides to scientific research but downright misleading. As noted above, the New Wave reductionists agree, and they take the resulting fragmentation of the psychological categories at the neural level as a reason to expect at least the dismantling, if not the outright elimination, of the psychological account. Fodor, by contrast, stresses that such a displacement of the psychological would result in a science that is needlessly impoverished *from the standpoint of explanation.* It would disassemble perfectly good, readily applicable psychological principles and replace them with a plethora of physical accounts about the micro-level details of diverse systems that would sacrifice all sense of the psychological regularities they exhibit.

The last step of this argument does not turn on anything special about mentality. Fodor's argument suggests that the same morals could apply at any level of analysis in science. He offers an analogy with arrangements within a different level of analysis in science in the service of a *reductio ad absurdum* argument. He compares proposals to replace various concepts in psychological principles with disjunctive summaries of all of their (possible) physical instantiations with proposals to replace the concept 'money' in the principles of economics with disjunctive summaries of all of the instantiations that money has taken in human history. Fodor submits that the former reduction is as pointless and forlorn as the latter (which would truly render economics a dismal science). The absence of anything other than the most multifarious mappings of psychological onto physical states offers no promise of usefully preserving the explanatory achievements of psychology, whether of the folk or scientific varieties. Controversy about whether psychology of either sort, in fact, has many explanatory achievements to be preserved divides participants in these debates along predictable lines [P. M. Churchland, 1989, ch. 1]. Fodor, presumably, thinks that psychology possesses at least enough explana-

tory grip to have motivated this argument in the first place. Even if Fodor's comparative optimism on this count is unjustified and Robert Cummins [2000] is right that most of psychology's principles are not explanatory laws so much as "effects," i.e., patterns in need of explanation, a version of Fodor's argument would still stand. For, whatever the status of psychological principles, the myriad mappings Fodor forecasts would obscure the coherence of patterns in need of explanation no less than that of any putative laws capable of supporting explanations, and in either case the resulting science would have fewer points of empirical leverage rather than more.

The problems with mapping the massive disjunctions that the multiple realizability of psychological states requires do not end there. Those disjunctions also leave the metaphysician at sea with respect to the ontological status of mentality — especially if, as Fodor [1981] entertains, functionalism characterizes mentality at such an abstract level that it might even permit non-material arrangements to instantiate mental states. Without deciding that question, Fodor's anti-reductionism does inspire parody of the early microreductionists' optimism. Highlighting the problems multiple realizability presents for reduction in science, Fodor [1974] advocates the "disunity of science as a working hypothesis."

The bleak reductive prospects Fodor foresees, in effect, insulate psychological theorizing and research. Psychology — and, presumably, any other science in which theories reign whose ontologies diversely map on to the entities that populate theories at lower levels — enjoys great theoretical and methodological leeway, essentially unencumbered by what might appear to be divergent evidence arising from research at lower levels of analysis. Inquiries at those levels merely concern the details of implementation [Fodor and Pylyshyn, 1988]. They will not bear in any integral way on the theories and principles at stake at the higher level. In short, psychology should proceed autonomously. (How all of this fits with Fodor's impatience with "special pleading" in behalf of some sciences is not obvious [1983, 105–106].)

Putnam, Fodor, and the functionalists' arguments have provided fertile ground for metaphysicians eager to preserve the integrity of the mental (and its causal integrity, in particular) without having to surrender their credentials as disciples of modern science. They also seem to permit the promotion of the metaphysics of mentality without risking breaches of materialism. All mental events, states, and (first order) properties are physical events, states, and (first order) properties, respectively. Armed with these multiple realizability arguments, non-reductive materialists can readily acknowledge that brains constitute a material platform on which mentality supervenes, but they need not concede any ground either to identity theorists or to reductionists, who even ponder the possibility of systematic connections between psychological and neural phenomena (typically referred to as "psycho-physical laws"). Finally, the range of possible instantiations and their details also seem to relieve non-reductive materialists from relying on the vagaries surrounding the notions of *emergence* and *supervenience* in order to make their cases. (See [Wimsatt 1997] and Wilson and Craver's "Realization" in this volume.)

Kim, for example, correctly notes that an assertion of "mind-body supervenience *states* the mind-body problem — it is not a solution to it" [1998, 14]. Claims about the emergence or supervenience of the mental are only promissory notes, at best. As with government deficits, the familiarity these notions enjoy and the ease with which contemporary philosophers speak of them does not obviate in the least what is an ever increasing need for their repayment. Non-reductive materialists view functionalism in the philosophy of mind and the multiple realizability argument to which it appeals as a first, major installment against their metaphysical debt.

The second influential anti-reductionist argument looks to the character of our mental experience rather than to technical difficulties associated with disentangling the myriad ways that some psychological phenomenon might be realized by various physical arrangements. This second argument, which is at least as old as Plato's *Phaedo*, holds that it is inconceivable that our conscious experience will ever be explicable or even describable in exclusively physical terms. With the advent of modern neuroscience, philosophers have formulated more finely tuned variations on this position. At a minimum, though, they all deny that any of the conceptions, theory, or language of the sort that inform contemporary physical and biological science will ever suffice to fathom the character of conscious experience. Conscious mental experience may invariably correlate with brain events of specifiable types, but these anti-reductive positions are united in maintaining that no possible amount of information about brains (or any other physical systems) of the type the sciences glean will capture the qualitative character of our conscious experience. These philosophers ([Jackson, 1982; 1986]; [Levine, 1983; 1997]; [Chalmers, 1995; 1996], etc.) have advanced assorted formulations of the features of our conscious experience at issue, but the general point is that even the most comprehensive accounts of the structures and operations of what look to be the relevant brain mechanisms will inevitably fail to convey "what it is like" to see the redness of a ripe tomato, to taste the subtle flavors of a Brussels sprout properly prepared, or to hear Callas' voice.

These anti-reductionists concede that psychological and neuroscientific research may illuminate the comparatively easy problems about mind — basically, those concerning mental contents and access to information, but they insist that these sciences will prove inadequate to the challenges that conscious experience entails. They will not solve the "hard problem" of consciousness [Chalmers, 1995; 1996]. With respect to qualia they will inevitably manifest an "explanatory gap" [Levine, 1983]. The point is not that the qualitative character of consciousness involves features for which we have yet to find any mappings into theories about brains. Rather, their advocates contend, those features are such that no mapping could ever adequately re-describe them. It is this *inevitable* failure to map qualia into theories about brain activities that is the alleged barrier to reduction.

Those non-reductive materialists, who want both the materiality of their cake and their irreducible ability to taste it as they eat it too, may worry that these philosophers' convictions concerning the latter may sabotage their status as materialists. At first blush, the resulting position looks every bit as much like property

dualism as it does like some form of materialism. But if consciousness reflects a level of complexity that is beyond our abilities (or our sciences' abilities) to comprehend, then it is at least possible that there are some physical properties that will ever remain obscure to us [McGinn, 1991]. Given how little of interest we know about our epistemic limitations, the resulting position would seem, at least, to encourage a search for a new fundamental theory of the intrinsic properties of the physical. Owen Flanagan [1992, 128] describes such positions as examples of the "new mysterianism," since from the standpoint of current science these putative intrinsic properties of the physical are mysterious, indeed. Claims about the inability of humans ever to learn more about such properties only introduce an added layer of mystery.

These various non-reductive materialists and New Wave reductionists agree that, whether the issue is *too many* possible connections between psychology and lower level inquiries (the multiple realizability objection) or *too few* (the explanatory gap objection), both constitute barriers to the standard reduction of psychology. Fodor thinks too many connections argue for granting psychology a comparative *autonomy* in its pursuits. Fans of consciousness and the new mysterians, in particular, think that too few (probably zero) connections leave an un-bridged *explanatory gap* and likely point to one that is unbridgeable. New Wave reductionists, by contrast, think that either of these two circumstances will lead at least to psychology's fragmentation [P. S. Churchland, 1983], probably to its dismantling ([P. M. Churchland, 1989]; [Bickle, 1998; 2003]), and perhaps even to its outright *elimination* in at least some of its sub-domains ([P. M. Churchland, 1979; 1989]; [P. S. Churchland, 1986]).

The following three sections take up problems with the New Wave view and with its prognostications about the elimination of psychology. Presenting an alternative account of these matters there and in the remaining sections will include not only richer, more fine-grained analyses of reductive possibilities in psychology and the cognitive sciences that are more faithful to the wide range of considerations that bear on scientific practice but also proposals for defusing both the multiple realizability objection (in section 8) and the explanatory gap objection about consciousness (in section 10).

5 NOT SO NEW WAVE REDUCTIONISM (AFTER ALL): MUTUAL PREOCCUPATION WITH THEORIES

New Wave reductionists have usefully criticized the standard model, exposing the limitations of its restrictive assumptions. The New Wave model's continuum of intertheoretic mapping allows for a range of intermediate relationships that more or less closely approximate the standard model's ideal. They provide little detailed guidance, though, about where approximate reduction collapses into telling discontinuity and about just where on that continuum controversial cases fall. Consider, for example, the putative reduction of Newtonian mechanics to the mechanics of relativity. Nagel [1961/1979, 111] offers this case of intertheoretic relations as an il-

lustration of the standard model of reduction. By contrast, Paul Feyerabend [1962] advances the very same case as a counter-example to the standard model's dictates. However praiseworthy the improvements are that New Wave reductionists have introduced, their analyses also retain at least two *problematic* commitments of the logical empiricists' conception of reduction. They are the first and fourth of four assumptions that anchor the New Wave position.

Standard model and New Wave reductionists assume that distilling sciences down to their *theories* is epistemologically unproblematic. This is the first of the two problematic commitments and the first of the four assumptions. A second assumption, which the New Wave and standard models also share, is that the scientific promise of exploring intertheoretic relations hangs on the set of *functions* (in the most generic sense) that facilitate the characterization of one theory in terms of another. These reductionists hold this view in common, regardless of what conception of theories they prefer, i.e., regardless of whether they construe theories as collections of related statements (logical empiricists) or as sets of mathematical models of characteristic structures [Bickle, 1998; 2003] or as configurations of synaptic weights in brains [Churchland, 1989]. (On this point, see [Endicott, 1998, 62–70].) The third assumption, reviewed in section 3 above, is that, whatever their conceptions of those functions, New Wave reductionists further hold that the relevant collections of those functions connecting particular pairs of theories can be ordered along a *continuum* according to the comparative fidelity of the resulting analogue of the reduced theory. And, finally, however they propose to calibrate and divide that continuum up, partisans for the New Wave model further assume that it provides a *single model of intertheoretic relations with undeviating implications* (depending on where any particular case falls on their continuum) that can account for all of the significant ways in which scientific theories might be connected. This is the second problematic commitment that they hold in common with the defenders of the standard model.

The first half of this section examines grounds for questioning the first assumption, i.e., the mutual commitment of the New Wave and the logical empiricists' models that, ultimately, theories and their relations to one another generally exhaust what is of epistemological interest about cross-scientific connections. In the course of advancing an alternative conception of these matters, sections 6 and 7 present reasons for rejecting the New Wave advocates and the logical empiricists' second mutual commitment that a single interpretation of their models of intertheoretic relations suffices to capture all of the notable intertheoretic connections that arise within science, i.e., the fourth assumption above.

New Wave reductionists often seem to assume, along with their logical empiricist predecessors, that the only epistemologically interesting features (at times they seem to think the only interesting features at all) of cross-scientific relations are the explanatory connections that hold between *theories*. That Paul Churchland's discussions (e.g., [P. M. Churchland, 1989, 48–50]) often exhibit a substantial interest in the reconstruction of a theory's laws is a corollary of this principle. Although New Wave accounts have eschewed the traditional preoccupations with

the *deduction* of the reduced theory's laws, Churchland's equipotent image of the reduced theory constructed within the framework of the reducing theory in an approximate reduction, nonetheless, remains almost exclusively concerned with the reconstruction of the *laws* (and other generalizations) of the reduced theory. That he would show such interest in theories' laws is also somewhat unexpected in light of Churchland's arguments for the advantages of his prototype activation model over the deductive nomological model of explanation. (See [P. M. Churchland, 1989, ch. 10]; [Churchland and Churchland, 1996, 257–264].) By contrast, Bickle [2003, 16] and Schaffner [1992, 329] hold that psychoneural reductions will not involve laws.

This may not be exactly what either the logical empiricists or the New Wave reductionists explicitly say, but it is unquestionably what their discussions of reduction regularly *imply*. Their preoccupation with theories entices New Wave reductionists in their philosophical commentary to minimize the epistemological import not merely of preserving but of cultivating multiple levels of explanation in science and of attempts to integrate the associated inquiries that occur at each level. Instead, their discussions focus overwhelmingly on intertheoretic relations and the putative deflationary implications of intertheoretic reductions for theory and ontology — and for the theories and ontology of psychology, in particular.[4]

They do so even though the neuroscientists and neural network modelers, whose work they discuss, often appeal to findings from higher level psychological sciences to support the neuroscientific and neurocomputational models they prefer. Crucially, those researchers look to these findings (and, by implication, to the methods and techniques by which they were generated) not merely for guidance (e.g., [Hirst and Gazzaniga, 1988, 276, 294, 304–305]) but for *support* for their favorite hypotheses. New Wave and standard reductionists alike concede a role to the special sciences in scientific *discovery*, but the latter, in particular, are clear that the context of discovery is not epistemologically decisive. Both standard and New Wave models' spotlight on the *reduction* of theories in cross-scientific contexts renders them largely insensitive to the contributions higher level sciences regularly make to the *justification* of the very scientific theories — particularly some in neuroscience — that they champion. Neuroscientists and neurocomputational modelers regularly cite psychological evidence in support of their proposals, yet New Wave reductionists often lose sight of all of this in their explicit philosophical analyses and especially in their examinations of the relations between psychology and neuroscience.

Terrence Sejnowski and Charles Rosenberg [1988], for example, appeal to experimental findings in cognitive psychology to support their neurocomputational proposal about the character of human memory. Their famous connectionist model

[4]Bickle [2003, 114 and 130] has conceded an ineliminable "heuristic" role to psychological theorizing and research. Parity of reasoning suggests that theories cast at the sub-cellular and even molecular levels within neuroscience, which Bickle unequivocally prefers, would eventually occupy the same status relative to even lower level theories in related areas of chemistry. (See [Bickle, 2003, 115 and 157]. Contrast [Craver, 2002].)

[1987] for transforming graphemes into phonemes, NETtalk, appears to pronounce written text aloud, once its output is fed into an acoustical synthesizer. Sejnowski and Rosenberg argue that NETtalk's processing and underlying connectionist architecture also offer valuable insights about memory. They advance as a *complementary* theory to prevailing theories in cognitive psychology a new proposal that stresses the *form* of memory representations. Instead of advancing their theory as a competitor that might correct, let alone eliminate, either Bower's encoding variability hypothesis [1972] or Jacoby's processing effort hypothesis [1978] (two hypotheses they explicitly cite), they explore how connectionist architectures could effect such processes and how such connectionist modeling will suggest grounds for the *elaboration* (and, presumably, the reconciliation) of these two cognitive hypotheses.

The first important point here is that the principal evidence for their theory that they cite is the substantial similarities between NETtalk's performance with a particular mnemonic task and the findings in experimental cognitive psychology about the performance of human subjects. As with human subjects, NETtalk's mnemonic performance displays a short-term advantage for massed practice with items and a longer-term advantage for distributed practice, which is to say that it displays the classic spacing effect. This is one of the oldest, best known, and most thoroughly explored findings in the psychology of human memory, dating back to the nineteenth century. It is, for example, discussed in the seminal work in experimental psychology on memory, viz., Hermann Ebbinghaus' *Memory: A Contribution to Experimental Psychology* [1964/1885].

The connections of Sejnowski and Rosenberg's project to research in psychological science does not stop there. One of the familiar experimental designs in memory research in cognitive psychology, viz., cued recall, provides the model for their tests of NETtalk's mnemonic abilities. Since NETtalk is incapable of generating free recalls, for comparative purposes Sejnowski and Rosenberg had to confine their attention to a small subset of the experimental literature on the spacing effect in humans, viz., that concerned with the demonstration of the spacing effect in experiments employing cued recall (e.g., [Glenberg, 1976]). Notably, their argument for the consequence of their findings concerning their neurocomputational model turns, first, on its consilience with findings about the performance of human subjects from experimental cognitive psychology and, second, on the resonances between their tests of NETtalk's performance and the designs employed in experimental work on cued recall [McCauley, 1996].

(For additional examples of the salient role that the integration of evidence arising from various experimental techniques employed in both psychology and neuroscience plays at the borders between these sciences, see [Bechtel and Mundale, 1999]; [Bechtel and McCauley, 1999]; [McCauley and Bechtel, 2001].)

Neuroscientists and neurocomputational modelers regularly cite psychological evidence in support of their proposals, yet the philosophical analyses of New Wave reductionists frequently overlook all of this, especially in their examinations of the relations between psychology and neuroscience. To make this charge is not

to belittle the importance of intertheoretic relations in these settings. Certainly, understanding theories and their relations to one another *is* critical for comprehending much of what goes on in science, however, if the only available accounts of science and the sciences were those of the standard model and its not-so-new-wave successors, a ready inference might well be that *theorizing* exhausts what is epistemologically consequential about scientific activity. The logical structures, material commitments, and mutual relationships of scientific theories are not the only topics relevant either to the justification of those theories or to progressive programs of scientific research. Making sense of theories always requires — in addition to discussions of the theories themselves — discussions of scientific practices and experimental designs, the evidence those practices and experiments generate, and the appraisals of that evidence's import. When that evidence arises from sciences operating at different levels of analysis, it requires at least some attention to the overall structure of science as well. To help clarify some of the issues at stake and to provide an analytical framework for much of what follows will require a sidebar at this point concerning the overall architecture of science and, more specifically, how the various levels of analysis (or explanation) are distinguished.

Talk of analytical or explanatory levels is rampant throughout the literature of the cognitive and neural sciences, but systematic characterizations, let alone precise ones, are rare [Hardcastle, 1996]. A group of criteria for locating levels of explanation among the sciences roughly converge — at least with respect to theorizing about the structural relations of systems. Some of these criteria look to what might be called levels of organization in nature, but ontological and epistemological considerations become rapidly intertwined as these analyses proceed.

Presumably, our most successful theories provide significant clues about the furniture of the universe. This suggests that levels of analysis in science correspond to levels of organization in nature. Typically, what counts as an entity depends on both the redundancy of spatially coincident boundaries for assorted properties and the common fate (under some *causal* description) of the phenomena within those boundaries. For example, both their input and output connections and their various susceptibilities to stains aid in identifying cortical layers in the brain. Emphasizing causal relations insures that explanatory theories in science dominate such deliberations. Wimsatt [1976] suggests that the frequency of items' causal interactions positively correlates with the proximity of the organizational levels at which those items occur. So, usually items at the same level of organization in nature causally interact most often. The greater the number of theoretical quarters from which these ontological distinctions receive empirical support, the less troublesome is the circularity underlying an appeal to *levels of organization* as criteria for their corresponding levels of analysis. Herbert Simon [1969], for example, notes that the amounts of energy necessary to hold systems together increase at progressively lower levels of organization. The forces that sustain the organizational integrity of molecules far exceed those that bind together a block of wood. (This is why experts in karate can break apart blocks of wood, but not molecules, with their hands.)

The *range* of the entities that constitute any science's primary objects of study and its principal units of analysis also offer grounds for distinguishing analytical levels. The lower a science's analytical level, the more widespread the entities it studies. For example, subatomic particles, discussed in physics, are the building blocks of all other physical systems (from atoms, galaxies, and molecules to brains, social groups, and more). By contrast, the minds that psychologists study are only uncontroversially accorded to some (indefinite) subset of biological systems and are parts of socio-cultural systems only.

Such mereological considerations are relevant in distinguishing analytical levels in science but not unqualifiedly so. Analytical levels partially depend upon viewing nature as organized into *parts and wholes*, however the pivotal question for the differentiation of analytical levels is whether or not the wholes are organized or simply aggregates of their parts [Wimsatt 1974; 1986; 1997]. One way of casting the epistemological issue reductionists and anti-reductionists battle about is whether any features of wholes resist explanation in terms of their parts, i.e., whether from an explanatory standpoint wholes are greater than the sums of their parts. The aim is to rule out an account of organizational levels (and, on these criteria, to thereby rule out an account of analytical levels) that tracks simple considerations of scale. For not all big things that have lots of parts (e.g., asteroids or sand dunes) are, comparatively speaking, highly integrated things. If one entity contains others as its parts and if explanations of (some of) its behaviors require appeal to further organizing principles beyond those concerned with those parts, then it occurs at a higher level of organization and likely points to a distinguishable analytical level.

Much of this is an attempt to explicate our intuitions about the *comparative complexity* of systems that arise at different analytical levels. Although complexity has no simple or single measure, we usually have little trouble making comparative judgments about such matters. Cells are more complex systems than crystals. Comparing the relative complexity of systems seems to generate a similar picture of analytical levels in science, with progressively higher levels handling progressively more complex systems.

The order of analytical levels also corresponds to the chronological *order in natural history* of the evolution of systems. Somewhat more roughly, it also corresponds in the history of modern science to the order in which the inquiries in question emerged as disciplines to be differentiated from natural philosophy, as measured by the use of distinctive terms to indicate independent disciplines and the founding of specialized university departments, journals, and professional societies. The lower a science's analytical level, the longer the systems it specializes in have been around. For example, the subatomic particles and atoms that are the principal objects of study in the basic physical sciences appeared quite soon after the Big Bang whereas the compounds and systems on which the biological sciences focus first began to appear (on Earth, at least) somewhere around two billion years ago. Developed nervous systems and brains and the minds that eventually seemed to have accompanied them, by contrast, are more than a billion years newer. And,

finally, cultural systems that the socio-cultural sciences study date from a few million years ago on the very most optimistic estimates and, perhaps, no more than some tens of thousands of years ago (on more exacting criteria). Figure 2 summarizes how this as well as considerations about range and complexity organize the analytical levels of science.

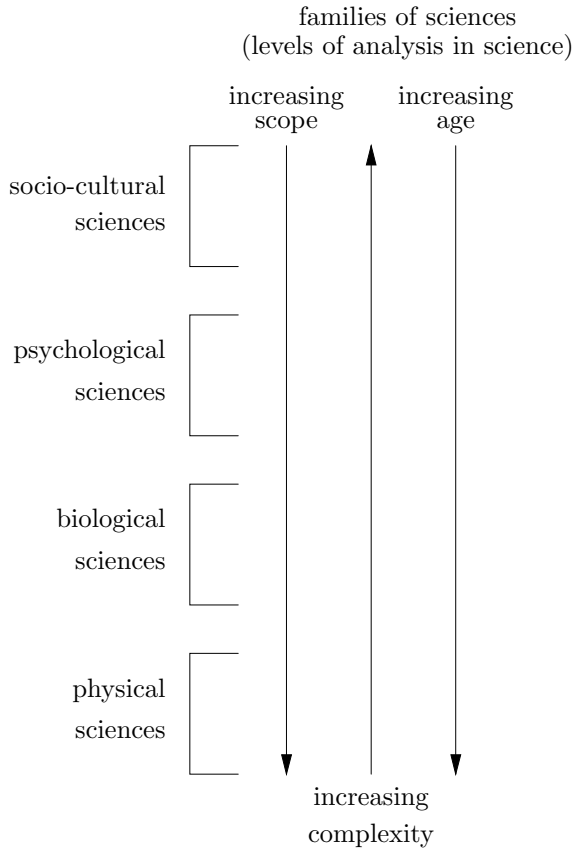

Figure 2.

Methodological considerations also segregate analytical levels but less systematically. Sciences at different analytical levels ask different questions, promote different theories, and employ different tools and methods. Theories at alternative explanatory levels embody disparate idealizations that highlight diverse features of the phenomena on which they concentrate.

The general assumption is that all of these various criteria converge on a grouping of the major scientific families into levels as portrayed in figure 2. Each of these families includes separate sciences that address specific domains; the physical sciences include such sciences as physics and chemistry; the biological sci-

ences include such sciences as molecular genetics and neuroscience, and so on. These sciences, in turn, contain multiple sub-levels [Mundale, 2001]. Employing principles of organization and scale, Churchland and Sejnowski [1992, 11] readily identify seven sub-levels within neuroscience alone (molecules, synapses, neurons, networks, maps, sub-systems, and the central nervous system overall). Figure 3 only begins to hint at the multitude of specialized sciences that have emerged within some of these families of sciences.

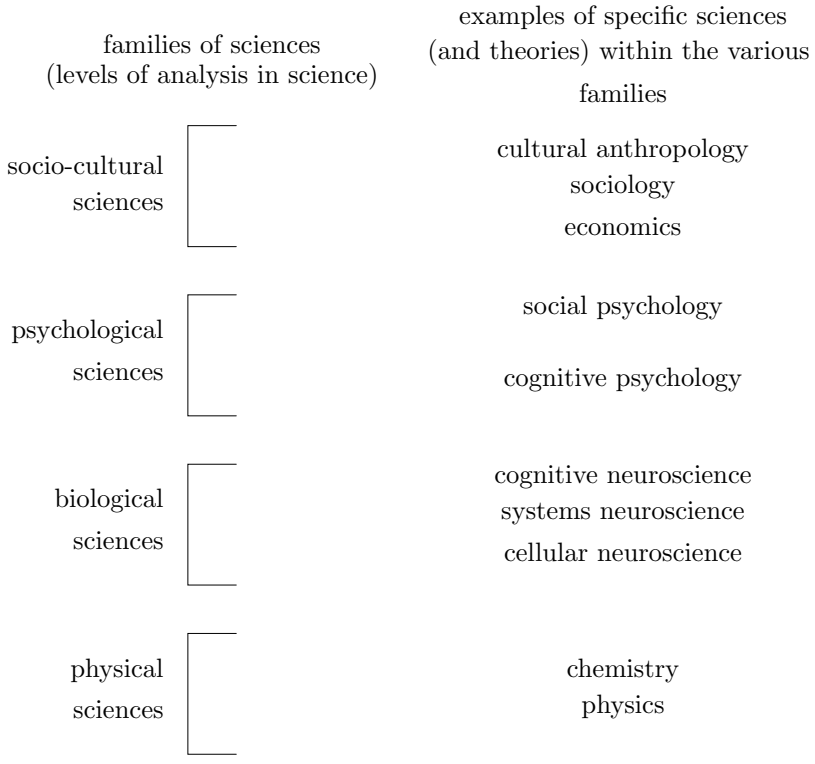

Figure 3.

In contrast to the broad consensus about the groupings of the families of the sciences that figure 2 illustrates, differentiating distinct analytical levels at this more refined register can prove more controversial. This is especially true when trying to sort things out early in some science's history, when decisive accomplishments that end up inaugurating long-standing research traditions that substantially define a sub-level have not yet been recognized as such. The publication of Ulric Neisser's *Cognitive Psychology* [1967] demarcated that sub-discipline within psychology by laying out its salient topics, approaches, theories, and findings and by pinpointing pivotal work by a variety of researchers that in some cases (e.g., [Bartlett, 1932] preceded the appearance of Neisser's book by decades.

6 NOT SO NEW WAVE REDUCTIONISM (AFTER ALL): INSISTENCE ON A SINGLE MODEL

The second problematic assumption that the New Wave reductionists share with the logical empiricists is that a *single model* (or, more accurately, in the New Wave case, a *single interpretation* of their model) can account for all of the (epistemologically) interesting connections between scientific theories. In the same spirit reductionists who subscribed to the standard model also defended a single model of intertheoretic relations, albeit an abbreviated one compared to that of the New Wave or, alternatively, one that assumed (falsely) that every case fell at or quite close to the endpoint of the New Wave continuum designating thoroughly smooth reductions. (Arguably, the logical empiricists' discussions of reduction in science never even countenanced the bumpy cases.) As is the case with the New Wave model, the logical empiricists' single model allegedly describes the critical dynamics connected both with theory succession over time within some science as well as with how the sciences hang together at any particular moment in scientific history, i.e., with intertheoretic relations in cross-scientific contexts.

Clearly, the assumption that a single version of intertheoretic relations (whether the formal standard model or a single interpretation of the New Wave continuum model) can capture the full range of worthwhile scientific connections rests upon their first (shared) problematic assumption that all of the relevant relations are ones or can be condensed to ones between theories. Arguments in the first half of the previous section about the psychological evidence to which researchers regularly appeal in support of their proposals about the structure and operations of the brain (and, thus, that New Wave reductionists regularly presume when they plump for the explanatory promise of some of their favorite candidate reductions) belie that first assumption. The most important consideration here is that making the case in behalf of theories at lower levels quite regularly depends, in part, on appeals to evidence that has been generated in higher level sciences by employing their characteristic methods and experimental techniques.

If these criticisms of the New Wave's first assumption are sound, then they also endanger the current assumption, since, as noted, it depends upon the first. If, as the first assumption asserts, sciences can be distilled down to their dominant theories, then it follows on the New Wave reductionists' account that their comprehensive continuum model of *intertheoretic* relations and their (single) interpretation of its methodological and ontological implications usefully apply to the full range of *cross-scientific* relations as well. This section will offer reasons for thinking that that conclusion is false and, thereby, disclose both why a commitment to a single interpretation of their model of intertheoretic relations is overly ambitious and why any putative reductions of *sciences* (as opposed to theories), including any purported reduction of psychology to neuroscience, are not only not disabling but, instead, a vindication of the "reduced" science and its ontology.

Although their models are less sophisticated than the New Wave models, standard reductionists exhibit some sensitivity to the issues at stake. For example,

even though he too aims to provide a unified treatment of both, in his discussions Nagel [1961/1979] distinguishes between "homogeneous" and "heterogeneous" reductions. Roughly, these options correspond (respectively) to cases in which the reduction functions are straightforward and clear and the terminology is consistent as opposed to those where they are less so. Attention to his discussion and the examples he supplies, however, reveals that Nagel's distinction closely tracks another, more important distinction, viz., one between modeling *progress within a particular science* and modeling *cross-scientific connections*. Nagel discusses the homogeneous reduction of a *theory* and the heterogeneous reduction not only of a theory but also of a *science*. The burden of the discussion that follows is to show that this latter distinction matters (even if Nagel's model mis-describes its implications) and that New Wave reductionists are remiss in ignoring it. Making that case will also uncover grounds for holding that anything anyone is even remotely tempted to describe as "the reduction of a science" will not — contrary to the claims of the New Wave reductionists — jeopardize its importance in the least, let alone lead to its displacement or elimination either in practice or in principle.

Note that *if* sciences could be legitimately distilled down to their theories as both the standard and New Wave models presume, then these putative reductions of *sciences* (say, reductions of perceptual psychology to neuroscience) would also be best understood as reductions of *theories*. The burden of the preceding discussion, however, has been that such a distillation involves unhelpful simplifications and, in particular, unhelpful simplifications about the *justification* of the lower level theories that reductionists prefer.

The homogeneous cases with respect to which Nagel utilizes talk of *theory* reduction concern the relations of successive theories within a single science. These are episodes in the history of science where a new superior theory in some science eclipses what had been the reigning theory. Thus, Nagel employs this language exclusively to explicate, for example, how Newtonian mechanics superseded Galileo's law of free fall [Nagel, 1961/1979, 338–339]. Nagel regards reductions of this sort as relatively unproblematic and comments that they "are commonly accepted as *phases in the normal development of a science*" (p. 339, emphasis added). For Nagel and the logical empiricists homogeneous theory reduction serves as the foundation for their account of scientific *progress*. The new reducing theory and the older reduced theory are continuous with one another logically and materially. Sciences progress through the discovery of new, more encompassing theories that explain everything their predecessors do and more. While the Galilean law only addresses free fall relatively close to the surface of the earth, Newtonian mechanics supplies laws that explain not only this and other terrestrial motions but celestial motions as well. According to the logical empiricists, the history of modern science is a story of the progressive accumulation of discoveries about the world organized by theories employing explanatory principles of increasing generality.

In his most extended treatment of these topics Nagel [1961/1979, 338–339] gives these homogeneous cases of scientific "development" little space. His principal concern is to address what he sees as the more difficult problems that surround

the heterogeneous reductions of *sciences*. The salient point here is how Nagel ends his discussion of the heterogeneous reductions. He chides both allies and opponents of the reduction of sciences for their mutual failure to acknowledge the need to qualify their claims *temporally*. He complains [1961/1979, 364, emphasis added] that " ... questions that at bottom relate to the strategy of research, or to the logical relations between sciences *as constituted at a certain time*, are commonly discussed as if they were about some ultimate and immutable structure of the universe." The aim here is to underscore Nagel's recognition that all such claims concern the various sciences "as constituted at a certain time."

Wimsatt [1976] explicated the critical distinction that lurked behind Nagel's separation of homogeneous and heterogeneous reductions. Nagel's homogeneous reductions of theories dealing with "phases in the normal development of a science" reliably concern what Wimsatt referred to as "*intra*level" relations, i.e., the relations over time between successive theories in some science. The relations of consecutive theories of greatest interest in this context are those between a reigning theory in some science and a competing theory (in the same science) that eventually supersedes it, e.g., in geology from 1950 to 1980, the relations between the theory of the earth's crust as a unitary structure that rises and falls and the theory of tectonic plates or in the neuropsychology of the late nineteenth century, the relations between David Ferrier's [1876] hypothesis that primary visual processing occurs in the angular gyrus and the view of Salomon Henschen [1893], among others, locating it in occipital cortex, instead. Intralevel relations concern theory succession. The crucial point is that these are the relations of competing *theories over time within a particular science operating at a single level of analysis*.

By contrast, Nagel's heterogeneous reductions of sciences pertain to what Wimsatt calls "*inter*level" relations, i.e., the cross-scientific relations between *theories that reign at the same time at different analytical levels in science* — for example, in biology the relations between the theories employed in population and molecular genetics or in the cognitive sciences the relations between psychological conjectures about a specialized capacity for the visual processing of human faces [Farah *et al.*, 1998] and proposals that the fusiform gyrus in the temporal lobe is where such processing is carried out in the brain [Kanwisher *et al.*, 1997].

Wimsatt proposes no longer examining just the logical relations of theories but also the impact of temporal considerations on construing the relations of theories both *within* and *between levels of analysis* in science (see figure 4). Careful scrutiny of these two kinds of contexts will reveal how much their methodological and ontological implications for theories and sciences contrast — especially when the prospects for intertheoretic mappings are bleak [McCauley, 1986; 1996]. Such scrutiny will also provide grounds for resisting the most extreme anti-psychology impulses of New Wave reductionists.[5] Ironically, to see why, it will be useful, first, to register one way in which the New Wave reductionists correctly capture

[5]Such impulses have gradually subsided in the work of some New Wave reductionists. Contrast, for example, P. M. Churchland [1979, esp. ch. 5] with Churchland and Churchland [1996, 219–220] and [1998, 79].

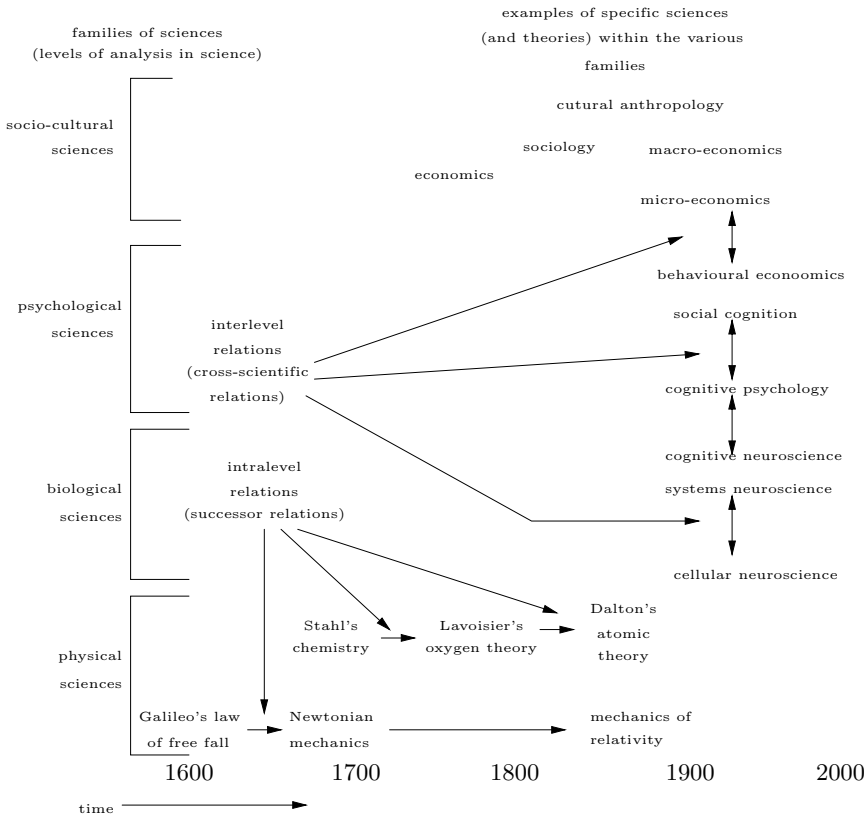
families of sciences
(levels of analysis in science)
examples of specific sciences
(and theories) within the various
families
socio-cultural
sciences
cutural anthropology
sociology
macro-economics
economics
micro-economics
behavioural econoomics
psychological
sciences
interlevel
relations
(cross-scientific
relations)
social cognition
cognitive psychology
cognitive neuroscience
biological
sciences
intralevel
relations
(successor relations)
systems neuroscience
cellular neuroscience
Dalton's
atomic
theory
Stahl's
chemistry
Lavoisier's
oxygen theory
physical
sciences
Galileo's law
of free fall
Newtonian
mechanics
mechanics of
relativity
1600
1700
1800
1900
2000
time

Figure 4.

how these two sorts of cases are *alike*. Both sorts of contexts *do* involve sets of possibilities that range across the New Wave continuum of goodness of intertheoretic mapping. Our ability to construct an analogue of one theory on the basis of a second theory's conceptual resources is unaffected by whether or not those two theories are successors within a single science or prevailing theories employed simultaneously at two different levels of analysis. On this front the two sorts of cases *are* the same.

Note that the point of contention does not concern the applicability of the New Wave continuum in both sorts of intertheoretic settings but rather New Wave partisans' insistence that their *single interpretation* of that continuum's implications will work equally well in both. Contrary to that view, the methodological and ontological implications of falling at some point or other on this continuum can differ substantially in the two different contexts, particularly where the probabilities for finding substantial connections between theories are meager. To help see why, consult figure 5. Introducing the distinction between interlevel (or cross-scientific) relations and intralevel (or successive) relations (portrayed in figure 4) yields *two* contexts in which the continuum of the goodness of intertheoretic mapping applies but, notably, in which its implications diverge when intertheoretic mappings are inauspicious.

Historically, philosophers spotted major cracks in the standard model of reduction when they pondered cases of poor mapping between successive theories in intralevel contexts. As Thomas Kuhn [1970] and Paul Feyerabend [1962] famously emphasized, contrary to the standard model's account of scientific progress, sometimes advances turn not on the discovery of more inclusive theories that are continuous with those that have preceded but rather on largely discontinuous theories, which offer little hope for constructing analogues of their predecessors. Contrary to the standard model of homogeneous reductions, the progress that such revolutionary episodes in the history of science involve does not readily lend itself to characterization in terms of accumulation. In the most extreme cases, the old theories and their ontologies are basically discarded.

New Wave reductionists have learned this lesson well. On their single interpretation of their model of intertheoretic relations the implications of a persisting inability to construct an analogue of one theory within the framework of the other are unvarying regardless of the context — sooner or later, they maintain, one theory's success demands the other's elimination. New Wave reductionists assume that the implications of falling at any particular point on their continuum are unaffected by *how and where* the pertinent theories are situated among the sciences. That contextual, pragmatic, problem solving, and (even) evidential considerations can or should bear on the interpretation of the ontological implications of the varied cases of unpromising intertheoretic mapping never enters the New Wave picture. New Wave advocates presume that any case of serious incommensurability will inevitably and uniformly result in one theory completely superseding another, permanently removing the latter from the scientific stage. In their philosophical proposals at least (see, for example, [Bickle, 1998, 30, figure 2.1]), they

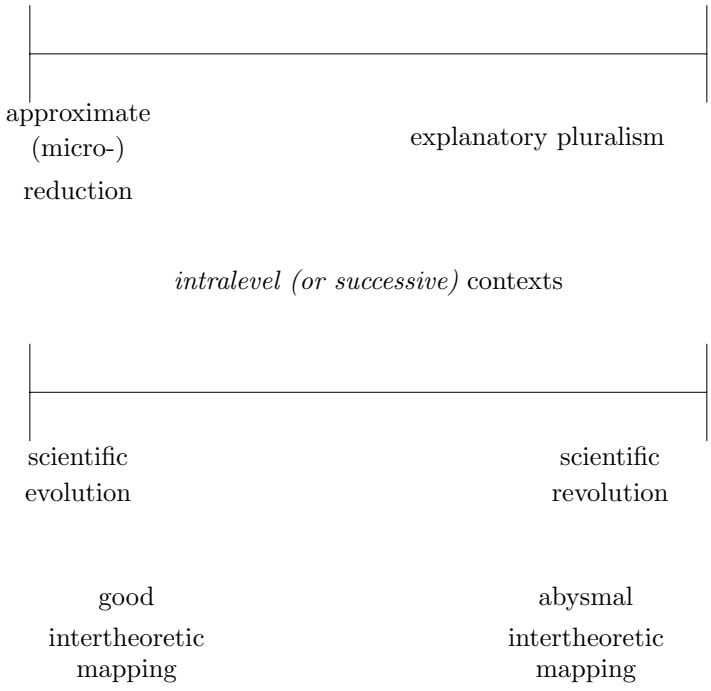

Figure 5.

anticipate this outcome, regardless of

- the amount of empirical support that each theory enjoys,
- the level of explanation in science that each theory occupies,
- the institutional health and longevity of the sciences in which the theories arise,
- the relative status and position of the theories within their respective sciences (for example, are either or both central theories that motivate progressive programs of research?), and
- the amount of fruitful interaction between each theory and other theories at explanatory levels *other* than those at which the two theories in question occur

to mention just some of the more prominent considerations that would influence the probability that one or the other will undergo elimination exclusively on the basis of such a mapping failure.

In fact, these sorts of large-scale, fell-swoop eliminations are an accurate prognosis in intralevel settings only. Where theories are substantially discontinuous in intralevel, successive contexts, the result, when the new theory triumphs, *is* something like a Kuhnian scientific revolution (though, see [Thagard, 1992]). The inability to map out a plausible analogue of the previously reigning theory (e.g., the theory of the bodily humors) employing the conceptual framework of the newly ascendant theory (e.g., modern theories of physiology and of disease) does lead to the elimination of the former theory and some of its accompanying ontology. This is true notwithstanding the persistence of its characteristic idioms in everyday parlance or as modified technical terms in scientific discourse [P. S. Churchland, 2002, 21]. Consider, respectively, the continued use of "choleric" and "sanguine" to describe personality types and the continued use of the term "planet" in the Copernican system or "module" in contemporary psychology [Fodor, 1983, ch. 1].

Reliably, the examples of theory elimination to which New Wave reductionists point — the theories of the crystalline spheres, alchemical essences, phlogiston, caloric fluid, the aether, phrenological faculties, etc. — concern a new theory superseding an older theory at the *same* level of analysis. They concern, in short, *intra*level relations. Because they often fail to distinguish them sharply, New Wave reductionists attribute the profile and outcomes characteristic of intertheoretic relations in intralevel settings to *both* sorts of contexts. Consequently, they anticipate theory elimination when sciences at two different levels of analysis provide disparate views of the same phenomena, i.e., in cross-scientific contexts where the dominant theories at two different analytical levels fail to agree and, therefore, where the prospects for productive intertheoretic mapping (and subsequent co-evolution) do not seem promising. According to the New Wave position, the inability to trace an analogue of the upper level theory within the framework of the lower level theory signals the inevitable elimination of the former by the latter. (Reliably, on the New Wave view it is the upper level theory that should go.) But because New Wave reductionists also subscribe to the (first) assumption that — for most epistemological and ontological purposes — sciences can be distilled down to their theories, the elimination of a theory in these interlevel contexts would be tantamount to the elimination of a *science*! At least some of the time, such going-out-of-business sales are just the result that they predict [P. M. Churchland, 1979; 1989, ch. 1].

The crucial point, however, is that the premier case to which New Wave reductionists wish to apply this moral of the elimination of theory (and, therefore, of a science), viz., the relations between theories in psychology and neuroscience, is precisely a case of *inter*level, cross-scientific relations ([McCauley, 1986]; cf. [Looren de Jong, 1997]). Historical, sociological, and normative considerations argue for why we should neither expect nor wish for the elimination of theories on such grounds in such cross-scientific contexts.

Historical and sociological considerations suggest that once sciences are up and running with research groups, journals, professional societies, university departments and the like — if, for no more reason than social inertia — eliminations of their theories do not occur solely (or even primarily) on the basis of their failure to prove consilient with prevailing theories at other analytical levels. This is particularly so when the putative elimination involves long running sciences and theories from wholly different scientific families (in the case at hand, the psychological as opposed to the biological sciences) that have, in fact, only begun to be very usefully interwoven. That interweaving has arisen as a result in large part of developing better research tools within the higher reaches of recent neuroscience such as the new brain-imaging technologies, PET and fMRI.

Typically, scientists do not look to an adjacent level of analysis in order to shear away its theories and wreak havoc with their ontologies. On the contrary, they are usually looking for help, hoping to find suggestive theoretical, methodological, or evidential resources. To repeat, reductionistic research strategies are probably the single most effective heuristic of discovery in the history of science. Interlevel influence and benefit depend upon forging such connections. In actual scientific practice confronting theoretical incompatibilities across analytical levels does not inspire triumphal campaigns of scientific conquest or usurpation. If anything, it initiates inquiries about points of possible cross-scientific connection. Their development occasions the "co-evolution" of theories at different levels of analysis of the sort that Patricia Churchland [1986; 2002] recommends and, eventually, the emergence of full-blown "interlevel theories" that aid the integration of scientific disciplines. (See [Maull, 1977]; [Darden and Maull, 1977]; [Bechtel, 1986a].) Although *some* of Churchland's commentary (e.g., [1986, 373]) on the probable fate of psychological theories on the basis of their relationships to developments in neuroscience reflect orthodox New Wave views, her discussions of the co-evolution of theories and an "integrationist strategy" [2002, 29] clearly demonstrate her interest in interlevel research.

Normative considerations suggest that this process does not ensue merely because of social inertia. The overall scientific enterprise would be woefully impoverished, if in the face of failures to map theories from two adjacent levels of analysis *at one particular moment in the theoretical evolution of each* (to reiterate Nagel's point) scientists decided to abandon not just the theory but all further investigation at the higher level. Lost would not only be sources of theoretical and conceptual novelty as well as testing procedures and experimental techniques but, frequently, ready access to treasure troves of evidence as well. The many benefits of an *explanatory pluralism* and the multiple analytical perspectives that accompany it and of the co-evolution of theories at different levels of analysis, which ordinarily results, would evaporate, if the failure of intertheoretic mapping in such cross-scientific contexts sufficed for the elimination of one or the other theory and (on the New Wave and standard model's first mutual commitment) the eventual enervation, if not elimination, of the entire science from which that theory springs.

Because the theoretical commitments of psychology (either from the various

areas of experimental psychology or from commonsense psychology) do not thoroughly square with the latest theories and findings in neuroscience is no reason to expect their impending elimination. In addition to the useful role that these notions play not only for practical purposes but in the workings of higher level social sciences, the history of modern science since the mid-nineteenth century provides no compelling precedent for holding either that these theories are prime candidates for elimination or that the psychological sciences, more generally, will likely close up shop (as a result of competition from neuroscience). The following clarification, however, is imperative. To make this claim is *not* to say that either psychological theories or the psychological sciences are autonomous of the prevailing theories (or sciences) at alternative levels (e.g., the neuroscientific) or that they remain in isolation from them or that they are uninfluenced by them. Nor is it to say that they ought to be — quite to the contrary! Eliminations of theory and ontology can occur at any level of analysis in science. The disagreement here is about the origins and the relative power of forces influencing such outcomes.

7 EXPLANATORY PLURALISM IN CROSS-SCIENTIFIC SETTINGS

The Churchlands [1996, 230–231] offer four putative counterexamples to the historical conjecture and normative proposal, offered above, that theory elimination in developed sciences does not (and should not) transpire primarily on the basis of profound theoretical conflicts across readily distinguishable levels of analysis. All four fail.

Responding to their second counterexample introduces no matters of philosophical principle. The putative counterexample concerns the elimination of the theory of caloric fluid by the kinetic theory of heat. The controversy, in this case, turns on the problem of distinguishing levels. In order to make the case that this example concerned an *inter*level setting, the Churchlands press a *controversial* historical claim, viz., that caloric was a *macroscopic* fluid. (This is, of course, in contrast to the uncontroversial claim that its putative action had macroscopically detectable effects). Yet a profound problem for their analysis here is that in her subsequent, detailed discussion of this very case, Patricia Churchland [2002, 21–23] offers compelling historical evidence that caloric was conceived as a *microscopic* fluid! This is the account that squares with the consensus among historians of science, and it undermines their earlier contention that it constitutes a counterexample to the explanatory pluralist's historical conjecture and normative proposal.

The first and the fourth of these alleged counterexamples are particularly startling. They concern theoretical arrangements in seventeenth and eighteenth century biology and in atomic chemistry respectively. The Churchlands note, in the first case, that "the conceptual framework of early biology ... was eventually displaced by an entirely new framework of biological notions ... regularly *inspired* by the emerging categories of structural and dynamical chemistry ... " and in the fourth that " ... most of the details of Dalton's atomism ... were *inspired* by higher-level chemical data concerning ... constant weight ratios experimentally revealed

. . . " [1996, 230–31, emphasis added]. (They go on to stress, quite correctly, that in the latter case the cross-scientific influence was *top-down*, not bottom-up.)

These alleged counter-examples startle, because they so clearly miss their target. These are counter-examples to the claim that prevailing theories in sciences operating at different levels of analysis remain in complete isolation from one another. Unfortunately, no one (not even Fodor) holds that view. To repeat the culminating claims of the previous section, to resist assertions about the elimination of whole theories and sciences in cross-scientific contexts is *not* to deny cross-scientific influences. On the contrary, a critical feature of explanatory pluralism is to highlight the prominence of such influences. After all, the benefits of the co-evolution of theories in science will not arise, if theoretical incongruities in interlevel contexts invariably demand either the elimination of theories and sciences, on the one hand, or — just as unhelpfully — their thorough or perpetual autonomy, on the other.

Explanatory pluralism *underscores* the on-going interaction of scientific enterprises carried out at the various analytical levels. All scientific explanation is partial explanation from the perspective of some analytical level or other. Scientific theories and the explanations they inspire are selective. They neither explain everything nor wholly explain anything (cf. [Polger, 2004, 203]. Explanatory pluralism denies that inter*theoretic* relations exhaust all of the selection pressures in the resulting co-evolutionary process and that those selection pressures are exerted exclusively from the bottom up (as the Churchlands' treatment of Dalton's atomism correctly headlines). Bickle comments, for example, that even the parade-case reduction of classical thermodynamics to statistical mechanics involved "mutual feedback" [2003, 11]. So, when the Churchlands tout the abilities of concepts and theories at lower levels (in the first case) or findings at higher levels (in the fourth case) to "inspire" developments at other analytical levels, what they have provided are not counterexamples but *illustrations* of the explanatory pluralist's account of cross-scientific dynamics that they take themselves to oppose. The plurality of explanations at multiple analytical levels is a necessary precondition for the "integrationist strategy" Patricia Churchland [2002, 29] advocates, in which developments (not just theoretical developments, incidentally) at various levels of analysis, both within and between the families of the sciences, inspire and enrich inquiries at alternative levels concerned with the same phenomena under descriptions of different grains. In short, these two of the Churchlands' putative counter-examples are utterly unconvincing.

The Churchlands' third alleged counterexample brings things full circle, since it explicitly addresses the *reduction of a science*. They [1996, 230] maintain that by reshaping our understanding of what light is Maxwell's electromagnetic theory showed:

> . . . the well-established conceptual framework of geometrical optics, while a useful tool for understanding many macro-level effects . . . to be a false model of reality when it turned out that all optical phenomena could be reduced to (i.e., reconstructed in terms of) the propagation of

> oscillating electromagnetic fields. In particular, it turned out that there is no such thing as a literal *light ray*. Geometrical optics had long been inadequate to diffraction, interference, and polarization effects anyway, but it took Maxwell's much more general electromagnetic theory to retire it permanently as anything more than an occasionally convenient tool.

They offer this episode in the history of science as a putative counterexample that will "contradict" the claim that *eliminations* of higher level theories and their ontologies in cross-scientific settings are not driven exclusively (or even predominantly) by incompatibilities with theories at lower levels [Churchland and Churchland, 1996, 224]. The case of geometrical optics fails to serve that purpose, however, because *the interlevel mapping here is comparatively good* (as the Churchlands themselves note). Consequently, it is the wrong sort of case. This case approximates the microreductive ideal. It falls on the *left* half of the upper continuum in figure 5, not the right! Nor does this case involve any fell-swoop eliminations of either theory or ontology. On the contrary, any putative *reductions of a science*, i.e., reductions in situations that involve *good* intertheoretic mappings in interlevel contexts, lead neither to the elimination or the displacement of an upper level theory, but rather to its *vindication*.

It will help to begin by pointing out a consideration that will *not* be employed in making that case. The Churchlands explicitly note that geometrical optics remains "an occasionally convenient tool," presumably for the purposes of rough calculation. (In light of its pervasive use in everyday contexts, to describe the convenience of geometrical optics as "occasionally" useful only may be a bit of an understatement.) Long abandoned theories are, however, often employed for their convenience in calculation. The principles and tools of traditional celestial navigation, for example, enable mariners to calculate their positions, yet they presume both geocentric and geostatic arrangements. So, the fact that geometrical optics remains the standard conceptual tool for many everyday calculations at the macro-level does *not* demonstrate that the theory and its ontology have not been eliminated from the canonical commitments of science.

The Churchlands' comments that "there is no such thing as a literal *light ray*" and that geometrical optics offers a "false model of reality" (and, of course, the fact that they offer this case as a counterexample in the first place) signal the candidacy of geometrical optics, in their view, for just that sort of elimination. Yet everything about their description of this case suggests that they finally regard the reduction of geometrical optics to electromagnetic theory as a prototypical case of a New Wave *approximate (micro-)reduction* of a technically false theory (in the same sense that Mendelian genetics or Newtonian mechanics are technically false theories). What the Churchlands show is that geometrical optics provides an obviously *incomplete* account of optical phenomena, not that it is glaringly false in the way that we regard, say, the phlogiston theory of combustion to be.

To repeat, the point about cases of approximate (micro-)reduction is expressly that the intertheoretic mappings they involve are *comparatively good*. The Church-

lands, after all, emphasize the ability of electromagnetic theory to reduce, i.e., reconstruct in its terms, *"all* [of the relevant] optical phenomena" that geometrical optics surveys — including, incidentally, precisely the phenomena still commonly referred to as "light rays" ([Churchland and Churchland, 1996, 230], emphasis added). There may not be any such thing as a literal light ray, but, quite *unlike* phlogiston, for example, there certainly is some thing to which the term "light ray" refers, viz., a direction perpendicular to that of the relevant waves of electromagnetic radiation.

According to the Churchlands' own description, then, this case falls in the *left* half of the *upper* continuum in figure 5. Apparently, like the classical gas laws, the principles of traditional geometrical optics are unable to match the precision and scope that the corresponding lower level theory provides. Also like the classical gas laws, though, they constitute "a useful tool for *understanding* many *macro*-level effects" ([Churchland and Churchland, 1996, 230], emphasis added). Arguably, conceptual tools that create an *understanding* of effects provide deeper, more valuable, cognitive insights than those which merely predict those effects or facilitate calculations. (This claim is a corollary of the prototype-activation model of explanatory understanding that Paul Churchland defends. See [P. M. Churchland, 1989, ch. 10]; [Churchland and Churchland, 1996, 257–258 and 264].) That geometrical optics manages to do all three testifies to its probity. The salient point here, though, is that *its approximate microreduction* to electromagnetic theory *adds* to that testimony.

Scientists' ability to construct relatively faithful analogues of the principles of traditional geometrical optics within the conceptual framework of electromagnetic theory amounts to even more compelling evidence of its integrity. That traditional geometrical optics — so far as it goes — basically squares with electromagnetic theory is an asset, not a liability. The heightened precision and generality and the added insights (concerning just such things as diffraction, interference, and polarization) in this domain and others that electromagnetic theory brings do not discredit the accomplishments of traditional geometrical optics. Nor did they bring research in geometrical optics, transformed and enhanced by its reductive integration into electromagnetic theory, to a halt. Consider the work of Lord Rayleigh and, for example, the computation of the Rayleigh limit.

Again, like the classical gas laws and Mendelian genetics and *quite unlike the theory of the bodily humors or the chemistry of Stahl*, the basic principles of geometrical optics are construed as broadly continuous with subsequent theory and regularly serve in educational programs for mastering the modern science. Not only does mapping an upper level theory comparatively well on to a successful lower level theory not impugn it, it vindicates it. That such upper level theories can serve as heuristics of calculation settles questions about their *practical* value straight away. To the extent that an upper level theory either has, in the past, provided sound principles that organize phenomenal patterns and offered explanatory insights that are *corroborated* by its reductive relationship with a more inclusive lower level theory or, in the present, continues to offer theoretical guidance and in-

spiration, productive experimental techniques, and bodies of evidence (to which, among others, theoreticians working at lower levels may appeal) establishes its fundamental contribution to the scientific enterprise and should discourage talk of its elimination. That some of these observations are generally of a piece with many of the Churchlands' comments at other points about such cases of approximate reduction (e.g., [P. M. Churchland, 1989, 47–50, 215]) renders their appeal to the case of geometrical optics as a putative counter-example to the explanatory pluralist's claim that the elimination of theories does not occur as the result of interlevel conflicts all the more puzzling.

Endicott argues that New Wave reductionism faces a dilemma with just such cases of approximate (micro-)reduction, since once the co-evolution of theories between levels begins, the notion of independently constructing an analogue of the reduced theory (within the framework of the reducing theory) on which that reduced theory has had no influence is either implausible or will look superficial and artificially abstemious from a historical standpoint. About cases just like the microreduction of geometrical optics, Endicott [1998, 67] comments "on the worst case scenario, new-wave construction is flatly contradicted by co-evolutionary facts; on the best case scenario, it is historically shallow and methodologically restrictive."

Fears about a *science's* elimination or enervation or dispensability on the basis of its alleged (standard or New Wave) reduction are unfounded. Reductions of psychological theories to theories in neuroscience are conceptual bridges that permit intellectual traffic to flow between these two sciences. Constructing such integrative relationships amounts to establishing the sort of infrastructure that facilitates important forms of scientific progress at both levels. Nor does endorsing either the power of reductive strategies in science generally or individual reductive proposals at the border of psychology and neuroscience specifically require any expectations about sciences that operate at higher analytical levels eventually being forced to conduct going-out-of-business sales. It never has been nor is there any reason to think that it ever will be the case that up-and-running sciences collapse on the basis of successful theoretical reductions, or that they could so collapse, or, especially, that they *should* so collapse. *Sciences* are not the sorts of things that get reduced or eliminated. Theories are.

When considering reductive relations between theories in psychology and neuroscience, it is critical to realize that the arguments that have been offered above for those last claims are not the familiar ones that either Anglo-American or Continental philosophers, who fashion themselves friends of subjectivity or intentionality or consciousness or individuality or contextuality or human freedom or the great tradition of humanism, propound. Contemporary naturalists, who are interested in the cross-scientific connections between the psychological and cognitive sciences, on the one hand, and neuroscience, on the other, are not attracted to bare philosophical conjecture in the domain of mind and brain. The Churchlands [1998] defuse five well-worn objections to the reduction of psychology to neurobiology that appeal, respectively, to sensory qualia, intentionality, complexity, freedom,

and multiple realizability. Bickle [1998] employs essentially the same strategy that the Churchlands do to dismantle the multiple realizability objection. Bickle also concurs with Kim [1989] that boosters of supervenience and non-reductive materialism have, unfortunately, misunderstood the inevitable property dualism their positions entail. Bryon Cunningham [2001a; 2001b] provides unified replies to anti-reductionist arguments that look to complexity and qualia, respectively.

The last three sections examine the confluences of views among naturalistically oriented philosophers interested in cross-scientific relationships where interlevel connections seem to promise a co-evolutionary integration of sciences and the eventual development of interlevel and interfield theories [Maull, 1977; Darden and Maull, 1977]. Like the emergence in the middle of the twentieth century of biochemistry at the border of two of the major families of sciences, the blossoming of cognitive neuroscience over the past few decades has signaled the development of an interlevel science at the border between the biological and psychological sciences. The development of brain imaging technologies has only abetted the prospects for even richer integration.

The Churchlands' comments on geometrical optics to the contrary notwithstanding, much of the time New Wave reductionists seem to agree with explanatory pluralists (and those whose work exhibits the morals of explanatory pluralism in the study of complex mechanisms and mechanistic explanation in science). They concur about both the prospects for and the implications of such scientific integration in the areas of the psychology-neuroscience interface that fall in the left half of the upper continuum in figure 5. At the very least, none of these philosophers subscribe to either the multiple realizability objection or the explanatory gap objection to the reductive integration of the psychological and neuroscientific. Their responses to these objections in large part take their inspiration from patterns in the history of research in science and in cognitive neuroscience, in particular.

8 MULTIPLE REALIZABILITY AS AN ARGUMENT FOR (NOT AGAINST) REDUCIBILITY

It is precisely concerning cross-scientific relations where sciences at adjoining levels of analysis offer fruitful explanatory proposals, where interlevel conflicts are neither plentiful nor weighty, and, thus, where intertheoretic mapping is promising and where the co-evolution of theories *and integrative empirical research* are underway (for example, in cognitive neuroscience) that contemporary accounts of reductive integration (whether New Wave or explanatory pluralist) most resonate with one another. In his most recent work even Bickle, "ruthlessly reductive" by self-description, who has argued most steadfastly for increasing deference to lower level theories and explanations as a principle that organizes scientific progress, now lauds the co-evolution of psychology and neuroscience, eschews the elimination of psychology, and recognizes its vital contribution to the characterization of higher cognitive functions ([1998, 141]; [2003, 128 and 130]). The complicated brain activation patterns that brain imaging technologies reveal would, for ex-

ample, be virtually meaningless without theoretical and methodological guidance from psychology [Mundale and Bechtel, 1996]. Bickle also allows that progress in neuroscience, at least, requires methodical research at "multiple levels" and that "psychological causal explanations still play *important heuristic roles* in generating and testing neurobiological hypotheses"[6] [2003, 178 and 114]; cf. [Craver, 2002]).

Because of their attention, first, to particular processes and mechanisms that both psychology and neuroscience address and, second, to the lessons from the history of research and practice in these and other areas of science, virtually all contemporary philosophers, who are at all sympathetic to reductionist projects, remain unimpressed by the two major philosophical objections to the reductive integration of psychology and neuroscience. Among the things that these naturalists hold in common is that the relevant sciences here have progressed to the point where, like the philosophies of physics and biology before them, the philosophies of psychology and neuroscience can no longer afford to prize philosophical cleverness or metaphysical comfort over empirical accountability and explanatory adequacy.

Initially, the two major replies to the multiple realizability objection might seem antithetical, since the first headlines how *often* multiple realizability arises in science, while the second focuses on how *infrequently* it arises at the psychology-neuroscience interface, at least once researchers are clear about the best models available in each science and about how coarse or fine-grained those models are. In fact, the two replies are closely related, with the second reply building on the first. The first reply denies that the anti-reductive conclusion follows from a premise affirming multiple realization. The second reply suggests that the range of *theoretically interesting* realizations may have been greatly exaggerated and that what few exist will prove eminently manageable.

---

[6] On the other hand, it would be easy to exaggerate the resonances between the various philosophical conceptions of cross-scientific relations that hold out promise of extensive integration.

Like the Churchlands' take on geometrical optics, Bickle sometimes seems to lose sight of the support an upper level theory enjoys in an approximate microreduction and of the contributions it can still make to on-going inquiries when he holds that model building ceases at the psychological level and that psychological models lose their causal explanatory status once reductive integration has commenced [2003, 110–111, 114]. So, Bickle argues that in the face of the reduction of the consolidation of declarative memories to the operations of the cellular and molecular mechanisms that drive the shift from E-LTP to L-LTP and that selectively preserve L-LTP in certain synapses [Lynch, 2000], "it seems silly to count psychology's "explanation" of consolidation as "causally explanatory," "mechanistic," or a viable part of any *current* scientific investigation *still worth pursuing*" [2003, 112].

Maurice Schouten and Huib Looren de Jong [1999] argue that such claims misread the on-going contributions to the specification of the underlying mechanisms of functional analyses in psychology (which can, among other things, implicate relations with socio-cultural phenomena external to the organism). They maintain, in effect, that downplaying psychological considerations in this way strikingly misjudges just how long the co-evolution of theories and entire sciences can go on, which, presumably, arises here from an underestimate of the range of psychological questions that are of interest, for example, concerning memory. Rather than worrying about settling attributions of causal responsibility, Schouten and Looren de Jong (like [Bickle, 2003, 31–40] advocate attending to the details of particular mechanisms. They argue that Bickle overplays the import of findings concerning LTP for the reduction of the wide range of findings and patterns that constitute the psychology of memory. (See [Bickle, 2003, 45–46].)

The first reply highlights the *pervasiveness* of lower level multiple realization of higher level phenomena throughout the sciences and, particularly, in cross-scientific contexts that involve uncontroversial reductions. Multiple realizability often arises even in the parade cases of the standard model, such as the reduction of the notion of temperature in classical thermodynamics to an account in terms of the kinetic theory of heat. Technically, though, what the kinetic theory readily reconstructs is temperature-in-a-gas. The Churchlands [1998, 78] note that "in a gas, temperature is one thing; in a solid, temperature is another thing; in a plasma, it is a third; in a vacuum, a fourth; and so on ... this ... just teaches us that there is more than one way in which energy can be manifested at the microphysical level." As Robert Richardson [1979; 1982] has emphasized, reductions in science are generally *domain specific*, and Nagel's comments on the character of the bridge principles linking the predicates of the reduced and reducing theories indicate that he already recognized this.

The goal of this first argument is to domesticate multiple realization. Philosophers should be less impressed with the multiple realizability objection once they appreciate the frequency with which multiple realizability occurs in nature. Although he confines his analyses to the logical empiricists' standard model of reduction, Kim [1989, 39] claims that the domain specific reductions, which result under such circumstances, abound throughout the sciences and not merely in psychology and that they are "reductions enough ... by any reasonable scientific standard and in their philosophical implications."

Still, might the anti-reductionist reply that the multiple realizability of *psychological* states is of a wholly different order? After all, what about Fodor's arguments that disjunctions of possible realizations of psychological states at the physical level need not be very large before they become unmanageable from a practical standpoint and unhelpful from the standpoint of explanation? The anti-reductionist can acknowledge widespread multiple realization in many other cross-scientific contexts but argue that the psychological case brings added problems all of its own.

The most troublesome problems here do not concern the psychologies of aliens or robots, since for many purposes the division of psychology and cognitive science into specialized sub-domains seems plausibly motivated on a variety of criteria in the same way that accounts of heat in gases, solids, plasmas, vacuums, and so on are usefully distinguished for some of our problem solving purposes in physical science [Mundale and Bechtel, 1996, 490]. Not even the multiplicity of non-human organisms' psychological states (recall the hunger of octopuses) presents the biggest challenge here. There may be fewer grounds for sub-dividing psychology by phylum or species in the light of the commonalities of their evolutionary heritages, but both the physical and the (probable) psychological differences look substantial enough in most cases that even this sort of specialization in psychology does not look utterly unmotivated. These forms of multiple realizability present challenges, but they are not insurmountable challenges, since the sub-division of some psychological concepts for some explanatory purposes is no more out-of-bounds than

is the occasional sub-division of the concept of heat in thermodynamics that the Churchlands raised.

The multiple realizability that Fodor's analysis points to is of a more fundamental sort. Sub-dividing psychology along the lines of basic material substrates or according to distinctions among species would circumvent many of the problematic disjuncts that would make for unwieldiness in the bridge principles of a reduction, but its sub-division between individuals or worse yet between the same individual at different times would drain away any possible interest in the putative reduction.[7] As William Bechtel and Jennifer Mundale [1999, 177] note, " ... even within a species brains differ. Even within an individual over time there are differences (neurons die, connections are lost, etc.). Thus, multiple realizability seems to arise within species (including our own) and even within individuals." The rarefied sub-divisions within a reduced psychology that these forms of multiple realization portend surely expose the futility of the reductionist's project. Apparently, psychological insights would be inundated and sink, lying lost along the bottom of vast oceans of neuroscientific detail.

Bechtel and Mundale, however, do not expound on this objection concerning the proliferation of possible neural instantiations of human psychological states in order to praise it, but rather to bury it. Multiple realizability arguments look plausible, first, because anti-reductionists have generally failed to attend to what scientists have ascertained to be the theoretically significant kinds at each analytical level (and especially at the level of neuroscience) and, second, because they have also consistently ignored whether the kinds they do discuss are cast at comparable grains.

On the first front, Bechtel and Mundale [1999, 203] remark that, as Leibniz observed three centuries ago, any two particulars will both resemble and differ from one another in an infinite number of respects, and there is no reason to expect things to be any different when comparing brain states and psychological states. Science is always concerned with ascertaining which resemblances and differences *matter* from the standpoints of explanation, prediction, and control. The aim is not to map each and every homespun category we may employ in these or any other domains, but rather to concentrate on those that our best explanatory theories spotlight [Hardcastle, 1996].

Such considerations may even partially neutralize the sting of Fodor's famous argument about the fruitlessness — for understanding economics — of a focus on the various instantiations of money. Attention to the limitations that particular material forms that money can take impose will disclose some eminently useful, though admittedly low level, generalizations about those forms' deployment within economies. For example, some transactions such as mortgage closings at banks and purchases of items stored in inside pockets of less scrupulous vendors' trench coats in alleys in large cities will almost never involve personal checks or credit cards or, at least at the mortgage closings, large amounts of cash. Thus, some

[7] Davidson [1970] employs different premises but the conclusion he draws, viz., that there are no psycho-physical laws, is functionally equivalent.

patterns in the economic domain may offer grounds for the fragmentation of the concept 'money' along these lines for certain limited, domain specific explanatory purposes.

On the second front, philosophers find ubiquitous multiple realizability in psychology because they regularly compare coarse grained psychological concepts with exceedingly fine-grained conceptions of brain states. The folk psychological notions that particularly interest philosophers are more coarse grained than most employed in experimental cognitive psychology, while the conceptions of brain states they discuss, Bechtel and Mundale argue, are much finer-grained than the ones practicing neuroscientists use in their theories. They comment [1999, 178] that "when a common grain size is insisted on, as it is in scientific practice, the plausibility of multiple realizability evaporates." The point is not that any particular grain is canonical, but only that, once philosophers compare psychological and neuroscientific blueprints of comparable scale, multiple realizability vanishes.

When Putnam [1967] examined hunger in humans and octopuses, his grain for type identifying psychological states was not especially fine. Certainly, such a broad extension of psychological types poses problems for the *functional* identification of psychological states, since the links to other mental states and to behaviors that are central to functional analyses differ profoundly between such radically different species. (For example, as the result of their hunger, octopuses never ponder a quick trip to the supermarket nor do they ever slap some peanut butter and jelly on some bread for a quick snack.) Still, given that evolution tends to conserve and extend existing mechanisms rather than create new ones, researchers could well end up type identifying even the neural mechanisms involved in hunger in the octopus and human, which would substantially defuse Putnam's intuitively plausible example. This is not to rule out the possibility of radically different ways of performing similar functions emerging in evolution. However, the point is precisely that when researchers discover evidence of multiple mechanisms for performing similar functions, such as alternative pathways for processing visual input in invertebrates and vertebrates, it provides an impetus for psychologists to search for functional (behavioral) differences that motivate the differentiation of types at the psychological level as well (cf. [Kim, 1972]).

Ascertaining compatible grains between research at two different levels is one of the most basic steps in the co-evolution of sciences. Getting the grains right between theoretically significant kinds makes all the difference. A variety of *successful research strategies* at the borderline of psychology and neuroscience, some of which have, by now, been *utilized for more than a century*, tacitly repudiate the multiple realization of theoretically relevant psychological states. This is not only true about the interpretation and the integration with models in cognitive neuroscience and cognitive psychology of recent findings from PET and fMRI research, where scientists have obtained generalizable results by employing sophisticated statistical techniques to analyze multiple images of multiple brains. Neuroscientists' inferences about the cognitive functions of various areas in unimpaired brains, on the basis of studies, such as Paul Broca's ([1861]– cited in [Bechtel and Mundale, 1999,

184]) classic work, on performance deficits and brain damage, have simultaneously proceeded unencumbered by worries about multiple realizability and proven one of the most fertile research strategies available.

The force of the first reply to the multiple realizability objection was to regularize it by stressing how often it arises in science. The force of this second reply, which explores the consequences of the first for scientific practice, is to accentuate how rarely multiple realizability presents any barrier to reduction in science, at least once we get clear about both the operative explanatory categories and their comparative grains.

Bechtel and Robert McCauley [1999] push this second reply one crucial step further. They argue for a conclusion that, if sound, not only deflates the anti-reductionists' multiple realizability objection but construes multiple realization as a platform for *defending* a version of the position that the objection was formulated to waylay! Bechtel and McCauley point out that for a host of reasons, beginning with ethical ones, the overwhelming majority of the research and experimentation in the history of neuroscience has been done on the brains of non-human animals. This is to say that neuroscientific research on the identification of brain areas and processes is done *comparatively*.

Korbinian Brodmann [1909/1994] used a variety of criteria to map the human brain into functionally distinguishable areas. These included attention to the gross anatomical features of brains as well as the examination of cortical micro-structure. Brodmann employed cytoarchitectural tools to demonstrate that cortex generally consists of six layers. He distinguished brain areas, in part, on the basis of the relative thickness of these layers (e.g., layer 4 was very thick in areas 1, 2, and 3, but much thinner in area 4) and the particular types of neurons (e.g., pyramidal cells) found in them. The important point is that Brodmann based his account of cortical layers on *comparative studies involving fifty-five species*. Brodmann also proceeded comparatively in his study of neuroanatomical features. In addition to his well-known map of the human cortex (figure 6), Brodmann produced maps for an assortment of other species, including the lemur, flying fox, rabbit, hedgehog, and others.

For Brodmann finding comparable areas in different species despite differences in brain shapes and in the relative location of areas was pivotal in identifying and building a case for functionally distinct areas in the human brain. In his work and that of other neuroscientists concerned with such questions, the multiple realization of some psychological function across species in related but different structures does not obstruct the identification of an area. On the contrary, it is the single most compelling type of evidence available for identifying an area in the human brain! Contrary to contemporary anti-reductionist orthodoxy, multiple realization across species is not a barrier to the mapping of some psychological function on to brains, rather it is the key to accomplishing such mappings.

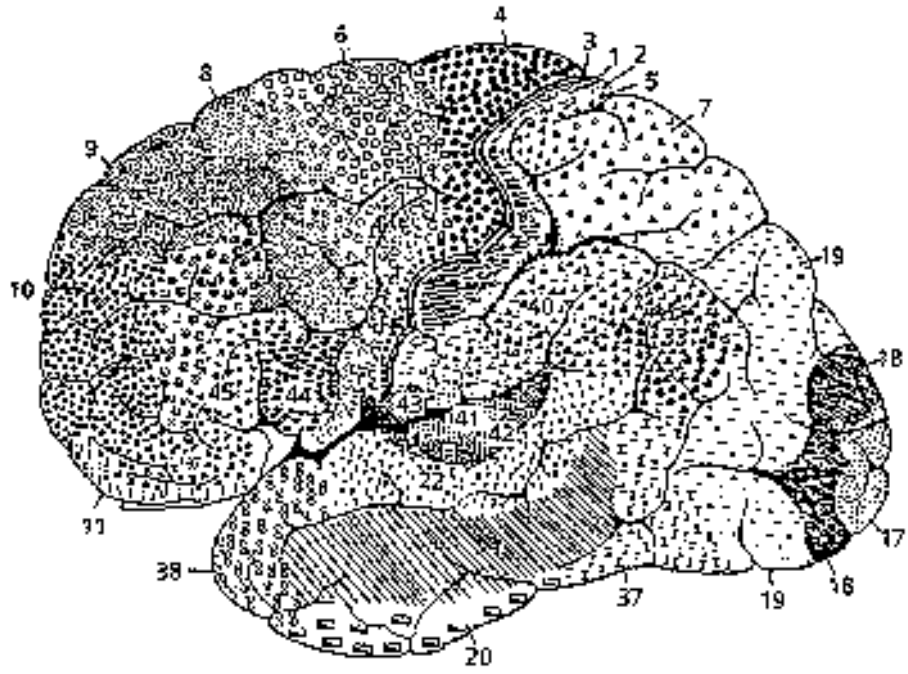

Figure 6.

9 MECHANISTIC ANALYSIS AS EXPLANATORY PLURALISM WRIT SMALL

Recasting multiple realizability as an aid, rather than a barrier, to the integrated development of the psychological and neuroscientific has encouraged some naturalistically oriented metaphysicians to reassess the relative promise of functionalism and the psycho-physical identity theory. This new perspective on multiple realizability suggests that even if the premises that inspire functionalism are true, they block neither the psycho-physical identity theory nor the reductive integration of a good deal of the psychological with the neuroscientific.

Thomas Polger [2004] accepts Bechtel and Mundale's argument about the decisive importance of ascertaining the grain at which various psychological and neuroscientific descriptions are cast when pondering the merits of the identity theory. A perfectly interesting version of the identity theory need only require that humans' mental states are identical to some of their brain states. It does not require that all mental states, especially those cast at coarse grains within psychology, are identical to one another or that all identities of psychological and neural states are cast at a single grain. To repeat, the point is *not* that some grain or other is canonical, but rather that preferences concerning grain in any given case have ev-

erything to do with added empirical accountability, explanatory accomplishment, and promise of productive extensions of research and of cross-scientific consilience — in just about that order of decreasing significance.

Polger argues that, finally, the only version of functionalism that matters for these controversies is one that maintains that functionally distinguishable mental kinds merit an independent metaphysical status. According to Polger [2004, 136], abstractness of its functionally discriminated kinds, relative to the categories deployed in the explanatory theories of biological science, is the critical condition such a version of functionalism must satisfy in order to outstrip the identity theory. He argues that none of the prominent accounts of the functional advanced in the literature can legitimately purchase that biological abstractness without sacrificing one or more of three other necessary conditions for a satisfactory conception that he specifies, viz., causal efficaciousness, objectivity, and synchronicity [2004, ch. 5, esp.177–178].

Polger's proposal [2004, 188] for resolving the complications, which the possibility of multiple realizations of psychological items introduces and which originally motivated functionalism in the philosophy of mind, echoes Bechtel and McCauley's [1999] argument. As with any other cross-scientific case, the conceptual and theoretical resources of psychology and neuroscience provide a variety of different sub-levels and associated grains at which the relevant events, states, structures, systems, and processes may be cast. Identity theorists have every right to search for patterns and mechanisms at different levels and grains on a case by case basis. That is in the service of sub-dividing these phenomena according to the lights of the best explanatory theories available. Polger holds that the considerations that most philosophers seem to think point toward functionalism will, in fact, just as readily square with a version of the psycho-physical identity theory. What he, ultimately, seems to anticipate is a theory of mind that is informed by the lessons of the explanatory pluralist's picture of cross-scientific relations.

Polger is clear that any version of the identity theory he envisions is not one that turns on consummating either standard or New Wave reductions.[8] He turns for inspiration, instead, to the recent literature in the philosophy of science on mechanisms and mechanistic explanation. (See [Bechtel and Richardson, 1993]; [Machamer, Darden, and Craver, 2000]; [Craver, 2001]; [Craver and Darden, 2001], and Wright and Bechtel's "Mechanism" in this volume.) This work has yielded useful tools for dealing with what would, from the standpoint of traditional discussions of intertheoretic reduction, seem to be the more rarified distinctions between analytical levels in the sciences. Focusing on the details of particular mechanisms, they offer a bottom-up approach that suggests, if anything, delineating analytical levels on a case-by-case basis. (See too [Bickle, 2003, 116–117].) These accounts of mechanisms in science mostly leave it to philosophers of science interested in the

[8]Because Kim's discussion [1998, ch. 4] does not explore recent accounts of cross-scientific relations, explanatory pluralism, or mechanistic analysis in the philosophy of science or of ongoing research programs in the relevant sciences, it only sketches in the most abstract terms a similar sort of negative case concerning the promise of anti-reductive materialism.

more traditional questions of reductionism to worry about whether these instances will serve up patterns capable of supporting new generalizations about analytical levels at these finer resolutions.

These analyses of mechanisms exemplify the general morals of the explanatory pluralism that, according to the model of intertheoretic relations that has emerged across the previous four sections, prevails in nearly all cross-scientific settings. (See figure 5.) They offer, in effect, an explanatory pluralism writ *small*. They provide *multi-level* causal explanations that "... explain by showing *how* an event fits into a *causal* nexus" [Craver, 2001, 68]. It is the discovery and the delineation of the mechanism and its operations that reveals what researchers are willing to count as causal.

Bechtel [1986b] and Carl Craver [2001, 62–68] argue that analyses of mechanisms require inquiries that adopt at least three different perspectives. Minimally, the analysis of mechanisms will include and integrate studies of them and their operations as *isolated*, as *constituted*, and as *situated*. Craver prefers to distinguish these three perspectives from fixed levels of organization in nature: "these are three different perspectives on [an] . . . . item's activity in a hierarchically organized mechanism; they are not levels of nature" [2001, 66]. To stress the salience of the details of particular cases and to be (justifiably) wary about drawing any generalizations on the basis of particular cases do not, however, impugn their relevance to judgments about the finer grained analytical levels in science that have traditionally informed discussions of reduction. That mechanistic analyses concentrate on the local does not preclude their ability to illuminate reductionists' concerns about analytical levels.

To study a mechanism *as an isolated system* involves formulating a hypothesis about the system's borders, offering a rough and ready characterization of the mechanism's activities, and investigating the character, frequencies, locations, etc. of the inputs and outputs that cross those borders. Descriptions of mechanisms in isolation set constraints on approaches that examine their constituents.

The study of mechanisms *as constituted*, of course, is the most obvious point of contact with traditional discussions of microreduction, its concern with mereological relationships, and traditional concerns about levels of organization in nature. Whether the components in question are items treated in theories at a recognizably lower level usually depends on the scale of the breakdown. Functionally speaking, though, their study, regardless of the scale of the breakdown, exploits a more fine-grained perspective. According to both explanatory pluralism and Bechtel and Craver's analyses of mechanisms, reductive explanatory strategies are a fundamental part of the explanatory story in science, but they are not the entire story.

As noted earlier, complex systems can also be studied in isolation, but they rarely, if ever, operate in isolation. Examining a mechanism's performance in the larger environment of which it is a part provides explanatory insights of a different order. In context, the mechanism is now construed as part of a larger economy, which is studied at a functionally higher analytical level. Studies of mechanisms *as*

situated provide information about the roles that the mechanism as a whole plays in a larger dynamic system, about the sources of its inputs and the recipients of its outputs, and about other mechanisms that are capable of producing or influencing its inputs or outputs. Scrutinizing mechanisms in context commonly calls for appraisals of the functions of relationships and, at least from the biological level on up, of cooperative, competitive, and selective considerations especially.

Neither explanatory pluralism, writ large, nor these mechanistic analyses Polger cites leave any room for the dire or dismissive conclusions New Wave reductionists sometimes draw about intertheoretic relations in *interlevel settings*. All of the complex mechanisms and systems that populate the biological, psychological, cognitive, and socio-cultural sciences are, from the standpoint of explanation, greater than the sums of their parts and demand study as isolated, as constituted, and as situated. Such multi-level study not only offers a richer account of these particular mechanisms and their operations, it also enhances our understanding of all of the systems engaged in a mechanistic hierarchy. After all, these three perspectives are always *relative* to the particular mechanism that is the focal object of study. The strategies and pursuits at each level regularly yield findings that are mutually reinforcing [Craver, 2002]. Because they countenance the full range of possible cross-scientific relationships, these versions of explanatory pluralism do not confine themselves to reductive analyses only. Reductive explanation is a valuable contributor, but these approaches also embrace forms of non-reductive explanation as well [Polger, 2004, 205–209]. For the explanatory pluralist, all explanations are partial explanations; all explanations are from some perspective, and all explanations are motivated by and respond to specific problems.

10 HEURISTIC IDENTITY THEORY AND THE EXPLANATORY GAP OBJECTION

Discovering and explaining mechanisms proceeds piecemeal. Parallel research at multiple levels leads to increasingly developed views of mechanisms' organization and operations. Nothing intrinsic to analyses of mechanisms entails any particular ontological commitments, and Polger is correct that the models of these philosophers of science are "neutral about the nature of the entities that figure in mechanistic explanations" [2004, 209]. Nothing, however, follows about the implicit ontological commitments of scientists' specific mechanistic proposals. Progressively more integrated accounts of the structure and functioning of mechanisms yields the increasingly more articulated connections between models and theories at different analytical levels in science that animate philosophical naturalists' ontological calls.

Hypothesizing about cross-scientific identities serves as one of the principal heuristics for promoting such integrative research and for provoking discoveries at both of the explanatory levels involved. Polger [2004, 210] points out that "the identity theory in fact has more explanatory resources than functionalism because it makes use of both contextual and constitutive explanations." Such hypothetical

identities are common means for enabling scientists working at one analytical level to explore and exploit the conceptual, theoretical, methodological, and evidential resources available at another. The primary motives that drive the initial formulation of such hypotheses, whatever their eventual ontological consequences, concern their capacities to advance empirical research.

The logic behind their use looks to the converse of Leibniz's law. Instead of appealing to the identity of indiscernables, this strategy capitalizes on the indiscernability of identicals. What is known about an entity or process under one description should apply to it under its other descriptions. Scientists do not advocate hypothetical identities because the two characterizations *currently* mirror one another perfectly. On the contrary, it is just because the characterizations do not seem to mirror one another perfectly that the hypotheses are of interest. The theories at each level ascribe features to the entities and processes the interlevel, hypothetical identities connect. Scientists check the applicability of the feature lists at the alternative levels of analysis, using related research at each explanatory level to stimulate discovery at the other.

If unresolved conflicts persist, scientists pursue further empirical inquiry to ascertain either which characterization should prevail or in what directions each needs to evolve. *That* research yields more narrow hypotheses about the systems and patterns engaged. Within their respective levels, such speculations suggest new ways of organizing old theoretical commitments and familiar facts and point to new directions for empirical investigation. Finally, some of the new arrangements that result, almost inevitably, will produce some new cross-scientific conflicts, which will likely begin this cycle anew.

If, on the other hand, these hypotheses meet with much success at all, they will lead even more directly to the development and elaboration of more extended proposals about the connections between the two explanatory levels. Crucially, scientists accept or reject these hypotheses for the same reasons that they accept or reject any other hypotheses in science, viz., their abilities to stand up to empirical evidence, to stimulate new research, and to foster the integration of existing knowledge.

For example, in a relatively brief period in the middle of the twentieth century, research integrating behavioral and physiological techniques had defeated the hypothesized identification of the visual center with V1 exclusively. However, rather than undermining the strategy of hypothesizing identities, determining the functions of cells in V1 by David Hubel and Thorsten Wiesel [1962] led to more hypothetical identities that were even more precise about the visual and the neural processes identified and about the additional brain areas involved. Beginning with their landmark work and continuing for the next three decades, the ongoing interplay of behavioral and neurophysiological research led to repeated revisions of the hypotheses about the brain areas implicated and about the conception of the information processing performed in vision. This co-evolutionary process not only preserved a plurality of explanatory perspectives but resulted in a refinement of both psychological and neural models. For over a century now, exploring the

empirical merits of the hypothesized identity of the visual center and the occipital lobe has generated both more and more detailed hypotheses about the activities in the brain with which various aspects of visual processing and experience should be identified. Moreover, these findings about the various brain areas involved in visual processing played a key role in inspiring new, ambitious, higher-level hypotheses at the neuropsychological and psychological levels about visual processing and the organization of the human cognitive system overall. Probably the most famous is Ungerleider and Mishkin's [1982] identification of two processing streams for visual information in the brain.

Bechtel and McCauley appropriate these lessons from the philosophy of science and from their instantiations at the interface of psychology and neuroscience, in particular, to formulate both a new version of the psycho-physical identity theory and replies to the multiple realizability objection (scouted at the end of section 8) and to the explanatory gap objection ([Bechtel and McCauley, 1999]; [McCauley and Bechtel, 2001]). According to their Heuristic Identity Theory (HIT) psychoneural identities are not the conclusions of scientific research but the hypothetical premises. The preceding discussions of explanatory pluralism and mechanistic analysis show why hypothetical identities of psychological and neural processes generate both new hypotheses and new avenues of research that serve to direct those hypotheses' development and elaboration. The differences between theories at these two levels encourage scientists to consider adjustments to their conceptions of the pertinent processes and structures in a reciprocal process of mutual fine-tuning. By way of illustration, Bechtel and McCauley review the history of proposals about the locus of visual processing in the brain from the late nineteenth through the late twentieth centuries (briefly touched upon above).

The way that HIT construes psycho-physical identities suggests a framework for responding to the objection to the reduction of psychology (and to the identity theory) that appeals to an explanatory gap concerning consciousness. On HIT's account of things, finally, an explanatory gap at the interface of psychology and neuroscience, whatever its basis, is simply an instance of a failure of reductive integration between two sciences operating at adjacent analytical levels. The gap will be closed the same way that other cross-scientific gaps have been closed in the history of science, viz., through the co-evolutionary integration of the sciences in the course of on-going cross-scientific research.

Kim nicely summarizes the general argument informing the objection[9]: "*it is*

[9]David Chalmers' version of the argument [1996, 115 goes as follows:

> Neurobiological approaches to consciousness ... can ... tell us something about the brain processes that are *correlated* with consciousness. But none of these accounts explains the correlation: we are not told why brain processes should give rise to experience at all. From the point of view of neuroscience, the correlation is simply a brute fact.
>
> ... Because these theories gain their purchase by *assuming* a link ... it is clear that they do nothing to explain that link.

For an extended reply to this specific formulation of the explanatory gap objection, see [McCauley and Bechtel, 2001].

the explanation of . . . bridge laws, an explanation of why there are just these mind-body correlations, that is at the heart of the demand for an explanation of mentality . . . it is evident that the Nagel reduction of psychology is like taking mind-body supervenience as an unexplained brute fact" [1998, 96]. Within the objection lurk two challenges: (1) the identity theory and reductive materialism need to explain the identities they propose and (2) any evidence that can be cited in support of such an explanation is also perfectly consistent with affirming no more than psychoneural correlations (and, thus, anti-reductionists fault both positions for their metaphysical presumption). The specific complaint about an explanatory gap concerning consciousness maintains that physicalist accounts provide no explanation, in particular, of how something psychic could *just be* something physical.

The second challenge amounts to arguing that, from the standpoint of the logic of confirmation, claims about the identity of two things are indistinguishable from claims about their correlation. As Kim [1966, 227] has put the objection: " . . . the factual content of the identity statement is exhausted by the corresponding correlation statement . . . There is no conceivable observation that would confirm or refute the identity but not the associated correlation." If the whole philosophical story about proposing psycho-physical identities were one about their confirmation, then this perfectly uncontroversial logical point would carry the day. HIT, however, maintains that cross-scientific hypothetical identities between psychological and neural processes are part of a multi-level scientific investigation of human mentality that involves much more than connections between theories' ontologies (or their confirmation). Claims about correlations and claims about interlevel identities are different conceptual animals that thrive in different theoretical habitats. Not only does this part of the objection overlook the fundamental contribution hypothetical identities make to scientific *discovery*, it does not even get the role of these identities right in the *justification* of scientific theories. Unlike merely noting correlations, advancing hypothetical identities occasions explanatory connections that *demand* empirical exploration. Cross-scientific identities make evidence available from other explanatory levels, and, as noted above, they disclose avenues of research for generating new evidence as well. Their critical contribution resides in their abilities to provoke and refine theories at both of the levels engaged. *Their success at this task is their vindication — not* the accumulation of some sort of evidence that would rule the corresponding correlation claim out of court. The whole point of the correlation objection is precisely that such evidence cannot exist!

HIT simply denies the assumption underlying the first challenge. Identities are not the sorts of things that ever need explanation. What matters about hypothetical cross-scientific identities is not how they should be explained (they can't be) but what they explain, how they suggest (and contribute to) other, empirically successful, explanatory hypotheses, and how they create opportunities for scientists who work at one explanatory level to enlist methods and evidence from alternative levels of explanation. *That* is why "we are not told why brain processes should give

rise to experience . . . " [Chalmers, 1996, 115]. Scientists show why some mechanism *constitutes* some phenomenon by exploring the empirical success of the wide range of predictions and explanatory connections that assumption generates. It is that empirical success that corroborates the constitutive hypothesis and tentatively justifies its assumption [Churchland and Churchland, 1998, 120–122]. But, of course, the tentativeness here is nothing special. It is the same tentativeness about justification that accompanies every empirical claim in science.

HIT underscores the fact that evaluations of proposed identities do not turn on confirming them directly. What, after all, could *that* possibly be [McCauley, 1981]? In empirical matters the evidence for an identity claim arises *indirectly* — primarily on the basis of the emerging empirical support for the explanatory hypotheses it informs. For example, if normal activities in V4 are identical with the processing of information about wave length, then serious abnormalities of particular types in the structure and functioning of V4 should yield abnormalities of particular types in subjects' color experiences. The point is that this hypothetical identity is an empirical conjecture that researchers can use both psychological and neuroscientific evidence not only to assess but to refine. Obtaining indirect corroborating evidence for identifying some neural process with some psychological function along such lines no more finalizes that identity than it would any other hypothesis in science. Nor does it establish that the function under scrutiny is either the sole or even the primary function these neural processes carry out. (So, in fact, whether V4 is even primarily concerned with the processing of color is a point of some controversy among researchers.) Moreover, all research of this sort is limited by scientists' abilities both to conceive of what stimuli might provoke responses in a neural area and to test those conceptions. Still, the more hypotheses of this sort the identity informs and the more successful those hypotheses prove, the more likely the identity will come to serve as an assumption the sciences lean upon rather than a bare conjecture in search of support [Van Gulick, 1997]. Such identity claims are, of course, no less conjectures still. They are, however, no longer simply *bare* conjectures. Nor, manifestly, are the identities that HIT surveys "brute," contrary both to other versions of the psycho-physical identity theory and to virtually all of these anti-reductionist critics.

Those critics miss both the sorts of considerations that motivate hypothetical identities in science and their fundamental contribution to the development of scientific explanations. An emphasis on the multi-level character of scientific research in the sciences of the mind/brain does not bar the explanatory pluralist from embracing type-identities of suitable granularity between mental processes and brain processes. The explanatory and predictive progress that such hypothetical identities promote is the best reason available for acknowledging such cross-scientific connections.

BIBLIOGRAPHY

[Bartlett, 1923] F. C. Bartlett. *Remembering*. Cambridge: Cambridge University Press, 1923.

[Bechtel, 1986a] W. Bechtel, ed. *Integrating Scientific Disciplines*. The Hague: Martinus Nijhoff, 1986.

[Bechtel, 1986b] W. Bechtel. Teleological Functional Analyses and the Hierarchical Organization of Nature, *Teleology and Natural Science*. N. Rescher (ed.). Landham, MD: University Press of America, 26-48, 1986.

[Bechtel and McCauley, 1999] W. Bechtel and R. N. McCauley. Heuristic Identity Theory (or Back to the Future): The Mind-Body Problem Against the Background of Research Strategies in Cognitive Neuroscience, *Proceedings of the Twenty-First Meeting of the Cognitive Science Society*. M. Hahn and S. C. Stones (eds.). Mahway, New Jersey: Lawrence Erlbaum Associates, 67-72, 1999.

[Bechtel and Mundale, 1999] W. Bechtel and J. Mundale. Multiple Realizability Revisited: Linking Cognitive and Neural States, *Philosophy of Science* 66, 175-207, 1999.

[Bechtel and Richardson, 1993] W. Bechtel and R. C. Richardson. *Discovering Complexity: Decomposition and Localization as Strategies in Scientific Research*. Princeton: Princeton University Press, 1993.

[Bickle, 1998] J. Bickle. *Psychoneural Reduction: The New Wave*. Cambridge, MA: MIT Press, 1998.

[Bickle, 2003] J. Bickle. *Philosophy and Neuroscience: A Ruthlessly Reductive Account*. Dordrecht: Kluwer Academic Publishers, 2003.

[Block, 1997] N. Block. AAnti-Reductionism Slaps Back,@ *Philosophical Perspectives* 11, 107-132, 1997.

[Bower, 1972] G. H. Bower. Stimulus-Sampling Theory of Encoding Variability, *Coding Processes in Human Memory*. A. W. Melton and E. Martin (eds.). Washington: V.H. Winston & Sons, 1972.

[Broca, 1861] P. Broca. ARemarque sur le Siêge de la Faculté Suivies d=une Observation "Aphémie", *Bulletins de la Société Anatomique de Paris* 6, 343-357, 1861.

[Brodman, 1909/1994] K. Brodmann. *Vergleichende Lokalisationslehre der Grosshirnrinde*. L. J. Garvey, trans. Leipzig: J. A. Barth, 1909/1994.

[Causey, 1977] R. Causey. *Unity of Science*. Dordrecht: Reidel, 1977.

[Chalmers, 1995] D. Chalmers. Facing Up to the Problem of Consciousness, *Journal of Consciousness Studies* 2, 200-219, 1995.

[Chalmers, 1996] D. Chalmers. *The Conscious Mind: In Search of a Fundamental Theory*. New York: Oxford University Press, 1996.

[Churchland, 1979] P. M. Churchland. *Scientific Realism and the Plasticity of Mind*. Cambridge: Cambridge University Press, 1979.

[Churchland, 1989] P. M. Churchland. *A Neurocomputational Perspective: The Nature of Mind and the Structure of Science*. Cambridge: The MIT Press, 1989.

[Churchland and Churchland, 1996] P. M. Churchland and P. S. Churchland. Replies from the Churchlands, *The Churchlands and Their Critics*. R. N. McCauley (ed.). Oxford: Blackwell Publishers, pp. 217-310, 1996.

[Churchland and Churchlan, 1998] P. M. Churchland and P. S. Churchland. Intertheoretic Reduction: A Neuroscientist's Field Guide, *On the Contrary*. Cambridge, MA: MIT Press, 65-79, 1998.

[Churchland, 1983] P. S. Churchland. Consciousness: The Transmutation of a Concept, *Pacific Philosophical Quarterly* 64, 80-93, 1983.

[Churchland, 1986] P. S. Churchland. *Neurophilosophy*. Cambridge: The MIT Press, 1986.

[Churchland, 1988] P. S. Churchland. Reduction and the Neurological Basis of Consciousness, *Consciousness in Contemporary Science*. A. J. Marcel, and E. Bisiach (eds.). Oxford University Press, Oxford, 273-304, 1988.

[Churchland, 2002] P. S. Churchland. *Brain-Wise: Studies in Neurophilosophy*. Cambridge: MIT Press, 2002.

[Churchland and Sejnowski, 1992] P. S. Churchland and T. J. Sejnowski. *The Computational Brain*. Cambridge: MIT Press, 1992.

[Craver, 2001] C. F. Craver. Role Functions, Mechanisms, and Hierarchy, *Philosophy of Science* 68, 53-74, 2001.

[Craver, 2002] C. F. Craver. AInterlevel Experiments and Multilevel Mechanisms in the Neuroscience of Memory,@ *Philosophy of Science*, supplement to volume 69, S83-S97, 2002.

[Craver and Darden, 2001] C. F. Craver and L. Darden. Discovering Mechanisms in Neuroscience: The Case of Spatial Memory, *Theory and Method in Neuroscience*. P. Machamer, R. Grush, and P. McLaughlin (eds.). Pittsburgh, PA: University of Pittsburgh Press, 2001.

[Cummins, 2000] R. Cummins. 'How Does It Work?' versus 'What Are the Laws?': Two Conceptions of Psychological Explanation, *Explanation and Cognition*. F. Keil and R. Wilson (eds.). Cambridge: MIT Press, 117-144, 2000.

[Cunningham, 2001a] B. Cunningham. The Reemergence of 'Emergence', *Philosophy of Science*, 68 (Supplement–Proceedings), S62-S75, 2001.

[Cunningham, 2001b] B. Cunningham. Capturing Qualia: Higher-Order Concepts and Connectionism, *Philosophical Psychology* 14, 29-41, 2001.

[Darden and Maull, 1977] L. Darden and N. Maull. Interfield Theories, *Philosophy of Science* 44, 43-64, 1977.

[Davidson, 1970] D. Davidson. Mental Events, *Experience and Theory*. L. Foster and J. Swanson (eds.) Amherst, MA: University of Massachusetts Press, 79-101, 1970.

[Ebbinghaus, 1964/1885] H. Ebbinghaus. *Memory: A Contribution to Experimental Psychology*. New York: Dover, 1964/1885.

[Farah *et al.*, 1998] M. J. Farah, K. D. Wilson, H. M. Drain and J. R. Tanaka. What is 'Special' about Face Perception? *Psychological Review* 105, 482-498, 1998.

[Ferier, 1876] D. Ferrier. *The Functions of the Brain*. London: Smith, Elder, and Company, 1876.

[Feyerabend, 1962] P. K. Feyerabend. Explanation, Reduction, and Empiricism, *Minnesota Studies in the Philosophy of Science, Volume III*. H. Feigl and G. Maxwell (eds.). Minneapolis: University of Minnesota Press. 28-97, 1962.

[Flanagan, 1992] O. Flanagan. *Consciousness Reconsidered*. Cambridge: The MIT Press, 1992.

[Fodor, 1974] J. A. Fodor. Special Sciences (or: The Disunity of Science as a Working Hypothesis), *Synthese* 28, 97-115, 1974.

[Fodor, 1975] J. A. Fodor. *The Language of Thought*. New York: Thomas Y. Crowell Company, 1975.

[Fodor, 1981] J. A. Fodor. The Mind-Body Problem, *Scientific American* 244, 114-123, 1981.

[Fodor, 1983] J. A. Fodor. *The Modularity of Mind*. Cambridge: The MIT Press, 1983.

[Fodor and Pylyshyn, 1988] J. A. Fodor and Z. W. Pylyshyn. Connectionism and Cognitive Architecture: A Critical Analysis. *Cognition* 28, 3-71, 1988.

[Glenberg, 1976] A. M. Glenberg. Monotonic and Nonmonotonic Lag Effects in Paired-Associate and Recognition Memory Paradigms, *Journal of Verbal Learning and Verbal Behavior* 15, 1-16, 1976.

[Hardcastle, 1996] V. G. Hardcastle. *How to Build a Theory in Cognitive Science*. Albany: SUNY Press, 1996.

[Henschen, 1893] S. Henschen. On the Visual Path and Centre, *Brain* 16, 170-180, 1893.

[Hirst and Gazzaniga, 1988] W. Hirst and M. S. Gazzaniga. Present and Future of Memory Research and Its Applications, *Perspectives in Memory Research*. M. S. Gazzaniga (ed.). Cambridge: The MIT Press, 1988.

[Hooker, 1981] C. Hooker. Towards a General Theory of Reduction, *Dialogue* 20: 38-59, 201-36, 496-529, 1981.

[Hubel and Wiesel, 1962] D. H. Hubel and T. N. Wiesel. Receptive Fields, Binocular Interaction and Functional Architecture in the Cat's Visual Cortex, *Journal of Physiology* (London) 160, 106-154, 1962.

[Jackson, 1982] F. Jackson. Epiphenomenal Qualia, *Philosophical Quarterly* 32, 127-36, 1982.

[Jackson, 1986] F. Jackson. What Mary Didn't Know, *Journal of Philosophy* 83, 291-95, 1986.

[Jacoby, 1978] L. L. Jacoby. On Interpreting the Effects of Repetition: Solving a Problem Versus Remembering a Solution, *Journal of Verbal Learning and Verbal Behavior* 17, 649-667, 1978.

[Kanwisher *et al.*, 1997] N. Kanwisher, J. McDermott, and M. M. Chun. AThe Fusiform Face Area: A Module in Human Extrastriate Cortex Specialized for Face Perception,@ *Journal of Neuroscience* 17, 4302-4311, 1997.

[Kemeny and Oppenheim, 1956] J. Kemeny and P. Oppenheim. On Reduction, *Philosophical Studies* 7, 6-19, 1956.

[Kim, 1966] J. Kim. On the Psycho-Physical Identity Theory, *Materialism and the Mind-Body Problem*. D. Rosenthal (ed.). Englewood Cliffs, New Jersey: Prentice-Hall, 80-95, 1966.

[Kim, 1972] J. Kim. Phenomenal Properties, Psychophysical Laws, and the Identity Theory, *Monist* 56, 177-192, 1972.

[Kim, 1989] J. Kim. The Myth of Nonreductive Materialism, *Proceedings of the American Philosophical Association* 63, 3, pp. 31-47, 1989.

[Kim, 1998] J. Kim. *Mind in a Physical World: An Essay on the Mind-Body Problem and Mental Causation*. Cambridge: MIT Press, 1998.

[Kuhn, 1970] T. Kuhn. *The Structure of Scientific Revolutions* (second edition). Chicago: University of Chicago Press, 1970.

[Levine, 1983] J. Levine. Materialism and Qualia: The Explanatory Gap, *Pacific Philosophical Quarterly* 64, 354-61, 1983.

[Levine, 1977] J. Levine. On Leaving Out What It's Like, *The Nature of Consciousness: Philosophical Debates*. N. Block, O. Flanagan, and G. Güzeldere (eds.). Cambridge: MIT Press, 1977.

[Looren de Jong, 1997] H. Looren de Jong. Levels: Reduction and Elimination in Cognitive Neuroscience, *Problems of Theoretical Psychology*. C. W. Tolman, F. Cherry, R. van Hezewijk, I. Lubek (eds.). New York: Captus Press, pp. 165-172, 1997.

[Machamer *et al.*, 2000] P. Machamer, L. Darden, and C. F. Craver. Thinking about Mechanisms, *Philosophy of Science* 67, 1-25, 2000.

[Maull, 1977] N. Maull. Unifying Science Without Reduction, *Studies in History and Philosophy of Science* 8, 143-62, 1977.

[McCauley, 1981] R. N. McCauley. Hypothetical Identities and Ontological Economizing: Comments on Causey's Program for the Unity of Science, *Philosophy of Science* 48, 218-27, 1981.

[McCauley, 1986] R. N. McCauley. Intertheoretic Relations and the Future of Psychology, *Philosophy of Science* 53, 179-99, 1986.

[McCauley, 1996] R. N. McCauley. Explanatory Pluralism and the Coevolution of Theories in Science, *The Churchlands and Their Critics*. R. N. McCauley (ed.). Oxford: Blackwell Publishers, 17-47, 1996.

[McCauley and Bechtel, 2001] R. N. McCauley. and W. Bechtel. Explanatory Pluralism and The Heuristic Identity Theory, *Theory and Psychology* 11, pp. 738-761, 2001.

[McGaugh, 2000] J. McGaugh. Memory: A Century of Consolidation, *Science* 287, 248-251, 2000.

[McGinn, 1991] C. McGinn. *The Problem of Consciousness*. Oxford: Blackwell Publishers, 1991.

[Mishkin *et al.*, 1983] M. Mishkin, L. G. Ungerleider, and K. A. Macko. Object Vision and Spatial Vision: Two Cortical Pathways, *Trends in Neurosciences* 6, 414-417, 1983.

[Mundale, 2001] J. Mundale. Neuroanatomical Foundations of Cognition: Connecting the Neuronal Level with the Study of Higher Brain Areas, *Philosophy and the Neurosciences: A Reader*. W. Bechtel, P. Mandik, J. Mundale, and R. S. Stufflebeam (eds.). Oxford: Blackwell Publishers, 37-54, 2001.

[Mundale and Bechtel, 1996] J. Mundale and W. Bechtel. Integrating Neuroscience, Psychology, and Evolutionary Biology through a Teleological Conception of Function, *Minds and Machines* 6, 481-505, 1996.

[Nagel, 1961/1979] E. Nagel. *The Structure of Science*. New York: Harcourt, Brace and World / Indianapolis: Hackett, 1961/1973.

[Neisser, 1967] U. Neisser. *Cognitive Psychology*. New York: Appleton-Century-Crofts, 1967.

[Oppenheim and Putnam, 1958] P. Oppenheim and H. Putnam. Unity of Science as a Working Hypothesis, *Minnesota Studies in the Philosophy of Science–Volume II*, 1958.

[Polger, 2004] T. Polger. *Natural Minds*. Cambridge: MIT Press, 2004.

[Putnam, 1967/1975] H. Putnam. Psychological Predicates, *Art, Mind, and Religion*. W. Capitan and D. Merrill (eds.). Pittsburgh: University of Pittsburgh Press / reprinted in *Mind, Language and Reality: Philosophical Papers* (volume 2). New York: Cambridge University Press, 1967/1975.

[Richardson, 1979] R. Richardson. Functionalism and Reductionism, *Philosophy of Science* 46, 533-558, 1979.

[Richardson, 1982] R. Richardson. How *Not* to Reduce a Functional Psychology, *Philosophy of Science* 49, 125-37, 1982

[Schaffner, 1967] K. Schaffner. Approaches to Reduction, *Philosophy of Science* 34, 137-47, 1967.

[Schaffner, 1992] K. Schaffner. Philosophy of Medicine, *Introduction to the Philosophy of Science*. M. Salmon, J. Earman, C. Glymour, J. Lennox, P. Machamer, J. McGuire, J. Norton, W. Salmon, and K. Schaffner. Englewood Cliffs, New Jersey: Prentice Hall, 310-344, 1992.

[Schouten and Looren de Jong, 1999] M. K. D. Schouten and H. Looren de Jong. Reductionism, Elimination, and Levels: The Case of the LTP-Learning Link, *Philosophical Psychology* 12, 237-262, 1999.

[Sejnowski and Rosenberg, 1987] T. J. Sejnowski and C. R. Rosenberg. Parallel Networks that Learn to Pronounce English Text, *Complex Systems* 1, 145-168, 1987.

[Sejnowski and Rosenberg, 1988] T. J. Sejnowsk and C. R. Rosenberg. Learning and Representation in Connectionist Models, *Perspectives in Memory Research*. M. Gazzaniga (ed.). Cambridge: The MIT Press. 135-78, 1988.

[Simon, 1969] H. Simon. *The Sciences of the Artificial*. Cambridge: MIT Press, 1969.

[Thagard, 1992] P. Thagard. *Conceptual Revolutions*. Princeton: Princeton University Press, 1992.

[Ungerleider and Mishkin, 1982] L. G. Ungerleider and M. Mishkin. Two cortical visual systems. D. J. Ingle, M. A. Goodale, and J. W. Mansfield (eds.), *Analysis of Visual Behavior*. Cambridge, MA: MIT Press, 549586, 1982.

[van Essen and Gallant, 1994] D. C. van Essen and J. L. Gallant. Neural Mechanisms of Form and Motion Processing in the Primate Visual System, *Neuron* 13, 1-10, 1994.

[Van Gulick, 1997] R. Van Gulick. Understanding the Phenomenal Mind: Are We All Just Armadillos? *The Nature of Consciousness: Philosophical Debates*. N. Block, O. Flanagan, and G. G⟨zeldere (eds.). Cambridge: MIT Press, 1997.

[Wimsatt, 1974] W. Wimsatt. Complexity and Organization, *PSA-1972*. K. Schaffner and R. Cohen (eds.). Dordrecht: Reidel, 67-82, 1974.

[Wimsatt, 1976] W. Wimsatt. Reductionism, Levels of Organization, and the Mind-Body Problem, *Consciousness and the Brain*. G. Globus, G. Maxwell, and I. Savodnik (eds.). New York: Plenum Press. 205-67, 1976.

[Wimsatt, 1978] W. Wimsatt. Reduction and Reductionism, *Current Problems in Philosophy of Science*. P. Asquith and H. Kyburg (eds.). East Lansing: Philosophy of Science Association. 1-26, 1978.

[Wimsatt, 1986] W. Wimsatt. Forms of Aggregativity, *Human Nature and Natural Knowledge*. M. Wedin (ed.). Dordrecht: Reidel, 259-293, 1986.

[Wimsatt, 1997] W. Wimsatt. Aggregativity: Reductive Heuristics for Finding Emergence, *Philosophy of Science* 64, S372-S384, 1997.

PERCEPTION PREATTENTIVE AND PHENOMENAL

Austen Clark

The conundrums of phenomenal character and consciousness have often motivated philosophers to study perception; and a particular area of study that will prove worthy of their attention is research into what are called "early" or "preattentive" perceptual processes. These are, roughly, processes that start at the transducers and end where selective attention has access to the results, and can select some favored few for further processing. These early processes are the ones most likely to be called "sensory"; they are at any rate simpler, and they make their appearance earlier than, the more sophisticated states that underlie perceptual judgments. If non-conceptual representation is employed anywhere in the system, it would be employed here. Animals that cannot muster the words for a perceptual judgement can nevertheless sense things. These simple sensory states are dear to the heart of those interested in phenomenal character, and I hope they find their interests piqued by preattentive phenomena.

To fit constraints of time and space this paper must set aside consideration of relations between perception and knowledge, and between perception and action, even though there is enormously interesting work underway on both fronts. We will mine single-mindedly the vein that leads into phenomenology. To narrow the topic even further, I will confine the discussion to early vision. It is still an enormous field, and it provides more than enough materiel with which to examine some recent discussions of phenomenal consciousness.

The reader should be forewarned: the architecture of early vision is surprising, bizarre, *weird*. I will describe some of the surprising and weird features of that architecture, and then consider how folk concepts of appearance and awareness might be applied to it. The results, like the architecture itself, are somewhat bizarre. In particular, the preattentive architecture challenges the idea – it puts enormous stress on the common sense notion – that "phenomenal character" is and must be coeval with awareness. Instead, the two split apart, with "phenomenal character" showing up first, in places as yet unoccupied by awareness. After reviewing some of the evidence, I shall argue that the most reasonable conclusion is that some states of preattentive sensing are states in which one is being appeared-to, even though one is entirely unaware of what appears, of any aspect of its appearance, or of being in the state of being appeared-to. So "phenomenal character" and consciousness fly apart; their association in "phenomenal consciousness" is a merely contingent conjunction of two distinct things. This puts some stress on our ordinary notions; one burden of the argument is to consider the least costly methods for stress relief.

Handbook of the Philosophy of Science. Philosophy of Psychology and Cognitive Science
Volume editor: Paul Thagard
General editors: Dov M. Gabbay, Paul Thagard and John Woods
© 2007 Elsevier B.V. All rights reserved

1 THE ARCHITECTURE OF EARLY VISION

Models of early visual processing produced in the last forty years have grown increasingly sophisticated and complex, and in detail they differ markedly from one another. But certain global features of the landscape are found reflected in all of them (see [Neisser, 1967]). They are:

1. That distinct visual attributes of stimuli are initially processed largely independently of one another, in parallel, for the entire visual field;

2. That these processing streams, "channels", or "feature maps" extract information of varying orders of complexity, have complicated linkages with one another, may reside in distinct portions of the visual nervous system, and in cases of brain damage may be selectively dissociated from one another;

3. That at some point there is a competitive selection process, whereby some results of some of these parallel processes are selected for additional, capacity-limited, processing;

4. That a result of winning that selection is that discriminations in the selected domain can be made more quickly and more accurately; that the additional processing allows for the resolution of details and structure that might otherwise be unresolved, and that there may be features that are resoluble only as a result of such selective attention; and

5. That sensory representations which are not selected do not gain all these benefits, though they may gain some.

The initial stage is high-bandwidth analysis (e.g., of many megabits per second) over the entire visual field, extracting "local" or low level visual features, with different "channels" carrying on the analysis of different visual attributes. For example, motion, hue, luminance, and shape are likely candidates for being distinct attributes, analyzed in distinct channels. Each such channel carries on a high bandwidth analysis of its favorite feature for the entire visual field.

The surprise here is that the separation of channels is done by attribute (or by distinct dimensions of variation of attributes), and not something that would seem, a priori, or to a computer engineer, to be more reasonable. So for example one might think that if there are parallel channels, it would make sense to separate them by the regions they scrutinize, by their ambit. One covers the upper right quadrant, one the lower, etc; and each reports on everything visible occurring within their sector. Or perhaps the jobs of the different modules could be divvied up more intelligently and more dynamically: the visual field might be "segmented" into likely objects, and a channel is given the job of analysing everything visible about that object. The end result of such a module would be a visual representation of that object complete with color, shape, motion, texture, size, and so on. In fact a visual representation of the object would look suspiciously like a picture

of the object. We add up all these visual representations, and we get a complete picture of everything one can see.

But it is evident that early vision is not organized in that way at all. Finding a *picture* at the terminus should arouse suspicions; right around the corner from it you will fall under the disturbing, lidless gaze of the mind's eye. Visual perception does not produce anything like a fully detailed picture of everything one can see, and visual representation is nothing like a picture. Or at least if it is, it is a very odd picture (or better, set of pictures) indeed: we would need a picture of nothing but the object's color, leaving out all details of its shape; another picture of just shape, but no color; one of motion, but not of contour; and so on.

There are many different strands of evidence that support each of the five assertions, and many different points of disagreement between different models on each of the five. So for example the evidence for the first point – that distinct attributes are processed independently, and in parallel – comes from various sources. The first is neurophysiological: the careful tracing of visual pathways and different cortical areas connected by those pathways. The organization is quite intricate, with numerous independent input channels, roughly three dozen distinct cortical areas devoted to visual processing, and staggering interconnections between them. Different cortical areas seem specialized for the processing of distinct attributes, with distinct areas tentatively identified for motion, color, and form.

A second strand of evidence is neuropsychological: particular types of brain damage can "knock out" a particular dimension of variations in visual appearance, while leaving others intact. Particularly compelling is evidence of double dissociation: some types of brain damage knock out A, while leaving B; while others knock out B, and leave A. Achromatopsia and akinesia show this dissociation: one type of central nervous system damage can leave a patient colour blind, but still able to perceive motion. Another impairs discrimination of motion, while leaving color perception intact.

A third type of evidence is purely psychological, and is based on various types of experiments, including, most prominently, reaction times in visual search tasks [Treisman, 1998; 1996; 1993; 1988; Wolfe; 1994; 1996a; 1996b]. In a visual search task a "target" is to be located as quickly as possible within an array of varying numbers of "distractors". Particular combinations of attributes of targets and distractors can make an enormous difference to the speed of such searches, and these differential times give evidence that some classes of features are processed independently of others. For example it is found that in some classes of features, a target with a unique value within that class can be located in a fixed (and brief) interval, no matter how many distractors are present. If the target is the red letter, and we have a single red H among green H's, blue T's, and green X's, the red H will "pop out", and reaction times for its identification will be low and basically constant, no matter how many distractors are present. Similarly if we have a single letter "X" among H's and T's, of whatever colors, the X will pop out. But if the target can be identified only by a conjunction of features – a single red H among green H's, red and blue T's and red and green X's – then search becomes much

more arduous, and reaction times increase linearly with the number of distractors. It is as if each item in the array must be examined in turn, and such examination takes some finite time. The latter is sometimes called "serial search", or simply "less efficient" search.

The relatively straightforward inference from this result is that feature families which allow "pop out" are processed independently of one another, and such processing is high-bandwidth, covering the entire visual field. From the results above, for example, color and orientation (which allow the diagonals of an X to pop out from the horizontals and verticals of H's and T's) would seem to be features processed independently of one another. Each allows pop-out. Only if a target is defined by a conjunction does search become less efficient.

Pop out indicates a special kind of "preattentive" processing of the feature in question. Attention is effortlessly drawn to the singleton, no matter how many distractors are present. The processing to do this must happen before attention arrives on the scene, and it must encompass the entire visual field, since pop out occurs no matter how many distractors are present. So features in that group seem to be favored with independent, high-bandwidth, full field processing, completed preattentively.

There are other tests for the independent processing of distinct families of visible attributes [Treisman, 1988; Wolfe, 1996b]. One is "effortless texture segmentation", or the ability to discriminate regions defined by contours across which the feature in question changes. There is some boundary across which feature F changes; and the segments of that boundary in turn form a continuous closed curve. In this case the entire region might "pop out", or be immediately noticeable. The feature difference defining the texture can be minuscule: minute differences in orientation, closure (whether we have full squares or just three sides), and so on. Whereas other differences one might think would suffice do not: N's cannot be preattentively segmented from regions of backward N's, E's and F's are texture-wise indistinguishable, and so on (see [Julesz, 1984]). The inference is that certain feature values are the subject of preattentive parallel processing, while others are not. Orientation and closure yes; the difference between E and F, no.

With these and other sorts of experimental probes available, a considerable amount has been learned about the different kinds and categories of features processed preattentively. Jeremy Wolfe [1996b] has provided a useful list of visual features that pass both of the two tests just mentioned. He calls them "basic features":

- color
- orientation
- curvature
- vernier offset
- size/spatial frequency/scale

- motion
- pictorial depth cues
- stereoscopic depth
- gloss
- various form primitives

So for example a curved line (with basic feature number three) can be picked out quite efficiently if it is surrounded by straight lines, however many they may be (that's "pop out"); and it allows preattentive texture segmentation: a region filled with curved lines is effortlessly seen to have a different texture than one composed of straight ones. All of the features on the list pass both tests. Some are perhaps unfamiliar. We get "vernier offset" by breaking a straight line, and giving the second segment a slight offset, or sideways jog, relative to the first. (Human vision is extraordinarily sensitive to vernier offsets, detecting them at visual angles smaller than the angle between adjacent receptors in the retina.) "Pictorial depth cues" include cues for occlusion, perspective, shape from shading, and others that artists might use to make a two dimensional canvas represent a three dimensional scene. The last ("form primitives") is the most controversial and complicated. For reasons that will become important later (see [Wolfe and Bennett, 1997]) it is not exactly "shape", but includes local features of contours and surfaces that might together yield "cues" or evidence for shape. Wolfe mentions, for example: line termination, intersection, and closure (e.g. is the contour a closed curve?). Others are less well attested: topological features such as holes, convergence, and perhaps containment.

Different entries in the list of basic features are, intuitively, of differing orders of complexity. It takes more visual processing work to discover whether a contour is a closed curve or not than to discover whether some segment of it is oriented vertically or horizontally. Some of the features might be attributes of two dimensional polygons, while others (gloss, depth, occlusion) require more than two dimensions. Even "size" is ambiguous between a two dimensional interpretation (visual angle) and a three dimensional one closer to common sense. Preattentive features must then have some ordered information linkages to one another. In order to isolate a distinct texture region, for example, one must find boundary segments across which there is a discontinuity in the values of some feature; such segments must be linked to one another continuously; and the defined contour must be closed, without significant gaps. The relations that the data points must bear to one another become increasingly complex as one ascends in these constructions.

An example might clarify the problem. As emphasized by Marr [1982], the analysis of boundaries and edges in the optic array is a major task of early vision. At the receptoral level the raw data consists of variations in the membrane potentials of approximately two hundred million receptors – 200 mega-pixels, if you like.

Variations in these membrane potentials are caused by varying numbers of isomerizations of photopigment molecules, which are in turn related to the intensity and wavelength of the electromagnetic radiation absorbed by the cell. An initial task is simply finding the local edges – the regions where there are large discontinuities of intensity between neighboring cells. Solving this problem requires "local differencing" operations – a fancy way of saying we need to perform hundreds of millions of subtractions (between intensity values of neighboring receptors) and have usable results every few milliseconds. (And we are downsizing, so you have to do the job yourself, and complete it while paying attention to something else.) It's a bit of a challenge! But it seems to be accomplished within the retina itself. Then these local differences need to be linked up into coherent "edges" – connected tracks, or series of points, across which we find discontinuities of intensity. Only such edges have "orientations" (which needless to say are useful to record) and "contour" features such as curvature, kinks, or corners.

But even if we have isolated a continuous track across which there is a robust discontinuity of intensity, it is unclear what such an edge means in distal terms. It might be a "luminance" edge or a "reflectance" edge. That is, the differences in intensity might be caused by differences in the amount or character of light falling on a uniform surface, or they might be caused by differences in the way different parts of a surface reflect light of a uniform character. Perhaps the edge is the edge of a shadow; perhaps it is a place where surface characteristics change, or perhaps it is place where one object occludes another. Sorting out what kind of edge each edge might be is the next major job of early vision, and this analysis bumps us up into higher levels of complexity. A two dimensional world would lack shadows, and no object in it could occlude another. (We might have polygons that intersect, but a true flat-lander has no concept that one lies on top the other, or that one is behind the other.) So early on the analysis of edges will force us to assign depths to the neighboring points; the edge acquires an orientation that is not simply two dimensional.

We confront, then, a second source of complexity in the relations among preattentive features: what might be called their geometry, or the spatial structure of the coordinate system in which their relations are defined. This too varies across different visual features. It is clear for example that some "boundary" or "edge" features can be specified determinately in a two-dimensional retino-topic coordinate scheme. So the discontinuities of intensity which define the border might be specified to obtain between two adjacent retinal locations (and we use some spherical coordinate system for the two dimensions of the retinal surface). The retinotopic coordinates of the point on which one's eyes are focused are always $(0, 0)$. But of course the eyes move in the head, the head moves on the trunk, and the trunk moves in the world; each degree of freedom spawns another potential coordinate scheme, from head centered to body centered, to, eventually, a full "allocentric" system, in which no part of the body has any privileged position. "Being to the left of" has different interpretations in these different schemes: the left half of the retina might or might not be receiving stimuli from points to the

right of one's nose, which might or might not be to the right of one's sternum. And in an allocentric scheme, the relation no longer has a fixed interpretation at all, with as many variants as there are observers on the scene.

These logical and geometrical orderings of features are usefully combined in what has been called a "feature hierarchy" (see [Treisman, 1988]), starting at the simplest and proceeding upwards into those that require more intricate relations and more intricate coordinate schemes:

- Point level features: all the data needed for ascription of the feature in question can be had given just one visible point.
- Merely local differences, across some span, in some feature dimension: "discontinuities of intensity". Requires relations between at least two points: but "edge" is still too sophisticated a term for this.
- Segment or edge level features: oriented and connected discontinuities of intensity, connected in a series, and so having some orientation. Here for the first time a connected series of local differences is treated *as* a a series or group, and has an orientation different from the span across which there are feature differences.
- 2d surface properties, with no depth: properties of a 2d surface, specifically, so these are properties ascribed to points on a plane, in varying ways; no shadows or depth yet; no points on the surface vary in distance from the observer.
- 2.5d surface properties, so called because each point on the surface is also assigned a unique depth coordinate. (It is not fully three dimensional, because for each 2d point there is only one depth coordinate, and not a full third axis. Such surfaces can be tilted relative to the observer, have bumps, convexities, and concavities, but no three dimensional volume.)
- Full 3d surfaces and volumes.

Different "basic features" make their first appearances at different points within this hierarchy. It is perhaps problematic whether there are any point level features, but color is often mentioned as a possible example. (The problem is that color perception seems always to require color contrast, and so requires relations between a point and its surround.) But color also shows up as a property of two dimensional surfaces (the famous color patch of sense-data theorists), a property of oriented surfaces (walls of the same color meeting at a corner), a property of three dimensional surfaces of objects, and a property of three dimensional volumes (as in colored mists or colored light). We must arrive at "edge level" features before any interpretation can be provided for orientation, curvature, vernier offset, and various of the form primitives; though features that might show up as mere features of edges (such as curvature, orientation, and closure) are instantiated at

higher points of the feature hierarchy as well. Size and motion require at least a two dimensional coordinate scheme. The depth suggested by pictorial depth cues requires more than two dimensions, though some may suggest only orientation or curvature of a two dimensional surface, so fit in a 2.5 dimensional space. But shadows and occlusion, stereoscopic depth, and gloss have their expected interpretation only if they are placed in a coordinate scheme that is three dimensional. The difference between a reflectance and an illumination edge can be introduced only at this point; prior to it one thing cannot occlude another. Note that "motion" can in fact implicate developments across *four* dimensions: a boxer described as "weaving" and "bobbing" has a spatio-temporal trajectory, with a prominent temporal axis (as well as the spatial ones) along which developments have a characteristic form. Anyone acquainted with the idea that rhythms are temporal "shapes" will understand the point (see [Goodman, 1977]). Features such as "glistening" or "shimmering" also require some temporal extension.

In short, features are not all of piece; some have a higher order of logical and geometrical complexity than others. The point may seem purely pedantic, but grasping it is essential in the clinic, since otherwise the existence and symptomatology of the various *visual agnosias* are entirely inexplicable. Several of these syndromes seem precisely to be loss of the capacity to detect higher-order features, even though lower-order ones are still intact. An agnosic patient may show no loss of visual acuity, as shown by perimetry tests; yet lose the ability to visually recognize objects by their shape. In associative visual object agnosia, a patient may have unimpaired perception of local segment features – as shown by the ability to draw objects or match pairs of pictures faithfully – yet fail to recognize the object by sight. (See [Farah, 1990, 57–60]). Instead the patient often volunteers a list of lower order form primitives, and then hazards a guess as what the object might be:

> The patient could not identify common objects present visually and did not know what was on his plate until he tasted it. He identified objects immediately on touching them. When shown a stethoscope, he described it as 'a long cord with a round thing at the end' and asked if it could be a watch. He identified a can opener as 'could be a key'. Asked to name a cigarette lighter, he said 'I don't know' but named it after the examiner lit it. He said he was 'not sure' when shown a toothbrush. Asked to identify a comb, he said 'I don't know'. When shown a large matchbook, he said 'It could be container for keys'. [Farah, 1990, 58]

What seems to be lost is the ability to recognize the overall shape of an object, with this capacity replaced by arduous feature-by-feature matching and hypothesis testing. Such patients can fail to recognize common objects (tea bags, rings, pens) which they can quite accurately draw.

An even more spectacular and specific agnosia is prosopagnosia, or the inability visually to recognize faces. The deficit can be quite specific to visual identification

of faces. Such a patient may have no sensory loss and can recognize other common objects visually, and furthermore can recognize people by other distinguishing characteristics (voice, clothing, gait, etc); but if the face is the only clue, even a spouse can be misidentified. Facial features involve shapes, orientations, sizes, and symmetries, and are clearly higher order. Humans seem to have a "fusiform face area" whose activity is highly correlated with face recognition, which is not too surprising given the biological importance of recognizing con-specifics. In terms of our feature hierarchy this is specialized feature analysis at a rather high level.

2 COMPETITION, SELECTION, ATTENTION

"Basic features" are preattentive: their detection is completed prior to the activation of selective attention. They are noticeable or salient "attention-grabbers"; they draw attention to themselves, and so must not themselves require attention in order to be detected. Indeed, such detection is often the basis on which the selection of selective attention is made: attention might be drawn to the letter "*h*" because it is the only red letter in the display. Furthermore, different sets of basic features must be detected simultaneously and in parallel, since at any given moment any one of them could grab attention in the same effortless, almost instantaneous fashion. A single red letter, or a moving one, or one that displays a unique orientation, or a single glossy one, ... and so on, could grab attention in (roughly) fixed time no matter how many distractors. So the inference seems robust: one must simultaneously and in parallel detect basic features of color, motion, orientation, glossiness, and so on.

For various reasons offered over the years, it is clear that mental life cannot forever proceed in such parallel, independent tracks. Sadly, at some point it becomes subject to central control, and acquires at the least the appearance of unity. That is, at some point there is a transition from high-bandwidth, data-driven, independent channels to something more centralized, of lesser bandwidth, and with a greater capacity to pre-empt the underlings. The contrast between processing "before" and "after" this transition point has been made in various ways. So as not to beg any questions, I shall simply dub the two sorts of processing "preattentive" (or "early") and "post-attentive" ("central"). There have been many disagreements over the years on how to characterize the differences between these two sorts of processing: parallel v. serial; bottom-up v. top-down; data-driven v. task-driven; exogenous v. endogenous; modular v. general-purpose. Likewise there have been many controversies over where the transition point is found: whether it is "early" or "late"; and why there is a transition point at all. Some have argued that the transition is needed because of the enormous amounts of data, and the lower ("serial") capacity of central processing; others that it may simply reflect the need for response selection: to choose *the* way to move now, since one cannot manage to send oneself to two distinct destinations. Psychologists continue to differ over the benefits of central processing, and over the fate of sensory representations that do not receive it. But, however else it is conceived, there is a consensus that the

point of transition is the point where selective attention selects some representations, and not others, for further processing. This point controls, or exerts some control over, the flow of information from one style of processing to another. Various homely analogies are used: it is a gateway, a filter, a movable window, through which some can pass, but most are barred.

Post-attentive visual perception differs in various ways from the preattentive variety. It is "capacity-limited" in the sense that it cannot simultaneously be applied to every portion of the ambient visual array. A subset or subregion must be selected, and even then the analysis of that portion takes some time. So if central processing must be used to search every visible item, search times increase linearly with the number of items. This is often called "serial" search, and its contrast with the fast, "parallel" search in pop-out is one of the more robust findings of experimental psychology. The serial appearance of post attentive processing may be deceptive (parallel processes with limited capacities could yield the same pattern of results), and later models have moved away from the notion that there is a dichotomy between "serial" and "parallel" processing, but the basic empirical contrast (now called "efficient" vs "inefficient" search) is unarguable.

If central processing is capacity-limited in this way, then it would be swamped if it had to deal with every visible feature simultaneously. Some selection must be made; the system must somehow establish priorities among the important or salient visibilia, and select which ones are first to receive the benefits of fuller analysis. Since resources for that fuller analysis are limited, they must be allocated. For the psychologist, this is the key function of selective attention: to select from a pool of potential candidates some favored few that are elected to receive further processing. This selective process is widely agreed to be competitive; the candidates are racing against one another. The winners of the competition are the ones that are selected. Some models even mention primary elections, followed by a general election (see [Desimone and Duncan, 1995]). The rules of the race – the terms of the competition – can vary, depending on the current needs and desires of the needy organism. The smell of food is more compelling if one is hungry, and (surprise!) so is its taste. Bird-watchers and flower-gardeners will notice very different things if they walk together through the woods. In these and other competitions the competitors are representations, produced by the various feature channels, of stimuli one has sensed. Winners are the ones selected, by selective attention, for further, central processing. There may not always be exactly one winner – perhaps the hungry organism wants to gobble down several things simultaneously – but the losers in the competition do not gain entry to, are not selected for, central processing.

Common sense ascribes various benefits to the practice of paying attention to what you are doing, and experimental psychology has confirmed many of them. For example, if you are paying attention to a particular visual region, you are more likely to notice novel stimuli there, you will respond more quickly, and your discriminations of features in that region will be more accurate. (See [Posner, 1994; Posner and Rothbart, 1992]). "Paying attention to the region" is thought

to be equivalent to devoting more processing power to representations of features and events in that region; and that additional processing allows quicker detection, more accurate discrimination, and better resolution of details and structure. Basic features can perhaps be resolved without this additional processing, as can gross morphology and some jumble of form primitives, but discerning the precise shape of the thing – the exact structure of relations among those form elements – seems to require the sort of central processing that can only be brought to bear if the target is selected by selective attention (see [Wolfe and Bennett, 1997]).

This idea can explain why paying attention to something visible seems to place that thing in a spotlight – a rather mysterious spotlight, it turns out, since it employs no physical light. Increasing the illumination on an object can indeed help one see it better, but so can increasing the processing power devoted to analysis of visual detail. As anyone who has used the world wide web knows, the more computer power one can throw at a picture, the quicker its details are resolved. The spotlight analogy is probably as old as the human capacity to direct artificial illumination on objects of interest, but perhaps internet usage will eventually provide it with a rival. Attending to something helps one see it better, not by directing occult illumination at the object, but rather by allocating additional processing to the analysis of visual features.

Competition for selection imposes an organization on the transition between preattentive and post-attentive perception. In any race there are winners, runners-up, and all those who also ran, surrounded by an even larger class of spectators, hangers-on, couch potatoes, and others who did not compete at all. In the spotlight metaphor we have an object or two in the focus; others in the penumbra, to which the spotlight could be slewed with minimal effort; and a larger class of things entirely out of its range. Similarly, the stream of consciousness is reputed to have a few phenomena in the foreground at any moment, and a somewhat larger but still finite collection of phenomena in the background. The latter are all the ones to which one is not directly attending, but of which one is "aware". Finally, surrounding that is a limitless but insensible netherworld of phenomena to which one cannot shift attention. As many authors have noted, it is clear that one is "aware" of things beyond those to which one is strictly paying attention; indeed the former often provide the next targets when attention shifts. A simple hypothesis is that this terminology reflects the fact that selective attention is allocated competitively. One focuses attention on the winners (they are in the focus, the foreground, the spotlight) but the runners-up are the ones of which one is "aware". The latter could have been selected for further processing, and would have been, but for the fact that something else came in first. In the next moment the currents might shift, and something of which one is currently merely "aware" might attract one's attention. If the currents are shifting rapidly in this fashion, the waters are choppy, and we call it "being distracted".

The more important point for our purposes is that if there is a selective process, the selection itself must be based in part on current information. Efficient allocation of scant central resources requires that it be at least somewhat sensitive

to current developments. An animal whose attention was directed a priori would fail to manage any contingency; its steerage would be competent only among the necessary truths. So some sensory information is used in order to select where attention will be directed next. This information is found behind the stage-works, so to speak: it is not in the spotlight, or anywhere out there on the stage, but is rather being used, by the evil little homunculus who directs the spotlight, to determine where the spotlight will go a moment hence. (A homunculus is by definition a little man, and this one is evil because he controls what *you* pay attention to, and you never gave him permission.) It is not information of which you are aware, but is rather information already received by that sub-personal agency who is right now determining what you will be paying attention to a moment from now.

Furthermore, if there is some regular process of selection, it is carried out by some sort of mechanism, which like any mechanism has normal operating parameters, limits, and vulnerabilities. Posner [1994; Posner and Petersen, 1990; Posner and Rothbart, 1992] has done the most to reveal some of the details of how selection itself works. There are three separable processes in shifting visual attention: disengaging it from the current focus, moving it, and then engaging it on a new focus. Each of these three subprocesses has its own normal operating parameters and areas of vulnerability. For example, it takes a certain amount of time to disengage and to shift attention, and subjects have tremendous difficulty perceiving events that happen in that interval. Change blindness and the attentional "blink" both exploit this parameter of attentional selection. Change blindness is not literally "blindness": there is no sensory loss, and the change is perfectly perceptible once it is noticed. What is blind to the change are the mechanisms of attention: the change occurs in such a way that it is difficult or impossible for it to attract attention. And if it *cannot* attract one's attention, one is "unaware" of it.

The attentional blink is a similar phenomenon which was discovered years earlier (see [Raymond *et al.*, 1992; 1995; Luck *et al.*, 1996]. Subjects are presented a very rapid sequence of visual stimuli, and asked to respond to particular targets. For example, they press a key each time a vowel appears in a stream of letters. If two targets are presented within a certain interval of one another, the odds of missing the second target increase. But the odds do not increase in the way one might expect: if the second target is immediately after the first one, it is more likely to be identified correctly than if there is a slight lag, of about a hundred milliseconds, before it is presented. The worst performance is found when several letters intervene between first and second targets; with lesser intervals, or greater ones, identification improves. The suggestion is that the hundred millisecond lag indicates an operating parameter of the mechanisms that disengage, shift, and re-engage attention: by that time attention is engaged with the task of responding to the first target, and it is difficult to disengage it to register the second target. Just as changes cannot be noticed during an eye blink, the target is not noticed during an "attentional blink".

Both change blindness and the attentional blink show that it is possible to exploit vulnerabilities in the mechanisms that shift attention, and find stimuli

that (given everything else going on) slip between the cracks of the selective process. These are stimuli to which one *cannot* shift attention. It is interesting that change blindness is called "blindness"; ordinary language makes scant distinction between events one does not sense and events one cannot notice, often treating them as equivalent. They are both events to which one is, in different senses of the word, "insensible". But change blindness and the attentional blink are not sensory failures at all; they are demonstrations of the vulnerabilities of an attentional mechanism, specifically the one which does the job of selection. If that mechanism has a hitch, then one cannot shift attention to a stimulus; one is insensible to it, and it is not overly misleading to say one is "blind" to it. It is a distinct kind of blindness from the sensory kind, and from the agnosia kind ("mind blindness"), but it is a third way in which one can fail to have normal visual experience.

And this points us to the most deeply surprising feature of the visual architecture described so far. Visual inputs are divvied up into discrete dimensions of independently detectable "features", and processed independently and in parallel in topographically organized "feature maps". Each such feature map covers the entire visual field, and so in a certain way it is "panoramic", but only in terms of its own, favorite feature. One map is a specialist for motion, anywhere; another for color, anywhere; and so on. After all these feature extractions one might expect to find some comprehensive representation of all of them, put together; a synoptic view. But there seems to be no such thing. Beyond the feature maps we find nothing but what seem to be scattered (but carefully directed) glimpses; episodic and focused representations, serving the needs at hand. Their objects are scattered in space and time, in a surprisingly thin distribution. Change blindness, the attentional blink, and other results suggest that after the feature maps perhaps all we have are these temporary representations of conjunctions, relations, and structured properties, marking the visual cruxes for current affairs of interest, but vanishing as soon as the interests change.

The system can afford to analyse conjunctions, relations, and more structured properties only if such analysis is necessary for some on-going action. Even then it has to pick and choose. So conjunctions and more structured features are extracted only for isolated and scattered portions of space-time. Post-attentive processing is not engaged in the construction of a panoramic view, a fully detailed picture of all that is visible, but is better thought of as expedient probing and sampling, with targets carefully selected. For example, as you sit down you might employ well directed glimpses, and rather evanescent representations, to find the edge of the table, the front edge of the chair, the center of the seat, the handle of the coffee mug. Each of these targets requires rather sophisticated analysis (edges of objects, relational properties, etc), but results can be discarded as soon they have been used.

It is only in early vision – in preattentive, sensory processes – that we find anything resembling panoramic views. Even there the resemblance is not great. Perceiving, or post-attentive processing, is a narrowly focused business of probing and sampling, constructing representations that are punctate and transient, serv-

ing the needs at hand. It yields at best a pile of glimpses. In this architecture, a comprehensive picture is nowhere to be found.

3 FEATURES AND PHENOMENAL PROPERTIES

At this point in the story the sentient organism is focusing attention on some salient visible feature, and is aware of others, which were also in the running in the competition for selection. Such awareness of a visible feature counts as an example of what is ordinarily called "visual experience". The animal is aware of something it sees. One would think therefore that the state qualifies as an example of "phenomenal consciousness". If you want to check its qualifications, here is the initial baptism for the latter, technical term:

> *P*-consciousness is experience. *P*-conscious properties are experiential properties. *P*-conscious states are experiential states, that is, a state is *P*-conscious if it has experiential properties. The totality of the experiential properties of a state are "what it is like" to have it. Moving from synonyms to examples, we have *P*-conscious states when we see, hear, smell, taste, and have pains. *P*-conscious properties include the experiential properties of sensations, feelings, and perceptions, but I would also include thoughts, wants, and emotions. [Block, 1997, 380]

A state in which one is aware of something one sees is, presumably, a state of visual experience. Any state in which one sees something without being aware of what one sees would not be a state in which one could be said to experience that thing, and Block quite consistently therefore excludes "subliminal perception" from the ranks of phenomenal consciousness, on the grounds that it is "perception without awareness" (see [Block, 1997, 393]). And if the architecture so far described has the wherewithal to yield states of visual experience, then it is thereby also yielding states of "phenomenal consciousness".

Which states in the architecture are the states of phenomenal consciousness? The question gives one pause. The issue becomes quite complex, because the notion of a "phenomenal" property is quite complex. In fact the word is multiply ambiguous, and those ambiguities contribute to the problem. It has simpler and less simple interpretations, often conflated with one another.

The simplest and perhaps oldest notion of "phenomenal property" derives from ancient thought about the problem of illusion. Sometimes things are not as they appear. The arrow looks bent when it is half in water; the mountains look blue from a distance. Although the study of illusion is not a large part of perceptual psychology, the terminology needed to describe illusions looms large in the philosophical traditions. And indeed there are well known and replicatable stimulus arrangements under which people are subject to visual illusions of various sorts. So one line in the Mueller-Lyer illusion looks longer than the other (even though are both of the same length); after adapting to a bright red light, white paper will

look green. In these situations we have some visual appearance which is illusory: it fails to correspond with how things really are. The ancestral home for phenomenal properties is found in these contexts. A phenomenal property is a property of appearance – of how things look, feel, sound, or in general, seem or appear. Such a property may or may not be veridical.

A slightly more recent analysis [Chisholm, 1957] first picks out a class of "verbs of appearance", with which various predicates for sensible properties can be attributed, even though the referent of the sentence fails to have that property. So "the paper looks red" can be true even though the paper is, in fact, white. "Looks" is then a verb of appearance, and "red" (or any other predicate for a sensible property), used within the context of any verb of appearance, attributes a phenomenal property. "Looks red" characterizes how the paper appears, and it can look that way even though it is in fact white.

The acid test for this older notion of phenomenal property is that such a property can characterize appearances even if those appearances diverge from reality, as they do, prototypically, in sensory illusions. It remains true that one line in the Mueller Lyer illusion continues to look longer than the other, even though measurement shows them to be the same length, and one knows them to be the same length. But the notion gets generalized to include properties of appearance even when they are veridical – even in those situations in which things *are* as they appear. So "looks red" comes to characterize the similarity in appearance across a large class of episodes: all those episodes in which something looks red but is not, and those in which something looks red and is red (see [Sellars, 1963]). All those episodes have a shared feature – a shared phenomenal property – characterizing how something appeared. In some of them that appearance is veridical, in others not.

This is the oldest and simplest notion of "phenomenal property", but there are many other more recent and more intricate variants. Sense-data theories complicated the issue by identifying phenomenal properties with properties of odd objects called sense data; and the "qualia" debates around functionalism further confused matters by abandoning the odd objects but retaining talk of their inscrutable properties. Many philosophers today think of qualia as properties of mental states, and some identify those properties with phenomenal properties.

One question is whether phenomenal properties are characteristics of sensible appearance, or characteristics of something else. The other alternatives these days are, in order of increasing specificity: they are properties of mental states; they are properties of conscious mental states; they are characteristics of being conscious of one's own mental states. The simplest notion, as already described, is that "looks red" is a paradigm phenomenal property. A mental state itself has no known visible appearance; no mental state literally "looks red", as far we know. There is presumably some property of a mental state (specifically a visual sensory state) in virtue of which something one sees looks red, but that property is not itself the property of looking red. Similarly, if one is conscious of something that looks red, there may be some property of the state of consciousness related to the fact

that the thing one sees look red, but the property of the state of consciousness is certainly not one of looking red. If it too is a phenomenal property, it is a different sort than the one that can be instantiated by things one sees.

The most intricate notion of "phenomenal" properties ties them directly to "what it is like" to have certain mental states. Block for example identifies "experiential" properties with the totality of "what it is like" to have a state. Notice that the bearer of the properties has shifted yet again. In order for there to be something it is like to have some mental state, one must be conscious of that state. Nagel [1974; 1979] thought this to be an analytic truth, and that he was characterizing what it is for a state to *be* a conscious mental state. There is something it is like to have a state if and only if that state is a state of consciousness. He used the phrase to point to the problem of consciousness, which he thought all the accounts standard at the time ignored. If we take Nagel at his word, these properties obtain in virtue of similarities and differences in what it is like to have various mental states. They are variations in how states of which one is conscious appear to the one conscious of them; variations in the manner of appearance of mental states to one who is conscious of those states.

Suppose there is some property red* in virtue of which when one has a mental state with red*, something out there looks red. Red* will still fail to be a phenomenal property of this, the most intricate level. Instead we need some property red++, which characterizes what it is like to have a red* mental state. These are characteristics of one's consciousness of one's sensory states, so they neither look red nor suffice to make something in the surroundings look red. Instead they are closer to an individual mode of appearance of a mental state with the red* property.

In short, phenomenal properties were kicked indoors (they became properties of mental states), and then kicked upstairs (they came to characterize what it is like to be aware of one's mental states). It is certainly true that on the latter view, a state that has a phenomenal property must be a state of consciousness – in fact, it must be a property of a state in which one is conscious of one of one's mental states. But by the same token, this sort of phenomenal property is entirely distinct from what we are saying of the white paper when we say it looks red. Something looks red. One has a state of sensing something which looks red. One is (sometimes) conscious of being in a state of sensing something which looks red. And perhaps there is something distinctive and special about what it is like when one is conscious of being in a state of sensing something which looks red. But whatever that last special and distinctive property is, it certainly cannot be "looks red"! It is a property of a different bearer, in a different context, to explain different facts.

For now I propose we focus on the simplest variety: being aware of something that looks red. It gives us problems enough. The phenomenal property is the traditional sort: a characteristic of how things appear. The prototypical test is finding some stimulus arrangement in which things appear to have properties they do not in fact have. There is no need to get fancy about the assertion that

sometimes things are not as they appear; "is", "is not", and "seems to be" need only their ordinary interpretations. Our target is, for example, a state in which one is aware of something which seems to be red but is not red. This counts as a state of "phenomenal consciousness" because it involves awareness. The latter involvement is also as simple as can be: one senses something, and is thereby made aware of something one senses. One need not be aware *that* the thing in question looks red; we might have just "thing" awareness, not awareness "that". Furthermore, there is no requirement that one also be conscious of seeing something or of seeming to see something. One might be aware merely of *what* one sees, and not also of the seeing of it. Perhaps the visual experience is so compelling that one is, as we say, lost in it: perhaps one is absorbed by the colours of the sunset, and is not at the same time conscious of seeing the sunset.

So the notion I want to focus is the very simplest of the ones that might qualify as "phenomenal consciousness". Conflating these different notions has a high cost: the different versions have different bearers and their truth conditions differ. There is no logical or conceptual connection between the characteristic of a sensory episode in which something looks red and the distinct characteristic of a mental state in which there is something it is like for me to be conscious of sensing something that looks red. The cost is that we might confuse failure to explain one of these with failure to explain all of them. In fact I will argue that the simplest variety – being aware of something that looks red – is already within our grasp. We cannot explain all the different intricacies of different varieties of phenomenal consciousness, but if the simplest variety – being aware of something that looks red – still counts, then perhaps we can explain *a* variety of phenomenal consciousness.

4 PREATTENTIVE PHENOMENAL PROPERTIES

The first step to showing this is to show that the traditional notion of phenomenal properties has instances within the preattentive domain. This is something of a surprise, particularly to those who think that phenomenal character requires consciousness. Some preattentively detected features can be shown to be phenomenal properties, in the traditional sense. There are preattentive illusions. To put it another way, there are sensory states that have phenomenal character in the traditional sense (character which, were one conscious of it, would yield perceptual reports such as "it seems to be moving" or "it looks like a triangle" or "it appears to have a shadow"), even though, (a) in these cases, nothing is moving, there is no triangle, and nothing is in shadow, and (b) the sensory states are preattentive, and hence unconscious. (The "hence" takes some showing.) In fact these locutions are the natural ones to apply to the three different examples of preattentive phenomenal properties that will be described: apparent motion, subjective contours (as in Kanizsa triangles), and pictorial depth cues. All three have been shown to allow efficient search (or "pop out") and can serve to attract attention, in an effortless fashion, to a target. So they seem to be detected prior to selection by

selective attention. Indeed, their detection serves to direct the selective processes of selective attention.

Apparent motion is the simplest to describe. Flashing dots on a screen can give the appearance of motion, depending on the exact parameters of angular separation, brightness, and timing. If the interval is too long, subjects will see one flash, then a second, distinct flash. If it is too short, subjects will see two distinct dots flash nearly simultaneously. But if it is just right, subjects will seem to see one dot that moves across the intervening space, landing at the point of what is in fact the second flash. Nothing in fact moves on the screen, but the display generates a compelling illusion of motion.

Description of apparent motion is a paradigm example of use of the traditional notion of phenomenal properties. "Apparent motion" is not real motion; it simply the contrary-to-fact appearance of motion. The dot that moves through the intervening space is a merely phenomenal dot; there is no luminescence moving across those portions of the screen.

So it is very interesting indeed to discover that apparent motion sustains "pop out" in multi-item displays (see [Wolfe, 1996b, 32–33]). For example, Ivry and Cohen [1990] used displays in which "*X*"s could appear to move on the screen horizontally or vertically. Apparent motion was induced by alternating frames on the computer screen, with the lit-up pixels spaced and timed appropriately. They showed that a single "*X*" which appeared to move horizontally could be picked out quite efficiently from any number of distractors appearing to move vertically, as long as the apparent motion was over fairly short distances. Long-range apparent motion resulted in less efficient search. In another experiment the targets and distractors were "forms" constructed by filling in eight of sixteen pixels in a four by four grid. In the short-range condition these "forms" could change from frame to frame, and the pixel colors could also change, yet a target appearing to move in a unique direction among apparently moving distractors could still be detected efficiently. Detecting the target in displays using longer range apparent motion was less efficient, requiring an item-by-item search. Ivry and Cohen suggest that detection of short range motion is a process of early vision, performed in parallel across the entire visual field, while long range motion "requires an attentional process in which each object in the display is examined in a serial fashion" [1990, 329]. Both are separable from the discrimination of form or color.

But apparent motion is merely phenomenal. Hence a merely phenomenal property – *appearing* to move horizontally – sustains pop-out and effortless texture segmentation. It sounds portentous, but that's because of all the extra freight that the term "phenomenal" has acquired over the years. It used to be a clean little skiff but is now encrusted with barnacles and algae.

Subjective contours provide an even more compelling example. A "Kanizsa" square is not in fact a square, but instead a stimulus array that is cleverly designed so that it creates the impression of a square; it *looks* like a square. (See figure 1.). At the vertices one places circles with cut-outs corresponding to the local details of each vertex. The circles with cut-outs are often called "pac men" because they

look like the hero in the "Pacman" arcade game. The figure creates a compelling illusion of continuous edges, running from vertex to vertex; and the illusory shape will also appear to be brighter than the background: a slightly brighter square, floating on top of the white paper. But if you examine it carefully you see that there are no line segments between the circles, and the white paper is the same shade throughout.

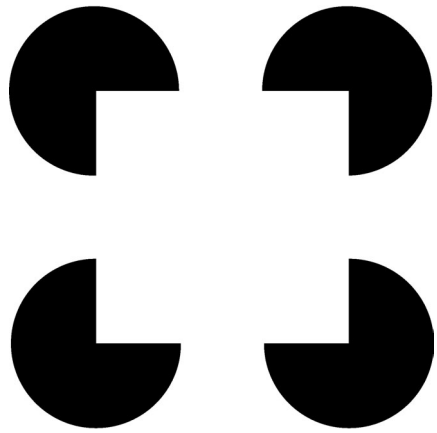

Figure 1. A Kanizsa square. A compelling illusion of something that appears to be a square, with continuous edges, lying on top of the four circles, and slightly brighter than the background.

The edges (which appear to be there, but are not) are called "subjective contours". They reveal again the importance of edge-detection in early vision; here the mechanisms seem a bit over-enthusiastic. It does not take much to seem to see an edge; just a few cues will do. As Ramachandran [1987] and others have pointed out, in natural environments of dappled light and heterogeneous reflectance profiles, camouflage breaking would be considerably easier if mechanisms for edge detection were quick on the draw, picking out boundaries on the basis of sparse data. Kanizsa figures show just how sparse the data can be.

Our terminology for "subjective contours" is another example of good, clean, traditional talk of phenomenal properties. A subjective contour is not real; it is merely the appearance of a contour. There is no real edge at that place; the appearance is an illusion. So it is a surprise to find that subjective contours can be detected in parallel, preattentive stages of vision. The simpler sort of experiment shows "pop out" for merely subjective polygons. For example, Davis and Driver [1994] used a stimulus array in which some portions had the "pac men" circles aligned so as to create the impression of a Kanizsa square, while in other portions the circles were rotated ninety degrees. In the latter condition the cut-outs are not correctly aligned, no potential edge is given two vertices, and no impression of a subjective square arises. They found efficient search and effortless segmentation

of the subjective square, regardless of the number of non-subjective (e.g., merely actual!) distractors. Enclosing the "pac man" circles with a boundary also disrupts the formation of subjective contours, and Davis and Driver showed that the same stimulus, but with pac-men encircled, also gave inefficient search. There was still one place where the alignment was correct, but without the subjective contours it had to be found by laborious, item-by-item scanning.

Davis and Driver [1998] provided further controls to show that subjective polygons can indeed by detected in parallel. That is, arrays that presented multiple subjective figures simultaneously, and in which the search target was defined as a feature in just one of those subjective figures, still allowed efficient search for the target. The target was a brown circle with cut-out, said cut-out in some trials aligned as a corner of a Kanizsa square. Non-targets were whole circles, which might or might not occupy the position of a corner for a potential square. The really clever bit was to present the notched circles by stereoscopic fusion, adjusted so that they would appear to lie either above or below the plane on which the subjective squares seemed to be located. If above the plane, then the subject would seem to see a notched circle, and through the notch of it the fourth corner of a subjective square would be visible. If below the plane, then the subject would seem to see a circle *occluded* by the corner of the subjective square. In the latter case it would not be clear to the subject that the circle was a notched circle: it would tend to appear as a full circle, with a portion occluded. It would appear to be occluded by a figure that is itself merely "subjective".

Results confirmed that, as Davis and Driver put it, "the presence of the notched large circle target is not immediately apparent, because of amodal completion behind the abutting subjective square" [Davis and Driver, 1998, 176]. That is, with the notched circles apparently lying in plane behind the subjective squares, search for a notched one is difficult and inefficient. Each circle near a subjective square must be examined in turn, since it might be a full circle, partially occluded by the subjective square that seems to lie on top of it. Whereas when presented to appear above the plane of the subjective squares, search for one notched circle is quite efficient. There were other important controls, but the gist is that multiple subjective figures are indeed coded in an obligatory, parallel fashion, as otherwise there would be no apparent occlusion, and no inefficient search.

The discussion of results invokes not only squares that are not there, but also three different apparent depths: above the screen, on the screen (where the unreal squares appear) or below the screen. As mentioned, this impression was produced stereoscopically: the "brown" circles were not really brown, but instead were the product of stereoscopic fusion of red and green circles, seen through red-green 3-D spectacles. This abuts, but does not encroach upon, my third category of examples: experiments showing preattentive parallel extraction of *pictorial* depth cues. A pictorial depth cue is a cue that can be presented on a flat piece of paper, without any special spectacles required, but which suffices to give an appearance of differing depths. Perspective and occlusion cues are well-known. Early work on machine vision at MIT explored how the the exact topology of intersections (Y

joints v. T joints) could be informative, and it was a target of analysis for early vision in Marr's models as well. (A T vs. a Y is a nice local feature that automatic mechanisms of early vision pick up; see [Enns and Rensink, 1991].) Shape from shading and surface contours are other, potent, indicators. Many of these yield pop out: occlusion, slant from texture, junction topology (T v Y junctions), shadows, and shape from shading, for example (see [Wolfe, 1996b, 39]).

In the two dimensional realm there are no shadows, but merely darker or brighter patches of the picture. Nevertheless, careful arrangement of the lighter and darker patches can give a compelling impression of depth. Not all arrangements will do; some might have inconsistent cues for the direction of the light source, and others are inconsistent with an opaque three dimensional object which occludes some of that light. The surface which the shadow falls upon must be a possible three dimensional surface, and the way the shadow falls upon it yields "shape from shading" cues about *its* shape. There are further subtleties in variations of brightness, reflections, glare, light scatter, and diffusion, but these can be ignored for now. A picture that fails to satisfy enough of these constraints will fail to "look like" a normal object.

Given a picture that does succeed in depicting a three dimensional object, lit up from one direction, casting a shadow, it is possible to rearrange its light and dark patches to yield a stimulus that has exactly the same surface areas at the same reflectance levels, but which fails to "look like" a three dimensional object. This was one of the manipulations that Enns and Rensink [1990] used to test preattentive sensitivity to pictorial depth cues. It was found that in the 3d panel – that is, in the panel presenting what *appear to be* three dimensional cubes – the one cube that appears to be lit from the side "pops out" and can be identified efficiently, independently of the number of distractors. But the corresponding target in the 2d panel is difficult to find, and requires each element to be scrutinized in turn. Somehow the visual system can effortlessly pick out the image with valid shape-from-shading cues, but must resort to serial search if the appearance of depth is eliminated. So shape from shading, despite its geometrical complexity, is detected in automatic, parallel, preattentive processes.

But just as a subjective contour is not a real contour, so the shape one seems to see when one looks at these pictures is not real. The object depicted appears to be three dimensional, while the picture is flat. As in our other examples, "pictorial depth cues" are cues to how things appear, not necessarily to how they are. And if the cue is *pictorial*, then it is derived from arranging patches on a two dimensional plane so as to look like a three dimensional object. In short, "looks like a cube" ascribes a phenomenal property, a characteristic of appearance. That figures so characterized can pop out preattentively hence yields another example of a preattentive phenomenal property.

5 PHENOMENAL CONSCIOUSNESS, BUT HOLD THE CONSCIOUSNESS

It bears emphasis that all the features so far discussed are preattentive: representations of them are completed prior to, and without, the assistance of any perceptual or cognitive capacities that might be summoned by the activation of selective attention. For if basic features are the ones allowing "pop out" and effortless texture segmentation, then these features are the ones that help determine which parts of the optic array *attract* attention. They are features to which attention is drawn. It follows that their extraction must be complete – the sensing of them must have already happened – *before* attention is drawn to the vicinity.

Motion, for example, is a terrific attention-getter; and results in the last section showed that merely apparent motion can be as well. The one area of the screen where things appear to be moving horizontally grabs attention, and can be identified in finite time, more or less independently of the number of other items that appear to be moving vertically. The task does not require deliberation or effort: the horizontal region is noticeable and salient, and attention is drawn to it, rather than being required for it. It is salient because it is the one region where apparent motion has a horizontal direction, rather than vertical.

It follows that the detection of the direction of apparent motion and of its singularity in one area must be completed prior to the moment that attention gets slewed to the area in question. The visual system can detect where things *appear* to be moving horizontally without one needing to focus attention on the candidates. As noted in section 3, the selection process itself requires some sort of data and some priorities in terms of which the selection can be made; and what these phenomena are taken to show is that basic features help provide the data in terms of which the selections of selective attention are made. To revert to the analogy of that section: this is information that is passed along offstage, to the little guy behind the curtains who controls where the spotlight is to go next. It is information he gets *before* the spotlight is moved, and here (given the imminent pop out) it is enough to determine where it will be moved.

Suppose the director yells "freeze!" at exactly that moment. Our homunculus has just received intelligence that there is an interesting development in quadrant delta. The subject of the experiment is *about to* have attention drawn to quadrant delta, because that is the one region where apparent motion has a horizontal direction, rather than a vertical one. The appearance, and the singularity of its features in delta, must have already been registered; and that is *why* attention is about to be drawn to region delta. I suggest that not only is the subject not yet attending to developments in delta, but also that at that moment the subject is not yet aware of developments in delta.

It is worth laying out the argument for this claim. The first (essential) premise is that the situation in question in one in which we are testing for "pop-out": it is designed to test for the involuntary, stimulus-driven, exogenous "grabbing" of attention by sudden presentations of novel stimulus arrays. The critical time interval occurs immediately after the presentation of the new stimulus array, and it

is brief enough that the subject does not have time to shift attention deliberately. Instead attention is about to be grabbed by something, at least briefly. The question is what directs that exogenous "grabbing" of attention.

Now if attention shifts involuntarily in this way, to an abruptly presented novel stimulus, then one only becomes aware of that stimulus after attention has shifted to it. Because it is novel, it wasn't even in the "background" previously; one could not have attended to it at any moment prior to the moment of its presentation. After that moment, one attends to it as quickly as possible. So if we have "pop-out" of some feature, then one is not aware of that popped-out feature until one focuses attention on it.

The other interesting fact about pop-out is that it is not random. Attention is drawn to the unique instance of basic feature F, which stands out among the many instances of basic feature G. This implies that information about the difference between F and G, and about the uniqueness of the occurrence of F, must be registered and received by the mechanisms of selective attention, which then use it to direct attention to that one instance of F. So if we have "pop out" of some feature, then the presence of the feature must be detected by the selection mechanism, and registered in the salience map as the next target, *before* attention can be directed to it. It is so registered *in order* to direct attention to it.

Here then are the two premises:

1. If a feature pops out, then one is not aware of it until one focuses attention on it.

2. If a feature pops out, then its presence is detected by the selection mechanism, and it is registered in the salience map as the next target, before attention can be directed to it.

From these we get a simple conclusion:

> If a feature pops out, it can be registered and detected without awareness.

This is unremarkable until one notices that (as argued in section 4) the features F and G can include various pairs of phenomenal properties. F is "apparent motion in a horizontal direction" and G is "apparent motion in a vertical direction". That this F pops out in a field of those G's indicates that the difference between "being appeared to horizontal-wise" and "being appeared to vertical-wise" can be registered and detected, and can serve to direct attention, even though one is not (yet) aware of any of the stimuli that appear in any of those ways, or of being appeared-to in any such way. But what we have, then, is a state in which the subject is being appeared-to, in a particular fashion, without thereby being aware of any aspect of that appearance, or of being in such a state. So we have a state of "phenomenal consciousness" without the consciousness. Being appeared-to and being aware-of split apart. We have the former without the latter. States of phenomenal consciousness are not states of one thing; the notion splits in twain. "Phenomenal" shows up – it can pop out – prior to consciousness.

The suggestion may seem outlandish and weird, so it may be worth reviewing the stages of the argument again, using one of the other examples. Talk of "subjective contours" clearly employs the traditional terminology of phenomenal properties. The pac-men create the appearance of a square, with bright contours; but there are no line segments on the page. So a Kanizsa square is a phenomenal square: a mere appearance-of a square. Now we abruptly present a subject a novel stimulus array in which there are many groups of notched circles, but in only one such group are they aligned so as to generate the Kanizsa illusion. That group will "pop out" – it will draw attention to itself, in more or less constant time, no matter how many distractor groups are also present. So, the psychologists conclude, subjective contours are basic features that are processed preattentively, in parallel over the entire visual field. The system registers the appearance of a subjective square in region delta, prior to, and without, the ministrations of selective attention. Indeed, it is because that appearance is unique that attention is about to be drawn to that region. So, in these delicate moments, the subject is being appeared-to in a determinate fashion – so as be presented the appearance of a Kanizsa square in region delta – yet is not aware of those contours or of that illusory square. Being appeared-to can occur in a determinate way without the subject being aware-of that way.

I earlier posed the question: where, in the architecture of visual processing, are we to find states that have all the properties of states of phenomenal consciousness? Now we have an answer. If states that have phenomenal character split apart from those that yield awareness, then there is no one place in that architecture in which states have all the properties of states of phenomenal consciousness. Furthermore, states in the preattentive portion of the architecture can satisfy the sufficient conditions for the traditional notions of phenomenal character: states therein can be states of being appeared-to in a determinate way. Phenomenal character is found early, in preattentive feature processing. States in virtue of which the subject is aware of features of the ambient array presumably arise later, in the "attentional" or "post-attentional" portions of the architecture. They are certainly not yet found in the preattentive bits.

6 PHILOSOPHICAL IMPLICATIONS OF DIVORCE

Phenomenal character can be entirely divorced from awareness. I will call this the "divorce" hypothesis. It has implications for theoretical arguments found within both philosophy and psychology. Take the philosophical discussions first. A widely held assumption, common to all sides of various philosophical debates, must be abandoned. If the divorce hypothesis is true, then there is no necessary connection between phenomenal properties and consciousness. What havoc this finding wreaks within recent discussions! For philosophers almost invariably assume that understanding phenomenal character and understanding consciousness are two labels for the same endeavor.

Consider this sampling of recent discussions. As exhibit number one, here is Chalmers describing "the hard part" of the mind-body problem:

> A mental state is conscious if there is something it is like to be in that mental state. To put it another way, we can say that a mental state is conscious if it has a *qualitative feel*–an associated quality of experience. These qualitative feels are also known as phenomenal properties, or *qualia* for short. The problem of explaining these phenomenal properties is just the problem of explaining consciousness. This is the really hard part of the mind-body problem. [Chalmers, 1996, 4]

"The hard problem" is a wonderful phrase for marketing purposes, and it has engendered a sprawling literature. But we can now see that it has been ambiguous since its inception. Is the hard problem the problem of explaining phenomenal properties, or is it the problem of explaining consciousness of phenomenal properties? They are not the same thing! So when Chalmers says "Color sensations stand out as the paradigm examples of conscious experience" [Chalmers, 1996, 6] one must say: not always. To understand the phenomenal properties of chromatic appearance, it is not necessary that we solve the problem of consciousness. It is conceptual error to claim that sensory appearance is necessarily tied to conscious experience. Similarly, it is simply a mistake to say

> the phenomenal concept of mind ... is the concept of mind as conscious experience, and of a mental state as a consciously experienced mental state. [Chalmers, 1996, 11]

"Phenomenal" and "conscious" are *two* concepts, and they can be divorced. The preattentive visual architecture provides examples in which they *are* divorced.

Now to be fair to Chalmers, these propositions are not entirely his fault; the notion of "phenomenal" is itself ambiguous, and there *are* readings in which phenomenal properties are by definition confined to states of awareness. (They might be understood as the manner of appearance of one's own mental states when one is conscious of them, for example.) Much of the blame can be laid at the feet of Thomas Nagel, author of the fecund phrase "what it is like to be a bat". Nagel took this to express a difference between states of consciousness and other sorts of states, yet it has been widely understood as picking out the sprawling class of states that possess phenomenal character. Now Nagel was successful in articulating one of the murky features of our notion of consciousness: there is something it is like to have state M if and only if one is conscious of M. So the misunderstanding is predictable: if "what it is like" characterizes a phenomenal property, then phenomenal properties become properties exclusively of conscious mental states.

The phrase "phenomenal consciousness" was introduced into recent discussions by Ned Block, and some of these assumptions were found in that introduction. For example, while contrasting phenomenal consciousness (P consciousness) and access consciousness (A conscious) he says:

> it is in virtue of its phenomenal content or the phenomenal aspect of its content that a state is *P*-conscious, whereas it is in virtue of its representational content, or the representational aspect of its content, that a state is *A*-conscious. ... The *P*-conscious content of a state is the totality of the state's experiential properties, what it is like to be in that state. [Block, 1997, 383]

He also urged that "phenomenal properties cannot be *P*-unconscious" [406]. As with Chalmers, one can chalk this up to a different notion of "phenomenal content": the sort that Block has in mind is not the simple sort we find while looking innocently at Kanizsa squares. Instead Block's sort must explain what it is like to be in the state, e.g., one's own mode of presentation of the mental state itself. In fact Block has more recently allowed that "phenomenality" can be divorced from consciousness, so he would allow amendments for both such claims.

7 PERCEPTUAL PSYCHOLOGY AFTER DIVORCE

The possibility of divorcing phenomenal character from awareness also has some startling implications for some of the models and theoretical disputes among perceptual psychologists. Candidates for divorce are far more numerous than the three described in section 4.

For one, there are other preattentive processes that clearly involve phenomenal properties as traditionally understood. The examples of "shape from shading" and pictorial depth cues raise a number of hotly contested issues. For example: to what extent can perceptual grouping processes – the capacity to perceive a set of elements *as* a group, a unity, or a structured whole – proceed preattentively? Is it only form primitives that can be sensed preattentively, or can more complex structural relations of overall shape also be found? One surprise with shape-from-shading was that features of such complexity could be handled in early, automatic, parallel processes; it had been thought that 3d shape was a higher-level feature whose analysis was complex, and done later. But perceptual grouping processes in general and these shape effects in particular invoke classical phenomenal properties. When an ambiguous figure shifts in aspect – it looked like a duck, and now it looks like a rabbit – what changes is nothing in the stimulus itself, but how it appears. And so all the results that suggest preattentive capacities for grouping and segmentation could be added as additional examples of potentially preattentive phenomenal properties.

Divorce also allows some otherwise odd results in priming and masking studies to be redescribed in simple terms. One relatively direct way to approach the study of perception without awareness is to "mask" some stimuli, and test whether they still have a "priming" effect of some kind. Masking is used to eliminate awareness of the stimulus; priming to show that it is, nevertheless, perceived. A stimulus might be presented for a very short duration, and immediately followed by a different

stimulus, which perhaps shares a boundary with the first one. With durations cut short enough, subjects will not report being aware of the first stimulus at all. It may nevertheless affect later perceptions in ways that indicate it was perceived.

Many such "priming" effects can be found; for the purposes of this paper the most interesting derive from the congruity or incongruity of colors. Complementary colors, or more specifically, opponent colors (red *v.* green, yellow *v.* blue) are counted as "incongruous". Whereas red and orange share a red component, and are "congruous". If one is asked to name the color of a patch, and an incongruous color is flashed along the way, there is interference in producing the name: subjects slow down, and errors increase. Flashing a congruous color does not produce the interference. A related effect is called "Stroop" interference. The task is to name the color of a patch, and the interference is produced by also presenting an incongruous color *word*. The color patch might be green, but if you flash the word "red" many subjects will call it "red". What is interesting is that both sorts of results are also found if the presentation of the interfering stimulus is masked to the point where subjects do not report being aware of it (see [Cheesman and Merikle, 1985; 1986; Cohen *et al.*, 1995].

Incongruity of color is a good example of what the tradition would call a relation between appearances. "Being a complement of" is a phenomenal relation. Interference is explained by the incompatibility of appearances. If this account is to be viable, then the subject must apprehend the appearances. We get a delayed response, for example, because first the screen looks green, then it looks red; and green and red are incongruous. Now suppose we mask the first stimulus to the point where the subject is at chance in guessing whether there was or was not a stimulus present. We get a reduced effect size, but one that is still statistically significant. The proferred conclusion seems, at first, shocking. A subject can be appeared-to greenly, without being aware of the green, or of the stimulus that appears green, or even of the fact that there was a stimulus at all. But with the wrinkles ironed out of our conceptual wardrobe, these results can be now be clothed in style. They are just another example of being appeared-to without being aware-of. The subject is being appeared-to in a determinate fashion – to get the results, the character of that appearance must be, specifically, *green* – but the subject is not aware of that appearance. The sensory state has phenomenal character of which the subject is not conscious. This shows, as definitively as anything could, that color sensations are *not* the paradigm example of conscious mental states. Here we have what seem to be chromatic sensory states, but (apparently) no awareness of the color of the masked stimulus. Both conjuncts are subject to some doubt, but if confirmed, we have a case of color sensation with no consciousness thereof.

Third and last in the list of potentially redescribable phenomena are some syndromes found in the neuropsychology clinic. Blindsight and the neglect syndrome provide some of the most compelling evidence for dissociations between perception and awareness. In both, perceptual capacities seem to be retained even though awareness of what is perceived is lost. The two syndromes pose rather different

research problems. In blindsight the surprise is not so much that damage to area V1 causes a scotoma – a region of visual perimetry within which the patient cannot report upon the features, or even the presence or absence of, visual stimuli. The surprise is that even in the absence of "acknowledged awareness", patients proved capable of successful forced choice discriminations: of location, direction of motion (vertical v. horizontal), form primitives (X v O), and, most surprising of all, color. The focus of research has been on how this residual perceptual function is possible given the damage to $V1$. In contrast, the retention of perceptual capacity in neglect is not so problematic, and there the issue is rather why there is such a puzzling loss of awareness of stimuli in portions of the visual field (see [Driver and Vuilleumier, 2001]).

To take these in turn. Residual chromatic discrimination in blindsight was demonstrated by an arduous procedure of forced-choice discrimination of pairs of stimuli presented in the patient's scotoma. The experiments are done in what Weiskrantz calls the "pure unawareness" mode: the patient will report no awareness of any feature or any event within the scotoma. They cannot tell when a stimulus is present or absent, much less anything about its features, and must be prompted to respond. Nevertheless, over a long, long series of what seemed to the patients to be forced-choice guesses, chromatic discrimination can be demonstrated, and – the real stunner – sensitivity to different wavelengths can be shown to be close in form to a normal "spectral sensitivity" curve (see [Stoerig and Cowey, 1989; 1991; 1992]). In daylight (photopic) intensities, with cones active, the visual system is most sensitive to wavelengths around 550 nm. In scotopic conditions, (night vision, with just the rods active), the peak sensitivity shifts to roughly 500 nm. Cowey and Stoerig demonstrated the same shift in spectral sensitivity for stimuli presented within the scotoma of blindsight patients, and they demonstrated that in photopic conditions the spectral sensitivity shows peaks and valleys similar to the normal curve, indicating the contribution of more than one type of cone. So even though patients may not be aware of the presence of visual stimuli in the scotoma, when prompted to guess, their responses indicate some capacities to discriminate among wavelengths similar to normal ones. As Weiskrantz remarks:

> The latter capacity – discrimination of colours – presses credulity to the limit, because in those tests – which by their nature were very time-consuming and lasted for several days – the subjects uniformly and consistently denied seeing colour at all, and yet performed reliably above chance, even between wavelengths falling relatively close together. Moreover, the fine-grained features of the spectral sensitivity curves of these subjects (carried out, again, by forced-choice guessing) suggested that wavelength opponency, that is, colour contrast, was intact [Stoerig and Cowey, 1989; 1991; 1992]. The subjects seemed to be able to respond to the stimuli that would normally generate the philosophers' favourite species of 'qualia', namely colours, but in the absence of the very qualia themselves! [Weiskrantz, 1997, 23]

I have avoided the term "qualia" in this paper, preferring here the traditional notion of "phenomenal properties", but even in the latter terms we get what might seem a paradoxical result. The shift between photopic and scotopic sensitivity is widely thought to explain a perceptual effect known as the "Purkinje shift". In the daytime the red flower of a rose might look as bright as its green leaves, but as dusk falls the relative brightness of different colors appears to shift. The red flower will come to look black, and the green leaves will appear noticeably brighter than the flower. This is a shift in sensible appearance, in how things look; the change in relative brightness is a change of phenomenal properties. But the standard explanation for it appeals to the differing spectral sensitivity curves of photopic and scotopic vision. Photopic sensitivity peaks at 550 nm, and the red and green surfaces are equally far down from that peak, and so appear equally bright. Scotopic vision is maximally sensitive at about 500 nm, and the green leaf happens to reflect light maximally in that range. Whereas nighttime vision is quite insensitive to wavelengths reflected at the red end of spectrum. So at dusk the green leaf will look brighter than the red of the rose.

The Stoerig and Cowey results indicate that these mechanisms are at least partially intact in blindsight. The wavelength discriminations that survive could manifest the same shift in sensitivity between scotopic and photopic conditions. So it is conceivable that even within a scotoma, the green leaf looks brighter than the red rose in scotopic conditions, while they look equally bright in photopic conditions. If, as argued earlier, it is possible to be appeared-to in a particular way without being aware of any aspect of that appearance, then it is possible that G looks brighter than R, even though the subject is not aware of G, aware of R, or aware of any feature of G or R. Perhaps blindsight is precisely a case in which being appeared-to is dissociated from being aware-of.

Blindsight might or might not entail the absence of qualia, depending on whether one considers qualia to be phenomenal properties or the awareness of phenomenal properties. The latter sort are missing, but the former might still be present. The discrimination data summarized in spectral sensitivity curves would be taken to show that surfaces that appear equally bright to this subject in the daytime will not appear so at dusk, and that the one which reflects a preponderance of wavelengths around 500nm will look brighter than the one that reflects longer wavelengths. Now we find those same discrimination data replicated for stimuli presented within the scotoma of a patient with blindsight. Suppose being appeared-to can be divorced from being aware-of. It might then still be true of a subject that one stimulus looks brighter than another, though that subject has blindsight, and is not aware of either stimulus, or of any aspect of their appearances. Perhaps the problem is not loss of the appearances at all, but instead loss of awareness of them. Put in terms of "qualia", perhaps these subjects still have chromatic qualia, but are not aware of them.

Hemi-neglect provides an even better example of this possibility. Unlike blindsight, neglect can be demonstrated even in the absence of any damage to primary sensory pathways [Driver and Vuilleumier, 2001, 40]. Recent accounts of neglect

attribute the symptomatology to damage to mechanisms of selective attention, rather than to sensory loss. Indeed, Driver and Vuilleumier [2001] propose that the visual processing in neglect is similar to the preattentive processing of normals. The conjunctions and dissociations of symptoms are quite intriguing, and provide a good review of the architecture of early vision.

The basic syndrome is as follows:

> a lack of awareness for sensory events located towards the contralesional side of space (e.g. towards the left following a right lesion), together with a loss of the orienting behaviors, exploratory search, and other actions that would normally be directed toward that side. Neglect patients often behave as if half of their world no longer exists. In daily life, they may be oblivious to objects and people on the neglected side of the room, may eat from only one side of their plate, read from only one end of a newspaper page, and make-up or shave only one side of their face. The spatial bias towards one side can also be apparent in many simple paper and pencil tests. When required to search for and mark all target shapes on a page, the patients may cancel only those towards the ipsilesional side. When bisecting a horizontal line, they may err towards that side, and when drawing from memory, or copying a picture, they may omit details from the contralesional side. [Driver and Vuilleumier, 2001, 40]

Neglect is not all or none, but is graded, with the probability of neglecting a stimulus increasing as it is presented further and further in the contralesional direction. But "directions" and "the left side" are relative to a reference frame, and the reference frames in neglect are both variegated and surprisingly sophisticated. Suppose (to simplify the discussion) we have a patient with damage to right parietal cortex, who tends to neglect "the left side". This is not a loss of visual sensitivity to the left side of space; it is not hemianopia. Patients can still see things on the contralesional side; the problem is difficulty in shifting and maintaining attention to events on that side. Patients with "extinction" for example might still be able to see and consciously report on a light isolated on the left side, if it has no competition on the right; the problem arises only when two lights are presented simultaneously, with the one on the right then "extinguishing" awareness of the one on the left. This can happen even if both lights are presented in the right (intact) hemifield. So the reference frame cannot be retinotopic. (The latter is clear anyway, since it would not explain why the left side of plates, newspapers, and one's own face might be ignored. A patient with hemianopia can compensate simply by shifting the eyes. Some neglect patients have no visual field cut at all!) Furthermore, stimuli affecting the same retinal point will be neglected to varying degrees, depending on the position of the eyeball, the head and the trunk. A stimulus might be neglected when the eyes are aimed forward in their sockets yet detected when they are shifted sideways, even though the stimulus is produced so as to affect the same retinal area on both occasions. Similar results are found with

head motions and with trunk motions. So "the left side" might be the left side of the head, the left side of the trunk, or the left side of external space [Driver & Vuilleumier, 2001, 46].

The spatial coordinate systems defining which side is "the left side" can be rather sophisticated! Even more intriguing is that some of these coordinate schemes must be object-centered: what is ignored is the left side of *an object*, and so it all depends on what the neglect patient sees *as* an object. For example, if a picture shows one plant that has two flowers, a neglect patient who is asked to copy it may copy only the flower on the right. But suppose we occlude parts of the picture, so that the same two flowers are shown detached from one another, on independent stems, and ask the patient to copy the picture. He or she will carefully draw two flowers, but neglect details on the left side of each one (see [Bisiach and Vallar, 2000, 481]).

What happens if an object is tilted? Unless one tilts the head, the "left side" of the object will be tilted relative to every body-centered coordinate scheme. Driver [1996, 318] showed that neglect patients continue to neglect the "left side" of objects, even when the object is tilted, and the axis of symmetry defining the left side of the object is tilted relative to every body-centered coordinate scheme. He calls this "axis-based" neglect.

Neglect must be something other than a sensory deficit. In order to determine which bits of an object make up its left side, one must be able to perceive its entire contour. To neglect the left side of a plate, the left side of the flower, or the left sides of the two flowers, one must have some perception of the midline of these objects, and so one must perceive their entire contours. So neglect seems to require intact perception of contours, combined with a deficit in allocating attention. Driver and Vuilleumier propose that the main deficit is a bias in the competitions for selective attention. Neglect is most clearly manifest when multiple items are competing for selective attention; the bias favors items on the ipsilesional side, in whatever reference frame defines that "side".

Grouping effects can decrease neglect, on this account, because they decrease the competition for selective attention. The left side of the entire plant will be ignored, if the entire assemblage is perceived as one plant; if instead its flowers are perceived as two detached flowers, the left sides of both such objects are neglected. Given our earlier concerns it is particularly fascinating to note that Kanizsa subjective figures can secure this advantage as well. Mattingley, Davis and Driver [1997] did experiments on neglect patients very similar to the Davis and Driver 1994 experiments on preattentive pop-out. In some cases the "pac men" were arranged so that the notches in the circles were aligned, and a subjective square would be apparent to normal viewers. In other cases the notches were misaligned, or the pacmen had circular borders, which would defeat the appearance of the subjective square. The results are dramatic:

> extinction was virtually abolished when such arcs were removed, as the bilateral events then yielded a single subjective object (a bright white rectangle, apparently superimposed on the black circles) due to modal

> surface completion. This suggests that extinction is reduced when the concurrent target events can be linked together into a single subjective object, becoming allies rather than competitors in the bid to attract attention. [Driver and Vuilleumier, 2001, 52]

Subjective contours and subjective objects are paradigm examples of entities whose properties are phenomenal, and so it is gratifying to find such creatures of appearance in a place where awareness is less than full.

To common sense a neglect patient seems oblivious or insensible to the contralesional side of the current goings-on, where "current" changes with the context. Yet now we see that what appears to be blindness can be produced by failures in the mechanisms that regulate competition for selective attention. If a subject is truly incapable of shifting attention to anything visible on the far left side, however "left side" is defined, then in many respects that subject is as good as blind to developments in that region. If a problem with selective attention implies that visual sensory states registering such developments cannot ever win in the competitive allocation, then the subject remains perpetually unaware of such phenomena. Yet sensory capacities remain; this would be a kind of blindness deriving from a failure in attentional mechanisms.

What we have, then, might be the full preattentive sensory capacities without the capacities to engage the selection that leads to post-attentive processing. Like the sensory states found in preattentive processing, neglect patients could have states in which they are being appeared-to in a determinate fashion, even though they are entirely unaware of those appearances. Unlike blindsight patients, the neglect patient might have no visual scotomata, and primary sensory channels might be entirely intact.

8 CONCLUSIONS

Different authors mean different things by the term "phenomenal property". The traditional notion is one used to characterize sensory illusions and perceptual effects, and is already familiar, well-behaved, and well-ensconced in the models of perceptual psychologists.

Empirical models of perception provide explanations of at least some of these phenomenal properties. Some of these explanations cite neural mechanisms as explananda.

Phenomenal properties as traditionally understood are almost invariably accompanied with awareness. That is, to speak plainly, usually people are aware of how things appear to them. They will tell you if you ask. Furthermore, any example of a phenomenal property identifiable from the first person perspective is necessarily one of which the author is aware. But there is no conceptual necessity – no necessity linking the concepts of – states of being appeared-to and states of being aware-of.

So it is conceptually possible that there exist states of being appeared-to which are not, in any way, conscious states. Some of the states identified in the preattentive processing of apparent motion, subjective contours, and pictorial depth cues show that this possibility is real. Some of the states in blindsight and hemi-neglect may be examples of the same thing.

It is therefore a conceptual error to identify the problem of understanding phenomenal properties with the problem of consciousness. If that conjunction is identified as "the hard problem" of consciousness, then the solution to the hard problem of consciousness, like the solution to the problem of life, will be seen in the vanishing of the problem.

Psychologists already employ tools that are adequate to the understanding of phenomenal properties as traditionally understood. They provide explanations of some perceptual effects, and for all we know some of these putative explanations could be sound.

Philosophers who use the term "phenomenal properties" to beat up on psychologists mostly mean something higher-order, such as: how mental states appear to someone who is conscious of them. These higher-order notions, while endlessly fascinating, have yet to be adequately clarified.

If we confine ourselves to the task of explaining phenomenal properties as traditionally understood, then *enormous* progress has already been made within the relatively short lifespan of the discipline of experimental psychology. An optimist can be forgiven for believing that the tools for its solution may already be at hand.

BIBLIOGRAPHY

[Allport, 1993] A. Allport. Attention and control: Have we been asking the wrong questions? A critical review of twenty-five years. *Attention and Performance*, 14(3): 183–218, 1993.

[Bisiach and Vallar, 2000] E. Bisiach and G. Vallar. Unilateral neglect in humans. In G. Rizolatti (ed.), *Handbook of Neuropsychology*. 2nd edition. Vol 1. Series Editors Francois Boller and Jordan Grafman. Amsterdam: Elsevier, 458–502, 2000.

[Block, 1997a] N. Block. On a confusion about a function of consciousness. In N. Block, O. Flanagan, and G. Güzeldere (eds.), *The Nature of Consciousness: Philosophical Debates*. Cambridge, Massachusetts: MIT Press, 375–415, 1997. (Originally appeared in *Behavioral and Brain Sciences*, 18: 227–288, 1995).

[Block, 2003] N. Block. Consciousness, philosophical issues about. In L. Nadel (ed.), *The Encyclopedia of Cognitive Science*. London: MacMillan. Volume 1, 760–770, v.

[Broadbent, 1958] D. E. Broadbent. *Perception and Communication*. Oxford: Pergamon, 1958.

[Chalmers, 1996] D. J. Chalmers. *The Conscious Mind*. New York: Oxford University Press, 1996.

[Cheesman and Merikle, 1986] J. Cheesman and P. M. Merikle. Distinguishing conscious from unconscious perceptual processes. *Canadian Journal of Psychology*, 40(4): 343–367, 1986.

[Cheesman and Merikle, 1985] J. Cheesman and P. M. Merikle. Priming with and without awareness. *Perception and Psychophysics* 36: 387–395, 1985.

[Chisholm, 1957] R. Chisholm. *Perceiving: A Philosophical Study*. Ithaca: Cornell University Press, 1957.

[Cohen *et al.*, 1995] A. Cohen, R. B. Ivry, R. D. Rafal and C. Kohn. Activating response codes by stimuli in the neglected visual field. *Neuropsychology*, 9(2): 165–173, 1995.

[Davis and Driver, 1998] G. Davis and J. Driver. Kanizsa subjective figures can act as occluding surfaces at parallel stages of visual search. *Journal of Experimental Psychology: Human Perception and Performance*, 24(1): 169–184, 1998.

[Davis and Driver, 1994] G. Davis and J. Driver. Parallel detection of Kanizsa subjective figures in the human visual system. *Nature*, 371 (27 October) 791–793, 1994.
[Desimone and Duncan, 1995] R. Desimone and J. Duncan. Neural mechanisms of selective visual attention. *Annual Review of Neuroscience*, 18: 193–222, 1995.
[Di Lollo *et al.*, 2000] V. Di Lollo, J. T. Enns and R. A. Rensink. Competition for consciousness among visual events: the psychophysics of re-entrant visual processes. *Journal of Experimental Psychology: General*, 129(4): 481–507, 2000.
[Driver, 1996] J. Driver. The neuropsychology of spatial attention. In H. Pashler (ed.), *Attention*. London: University College London Press, 297–340, 1996.
[Driver and Vuilleumier, 2001] J. Driver and P. Vuilleumier. Perceptual awareness and its loss in unilateral neglect and extinction. *Cognition*, 79: 39–88, 2001.
[Duncan, 1998] J. Duncan. Converging levels of analysis in the cognitive neuroscience of visual attention. *Philosophical Transactions of the Royal Society of London*, B 353: 1307–1317, 1998.
[Enns, 1990] J. T. Enns. Three dimensional features that pop out in visual search. In D. Brogan (ed.), *Visual Search*. London: Taylor and Francis, 37–45, 1990.
[Enns and Rensink, 1990] J. T. Enns and R. A. Rensink. Influence of scene-based properties on visual search. *Science*, 247 (Feb. 9, n4943): 721–723, 1990.
[Enns and Rensink, 1991] J. T. Enns and R. A. Rensink. Preattentive recovery of three dimensional orientation from line drawings. *Psychological Review*, 98(3): 335–351, 1991.
[Enns and Rensink, 1993] J. T. Enns and R. A. Rensink. A model for the rapid interpretation of line drawings in early vision. In D. Brogan, A. Gale, and K. Carr (eds.), *Visual Search II*. London: Taylor and Francis, 73–90, 1993.
[Farah, 1990] M. Farah. *Visual Agnosia: Disorders of Object Recognition and What They Tell Us about Normal Vision*. Cambridge, Massachusetts: MIT Press, 1990.
[Fodor, 1983] J. A. Fodor. *The Modularity of Mind*. Cambridge, MA: MIT Press, 1983.
[Goodman, 1977] N. Goodman. *The Structure of Appearance*. 3rd edn. Boston: Dordrecht Reidel, 1977.
[Humphreys *et al.*, 1994] G. W. Humphreys, N. Keulers, and N. Donnnelly. Parallel visual coding in three dimensions. *Perception*, 23(4): 453–470, 1994.
[Ivry and Cohen, 1990] R. B. Ivry and A. Cohen. Dissociation of short and long-range apparent motion in visual search. *Journal of Experimental Psychology: Human Perception and Performance*, 16(2), 317–331, 1990.
[Julesz, 1984] B. Julesz. A brief summary of the texton theory of human vision. *Trends in Neuroscience*, 7: 41–45, 1984.
[Julesz and Bergen, 1983] B. Julesz and J. R. Bergen. Textons, the fundamental elements in preattentive vision and perceptions of texture. *Bell Systems Technical Journal*, 62: 1619–1646, 1983.
[Livingstone and Hubel, 1987] M. S. Livingstone and D. H. Hubel. Psychophysical evidence separate channels for the perception of form, color, movement, and depth. *Journal of Neuroscience*, 7: 3416–3468, 1987.
[Luck *et al.*, 1996] S. J. Luck, E. K. Vogel and K. L. Shapiro. Word meanings can be accessed but not reported during the attentional blink. *Nature*, 383: 616–618, 1996.
[Marr, 1982] D. Marr. *Vision*. San Francisco: W. H. Freeman and Company, 1982.
[Mattingley *et al.*, 1997] J. B. Mattingley, G. Davis and J. Driver. Preattentive filling-in of visual surfaces in parietal extinction. *Science*, 275: 671–674, 1997.
[Milner and Rugg, 1992] A. D. Milner and M. D. Rugg. *The Neuropsychology of Consciousness*. London: Academic Press, 1992.
[Nagel, 1974] T. Nagel. What is it like to be a bat? *Philosophical Review*, 83: 435–50. Reprinted in Nagel, 1979, 165–180, 1974.
[Nagel, 1979] T. Nagel. *Mortal Questions*. Cambridge: Cambridge University Press, 1979.
[Neisser, 1967] U. Neisser. *Cognitive Psychology*. New York: Appleton Century Crofts, 1967.
[Posner, 1994] M. I. Posner. Attention: the mechanisms of consciousness. *Proceedings of the National Academy of Science USA*, 91 (August): 7398–7403, 1994.
[Posner and Petersen, 1993] M. I. Posner and E. E. Petersen. The attentional system of the human brain. *Annual Review of Neuroscience 13*: 25-42, 1990.
[Posner and Rothbart, 1992] M. I. Posner and M. K. Rothbart. Attentional mechanisms and conscious experience. In A. D. Milner and M. D. Rugg (eds.), *The Neuropsychology of Consciousness*. London: Academic Press, 91–111, 1992.

[Ramachandran, 1988] V. S. Ramachandran. Perceiving shape from shading. *Scientific American*, 259: 76–83, 1988.

[Ramachandran, 1987] V. S. Ramachandran. Visual perception of surfaces: A biological theory. In S. Petry and G. E. Meyer (eds.), *The Perception of Illusory Contours*. New York: Springer Verlag, 93–108, 1987.

[Raymond *et al.*, 1992] J. E. Raymond, K. L. Shapiro and K. M. Arnell. Temporary suppression of visual processing in an RSVP task: an attentional blink? *Journal of Experimental Psychology: Human Perception and Performance*, 18: 849–860, 1992.

[Raymond *et al.*, 1995] J. E. Raymond, K. L. Shapiro and K. M. Arnell. Similarity determines the attentional blink. *Journal of Experimental Psychology: Human Perception and Performance*, 21: 653–662, 1995.

[Rensink and Enns, 1995] R. A. Rensink and J. T. Enns. Pre-emption effects in visual search: Evidence for low level grouping. *Psychological Review 102*(1): 101-130, 1995.

[Rensink, 2003] R. A. Rensink. Visual attention. In L. Nadel (ed.), *The Encyclopedia of Cognitive Science*. London: MacMillan, vol 4., 509–515, 2003.

[Rensink, 2000] R. A. Rensink. Seeing, sensing, and scrutinizing. *Vision Research*, 40: 1469–1487, 2000.

[Sellars, 1963] W. Sellars. *Science, Perception, and Reality*. London: Routledge & Kegan Paul, 1963.

[Stoerig and Cowey, 1989] P. Stoerig and A. Cowey. Wavelength sensitivity in blindsight. *Nature*, 342: 916–918, 1989.

[Stoerig and Cowey, 1991] P. Stoerig and A. Cowey. Increment threshold spectral sensitivity in blindsight: Evidence for colour opponency. *Brain*, 114: 1487–1512, 1991.

[Stoerig and Cowey, 1992] P. Stoerig and A. Cowey. Wavelength sensitivity in blindsight. *Brain*, 115: 425–444, 1992.

[Treisman, 1988] A. Treisman. Features and objects: The fourteenth annual Bartlett memorial lecture. *Quarterly Journal of Experimental Psychology*, A 40: 201–237, 1988.

[Treisman, 1993] A. Treisman. The perception of features and objects. In A. Baddeley and L. Weiskrantz (eds.), *Attention: Selection, Awareness, and Control: A Tribute to Donald Broadbent*. Oxford: Clarendon Press, 5–35, 1993.

[Treisman, 1996] A. Treisman. The binding problem. *Current Opinion in Neurobiology*, 6: 171–178, 1996.

[Treisman, 1998] A. Treisman. Feature binding, attention and object perception. *Philosophical Transactions of the Royal Society of London*, B 353: 1295–1306, 1998.

[Treisman and Gelade, 1980] A. Treisman and G. Gelade. A feature-integration theory of attention. *Cognitive Psychology*, 12: 97–136, 1980.

[Umilta, 2000] C. Umilta. Visuospatial attention. In G. Rizolatti (ed.), *Handbook of Neuropsychology*. 2nd edition. Vol 1. Series Editors Francois Boller and Jordan Grafman. Amsterdam: Elsevier, 393–425, 2000.

[Weiskrantz, 1997] L. Weiskrantz. *Consciousness Lost and Found*. Oxford: Oxford University Press, 1997.

[Wolfe *et al.*, 1989] J. M. Wolfe, K. R. Cave and S. L. Franzel. Guided search: an alternative to the feature integration model for visual search. *Journal of Experimental Psychology: Human Perception and Performance*, 15(3): 419–433, 1989.

[Wolfe *et al.*, 1992a] J. M. Wolfe, S. R. Friedman-Hill, M. I. Stewart and K. M. O'Connell. The role of categorization in visual search for orientation. *Journal of Experimental Psychology: Human Perception and Performance*, 18(1): 34–49, 1992.

[Wolfe, 1994] J. M. Wolfe. Guided Search 2.0: A revised model of visual search. *Psychonomic Bulletin and Review*, 1(2): 202–238, 1994.

[Wolfe, 1996a] J. M. Wolfe. Extending guided search: why guided search needs a preattentive "item map". In A. F. Kramer, M. G. H. Cole, and G. D. Logan (eds.), *Converging operations in the study of visual selective attention*. Washington, DC: American Psychological Association, 247–270, 1996.

[Wolfe, 1996b] J. M. Wolfe. Visual search. In H. Pashler (ed.), *Attention*. London: University College London Press, 13–73, 1996.

[Wolfe and Bennett, 1997] J. M. Wolfe and S. C. Bennett. Preattentive object files: shapeless bundles of basic features. *Vision Research*, 37(1): 25–43, 1997.

CONSCIOUSNESS: PHENOMENAL CONSCIOUSNESS, ACCESS CONSCIOUSNESS, AND SCIENTIFIC PRACTICE

Uriah Kriegel

1 INTRODUCTION: CONSCIOUSNESS AND THE PHILOSOPHY OF SCIENCE

Analytic philosophy of mind is a relatively new branch of philosophy. Some treat it as an extension of the philosophy of language, as concerned with the relation between representation and reality. Some treat it as a chapter of metaphysics, concerned with the categories of mental existence (mental states, mental substance, mental causation, etc.). And some treat it as a chapter of the philosophy of science, namely, the chapter concerned with the philosophy of psychology and the cognitive sciences.

A classic example of the last approach is Fodor's early work on psychological explanation [Fodor, 1968], the language of thought [Fodor, 1975], and the modularity of mind [Fodor, 1983], which has been extremely influential, in content but also in style. The *style*, or *form*, of argument was always this: cognitive science presupposes such-and-such psychological structure; therefore, (plausibly) the mind exhibits the structure in question. Thus, we were told that thought is conducted in a language-like medium, featuring syntactically structured and semantically evaluable items, because this is presupposed in "the only cognitive science we've got."

Curiously, although work in the philosophy of mind from the angle of the philosophy of science has been extremely influential in discussions of the representational and functional aspects of mental life, it has seen less influence in discussions of phenomenal aspects. What (relatively) little work on consciousness from the angle of philosophy of science has been done has rarely been of great influence on the mainstream discourse in an otherwise vibrant realm of consciousness research.

One rare exception is Ned Block's [1995] "On a Confusion About the Function of Consciousness," which did impress a lasting and deep mark on mainstream issues concerning philosophical theories of consciousness. In that paper, Block argues, in essence, that current scientific practice in consciousness studies has been targeting the wrong phenomenon. After distinguishing between what he calls

Handbook of the Philosophy of Science. Philosophy of Psychology and Cognitive Science
Volume editor: Paul Thagard
General editors: Dov M. Gabbay, Paul Thagard and John Woods
© 2007 Elsevier B.V. All rights reserved

access consciousness and *phenomenal consciousness*, Block considers carefully the main scientific approaches to consciousness and argues that they can only be taken to shed light on the former, not the latter.

In this chapter, I want to engage Block's argument and ultimately vindicate current scientific practice. I will argue that although phenomenal consciousness and access consciousness are logically independent, as Block indeed claims, there is still some intimate connection between the two: access consciousness is a dispositional property, and when its categorical basis is correctly identified, it is seen to be a component of phenomenal consciousness. That is, although access consciousness is separate from phenomenal consciousness, its categorical basis is not. This would vindicate current scientific practice in consciousness studies, in that it would suggest that in studying the phenomena they do, cognitive scientists are targeting an essential component of phenomenal consciousness, namely, the categorical basis of access consciousness. Once the case is made for this general meta-theoretical approach, a number of case studies from empirical work in cognitive psychology and neuroscience will be examined in its light. Hopefully, this exercise will contribute to the elucidation of the conceptual foundations of consciousness studies in the cognitive sciences.

2 THE PHENOMENOLOGICAL STRUCTURE OF PHENOMENAL CONSCIOUSNESS

In this section, I will present — somewhat dogmatically — a specific conception of the phenomenological structure of phenomenal consciousness. This conception is by no means uncontroversial, but as I will argue toward the end of the section, the particular way in which it is controversial should not affect the main argument of this paper. Elsewhere, I have argued in greater detail for this conception (see mainly [Kriegel, 2004], but also [Kriegel 2003b; forthcoming]); here I will only recapitulate on its main tenets.

Looking at the blue sky on a sunny summer day, I have a distinctive conscious experience of it. This experience has many properties, but the one property we find scientifically mystifying is its *phenomenal character*: there is something it is like for me to have or undergo this experience [Nagel, 1974]. More specifically, there is a bluish way it is like for me to have my sky experience. The specific phenomenal character of my conscious experience is given by this bluish way it is like for me to have it.

As Levine [2001] has suggested, there is a distinction to be made between two aspects, or components, of this "bluish way it is like for me." On the one hand, there is the *bluish* aspect. On the other hand, there is the *for-me* aspect. Let us call the former (the bluishness) *qualitative character* and the latter (the for-meness) *subjective character*. Here is how Levine [2001, 6-7] puts it:

> There are two important dimensions to my having [a] reddish experience. First, ... there is something it's like for me to have this experi-

> ence. Not only is it a matter of some state (my experience) having some feature (being reddish) but, being an experience, its being reddish is "for me," a way it's like *for me*... Let's call this the *subjectivity* of conscious experience. The second important *dimension* of experience that requires explanation is qualitative character itself. Subjectivity is the phenomenon of there being *something* it's like for me to see the red diskette case. Qualitative character concerns the "*what*" it's like for me: reddish or greenish, painful and pleasurable, and the like.

We may construe phenomenal character as the compresence of qualitative character and subjective character.[1]

To say that my experience has a bluish qualitative character is to attribute to my experience the property of exhibiting a certain specific sensuous quality. It is not to say that the property in question is irreducible, or intrinsic, or inexplicable. It is merely to assert the *existence* of that property.

To say that my experience has a subjective character is to say that I am somehow *aware* of my experience. Conscious experiences are not sub-personal states, which somehow take place *in* us and which we "host" in an impersonal sort of way, without being aware of them. Mental states we are completely unaware of are unconscious states. So when I have my conscious experience of the sky, I must be aware of having it. In this sense, my experience does not just take place *in me*, it is also *for me*. Again, by asserting the existence of the property of subjective character, I do not mean to imply that it is irreducible. (Indeed, Elsewhere I defend a reductive account of subjective character [Kriegel, 2003a; Forthcoming].)

The notion that something like subjective character is indeed a crucial component of phenomenal character would be readily resisted by many. It is surely not uncontroversial. Let me therefore say a little more about how I conceive of subjective character, and finally why its postulation will not distort my discussion of Block's argument in '*Confusion*'.

Subjective character is a property conscious states have in virtue of the subject's awareness of them. This sort of awareness we have of our concurrent conscious experiences is clearly a somewhat elusive phenomenon. After all, normally we are not focused on our ongoing experiences, but rather on the world experienced therethrough (if you will). When I look at the sky, I am focused on the sky, not on my experience of it. And yet I am not altogether unaware of my experi-

[1] It is interesting to note that philosophers seem divided on which of the two components of phenomenal consciousness is the more fundamental. Evidently, those philosophers who offer representational theories of consciousness [Dretske, 1995; Tye, 1995; 2000] seem to conceive of qualitative character as more basic. Those who offer higher-order monitoring theories [Armstrong, 1968; Lycan, 1990; Rosenthal, 1997; Carruthers, 2000] seem to focus on subjective character. Rosenthal [1991] has argued explicitly in favor of distinguishing consciousness from what he calls "sensory quality," which appears to be the same as what we have called here "qualitative character." I have in several places defended what I call the *self-representational theory* of consciousness [Kriegel, 2003a; 2003b; 2006; forthcoming], according to which (approximately) a mental state is conscious in virtue of representing itself. This sort of view also conceives of consciousness as primarily a matter of subjective character.

ence. If I were completely unaware of my experience, it would not be a *conscious* experience.[2]

This raises the challenge of how to account for the elusive awareness we have of our ongoing experiences. One way of doing that is in terms of what I have called elsewhere [Kriegel, 2003b; 2004] *intransitive self-consciousness*, which is to be distinguished from *transitive self-consciousness*. To see the distinction, consider the following two reports:

1. I am self-conscious of my sky experience.

2. I am self-consciously experiencing the sky.

In the first report, self-consciousness appears as a verb, which takes the experience as an *object*. In the second report, self-consciousness appears as an adverb, which merely *modifies* the experience term. The former reports the occurrence of a second-order state that makes me focused on my experience. The latter reports the occurrence of a first-order state that makes me focused rather on the sky, and makes a comment on the *way* I am having this first-order state. The comment it makes is that I have the experience in a self-conscious sort of way.

But how are we to understand this self-conscious sort of way I am focused on the sky? My suggestion — developed, again, elsewhere [Kriegel, 2004; Forthcoming] — is that we construe this in terms of a distinction between *focal* and *peripheral* self-consciousness or self-awareness. The distinction between focal and peripheral awareness is widely applied to *perceptual* awareness. Thus, I can say that I am now focally visually aware of the laptop before me and peripherally visually aware of an ashtray on the corner of my desk. Likewise, when I listen to a piano concerto I am focally auditorily aware of the piano and peripherally auditorily aware of the cellos.[3] My claim is that the same distinction extends to second-order awareness and/or self-awareness: I can be focally aware of my experience of the piano concerto, as when I explicitly and deliberately introspect my ongoing auditory experience, or I can be merely peripherally aware of my experience of the concerto, as when I am focused on the pianist's interpretation of the piece but am nonetheless aware in that elusive and unimposing way of undergoing the experience. For further and more detailed exposition of this approach to the subjective character, or for-me-ness, of conscious experience, I refer the reader to Kriegel 2004.

As noted above, the very existence of this sort of elusive awareness, or for-me-ness, can be readily called into question. I cannot within the confines of the present paper present the full case for the psychological reality of subjective character (but

[2] And on a certain conception of the ordinary meaning of "experience," it would not be an experience at all. According to this conception, "experience" is always, and trivially, conscious — because this is just how the term works. In this paper, I wish to remain neutral between the view that this is so [Strawson, 1994] and the view that it is not [Carruthers, 1989; 2000].

[3] In fact, the case of auditory awareness is a better example than that of visual awareness, since it does not depend of the physiological structure of the organ. In visual awareness, the distinction between the focal and the peripheral is determined by the structure of the eye, in particular the fact that the fovea determines what will and what will not be focally seen.

see [Kriegel, 2004]). However, this should not matter overmuch to the discussion below. As we will see in the next section, Block admits the existence of what he calls "me-isness," which we would be pardoned to consider fundamentally the same as the "for-me-ness" discussed above. Block is perfectly happy, then, to allow for the fact that phenomenal experiences have a mine-ness built into them, a sort of immediate, built-in, and perhaps pre-conceptual self-ascription. His argument is therefore supposed to go through despite the existence of something like subjective character or for-me-ness. The argument does not depend on there not being such a thing in the phenomenology of conscious experience. Indeed, *how* the argument goes through despite that fact is something Block discusses explicitly. Let us turn now to consider Block's argument.

3 PHENOMENAL CONSCIOUSNESS, ACCESS CONSCIOUSNESS, AND SUBJECTIVE CHARACTER

According to Block [1995], cognitive scientists who work on consciousness are wrong to study the phenomena they do, because these phenomena can only shed light on a something that is inessential to phenomenal consciousness.

Block distinguishes two notions of consciousness: phenomenal consciousness and access consciousness.[4] Phenomenal consciousness is defined in terms of what it is like for the subject to have the conscious experiences she does. It is what generates the "explanatory gap" [Levine, 1983] or the "hard problem" [Chalmers, 1995], and is thus what we are truly struggling to understand. Thus, Block [1995, 382, italics original] writes: "I mentioned the explanatory gap partly by way of pointing to P-consciousness (i.e., phenomenal consciousness): *that's* the entity to which the mentioned explanatory gap applies."[5] Access consciousness, by contrast, is defined in broadly functionalist terms: a mental state is access-conscious just in case "it is poised for free use in reasoning and for direct 'rational' control of action and speech" [Block, 1995, 382].

It appears that the distinction itself is sound, at least as a *conceptual* distinction. A mental state is phenomenally conscious just in case there is something it is like for its subject to have it. There are no obvious conceptual ties between that and poise for free use in reasoning and action control. But Block's claim is that these are not only two separate *concepts*, but also two separate *properties*, and that current scientific research into consciousness focuses on the property of access consciousness *at the expense of* the property of phenomenal consciousness. He writes (Ibid.; italics original):

> . . . it is not easy to see how current approaches to P-consciousness [i.e., phenomenal consciousness] *could* yield an account of it. Indeed, what

[4] Actually, Block further distinguishes these two phenomena from what he calls *monitoring consciousness* and *self-consciousness*. But his argument focuses primarily on phenomenal and access consciousness.

[5] Here, and in what follows, I am quoting from the reprint in Block *et al.* [1997].

> passes for research programs on consciousness just is a combination of cognitive psychology and explorations of neurophysiological syndromes that contain no theoretical perspective on what P-consciousness actually is.

According to Block, current scientific research can only shed light on access consciousness, but access consciousness is not the source of the mystery of consciousness. So current scientific theories of consciousness do not contribute toward the demystification of consciousness. Only theories that would target phenomenal consciousness might possibly do that. We may call this the *access thesis*: Current scientific research focuses on access consciousness *instead of* phenomenal consciousness.

The overall argument of Block's paper is thus the following: 1) There is a difference between phenomenal consciousness and access consciousness; 2) Cognitive-scientific studies have targeted access consciousness; therefore, 3) Cognitive-scientific studies have failed to target phenomenal consciousness.

Some philosophers have argued that phenomenal and access consciousness are in reality not entirely separate properties, but rather entertain certain conceptual, internal, or otherwise non-contingent relations (see, e.g., [Dennett, 1995; Clark, 2000]). This is not the line I would like to take here, although it has certain similarities which will come out momentarily. The line I will pursue is that there is a conceptual or internal connection between phenomenal consciousness and the *categorical basis* of access consciousness. In other words, I will argue against the inference from the premises to the conclusion in Block's overall argument. This is different from the common tack, of arguing against Premise 1.

The first thing to note about access consciousness is that it is a *dispositional property*. Nothing has to actually happen with a mental state or event for it to qualify as access-conscious: the state or event need not be actually access*ed*; it only needs to be access*ible*. That is, in order to become access-conscious, a mental state need not be *actually* used in the control of reasoning and action; it need only be *poised for* use in such control.

The problem is that the property investigated in the cognitive sciences under the heading of consciousness studies is clearly *not* a dispositional property. When a mental state becomes conscious, there is something very real and categorical that happens to it and to the subject's relation to it. This suggests that Block is quite right to distinguish the *property* of phenomenal consciousness from the *property* of access consciousness. But it also suggests that there is something wrong in taking access consciousness to be the object of scientific investigation in consciousness studies.

Plausibly, dispositional properties are normally surrounded by two kinds of closely related non-dispositional property. There are, first, what may be called *manifestational* properties: the properties of manifesting the dispositions in question. Thus, mental states are often not only *poised for* use in reasoning and action control, but actually *are* so used. They then instantiate the manifestational property corresponding to access consciousness.

More interestingly, dispositional properties are often taken to require, as a rule, *categorical bases*. A categorical basis is a non-dispositional, occurrent property that accounts for and grounds certain dispositions. When a wine glass is fragile, its fragility cannot be a brute and inexplicable property. The fact that the glass is fragile is not an irreducible, *sui generis* fact. On the contrary, it must be possible to explain *why* the wine glass is fragile in terms of the physico-chemical properties of the glass it is made of. The glass is fragile – it is disposed to break under relatively lax conditions — *because* , or *in virtue of*, its physico-chemical constitution. Its particular constitution is thus the *reason for* its fragility — the *reason why* it is fragile. In this sense, the glass' physico-chemical constitution is the categorical basis of the glass' fragility.

Some philosophers have argued that all properties are at bottom dispositional (e.g., [Shoemaker, 1979]), a bundle of causal propensities and nothing more. If so, dispositional properties do not in fact require any categorical, non-dispositional bases. A full discussion of this view of properties — a sort of functionalism about everything — and of the problems attending it will take us too far afield. Perhaps it suffice that we recite here Russell's clever condemnation of this view as providing us nothing more than "a causal skeleton of the world." In this paper, I will assume that this "pure disposition view" is incorrect, and that dispositional properties do require categorical bases, that is, non-dispositional properties whose instantiation *explains why* the dispositional ones are instantiated.[6]

This applies, of course, to access consciousness. When a mental state is access-conscious, it must also have a categorical property *in virtue of which* it is access-conscious. There must be an explanation *why* the state is poised the way it is for free use by the subject, an explanation appealing to non-dispositional properties that *account* for the state's disposition to be freely used in that way.

What could this property be? A natural suggestion is that it is subjective character. The reason why the state is poised for the subject's free use in reasoning and action control is that the subject becomes *aware*, in that elusive and peripheral manner discussed in the previous section, of the state. Once the subject is aware of having her state, if ever so peripherally and dimly, she can freely make use of it in reasoning and action control. Thus the state's free usage to those ends can be *explained* in terms of its for-me-ness or subjective character. It appears that the subject's awareness of her conscious state is the *reason why* the state is poised for use in reasoning and action control — the *reason for* the state's poise for such use. It is *because* (or *in virtue of* the fact that) I am aware of my bluish experience of the sky that the experience is poised to be freely used in my reasoning about the consistently nice weather and in guiding my vacation plans. That is to say,

[6]Not that the view that dispositional properties do not always require categorical bases can be held without holding that all properties are dispositional. One might hold that some properties are dispositional and some categorical, and while most dispositional properties are anchored, or based, in corresponding categorical properties, some are not. This view appears somewhat unmotivated. In any case, I will not discuss these issues in the present paper. Instead, I will *assume* the correctness of the traditional view that all dispositions are based in categorical properties.

the categorical basis of access consciousness seems to be subjective character, the subject's special, peripheral awareness of her conscious state. This awareness is what *explains*, in non-dispositional terms, a conscious state's disposition to be freely used in reasoning and action control.

In other words, what explains the subject's ability to use conscious information freely is the fact that she is aware of that information. If she was not aware of it, so that the information was merely sub-personal, she would not be in a position to use it so freely.

It might be objected that categorical bases are to be found at the scientific, "micro" level, whereas subjective character is a property at the commonsensical, "macro" level. But the supposition that a categorical basis *must* be a micro property is misguided. Consider explosiveness: the property of being disposed to explode. It has a very clear categorical basis at the macro level — the property of containing gun powder. To be sure, it has *also* a micro-level categorical property, namely, its containing potassium nitrate. And that is because gun powder is effective in the way it is due to the chemical nature of potassium nitrate. But it does not follow from the fact that something is explosive in virtue of containing potassium nitrate that it is not also explosive in virtue of containing gun powder (in the same sense of "in virtue," the categorical-basis sense). Although the citation of potassium nitrate accounts for the explosiveness in a deeper and fuller way, the citation of gun powder accounts for it as well.[7]

My contention, then, is that the subjective character of conscious experience is the *categorical basis* of its access consciousness. Now, as I claimed in the previous section, the subjective character of a conscious experience is a component — indeed an essential component — of its phenomenal character. Happily trivially, for-me-ness is a component of what-it-is-like-for-me-ness.

If this is indeed the case, and a component of phenomenal consciousness is the categorical basis of access consciousness, then this can be taken to vindicate current scientific practice. The claim would be that, in conducting the studies they do, scientists are probably not targeting what Block identified as access consciousness, but rather the categorical basis of access consciousness. Since the categorical basis of access consciousness is a component of phenomenal consciousness, by targeting the categorical basis of access consciousness scientists are ultimately pursuing the study of phenomenal consciousness.

It is quite common in the history of science that scientists labor around a dispositional property by way of trying to learn about its categorical basis. Thus, for centuries geneticists have been studying hereditary properties, which are dispositional, by way of trying to reach an understanding of their categorical basis, which we have only recently identified as DNA. A good example is the so-called Huntington Disease, an inherited neurological degenerative disorder characterized by loss of striatal neurons. Research into Huntington Disease has led to the discovery in 1993 that a mutation of the CAG gene — a mutation in which the triplet repeats

[7]For discussion of this point, I would like to thank Orlin Vakarelov, Stephen Biggs, Farid Masrour, and Keith Lehrer.

at least 42 times (as opposed to between 11 and 34 times in the normal case) — is what causes the disease (see [Huntington's Disease Collaborative Research Group, 1993]). In this case, what drove the study of the disposition to suffer from Huntington Disease is a practical interest in identifying what causes the disposition's manifestation. But what motivates the study is beside the point. The important fact is that the study resulted in learning more and more about the properties of the disposition's categorical basis, until the particular gene responsible for it could be singled out.

Research into access consciousness can be seen in a similar light. By looking at what causes this disposition's manifestation (the manifestation being the actual use of a state in reasoning and action control), more and more can be learned about the properties of the disposition's categorical basis, namely, subjective character.

One advantage of this view on the relation between access and phenomenal consciousness — that far from being completely independent of each other, the latter (or a component thereof) is the categorical basis of the former – is that it accounts for the functional role of phenomenal consciousness. Indeed, it accounts for there *being* a functional role to phenomenal consciousness.

A problem with Block's distinction is that any function we may wish to attribute to phenomenal consciousness would be more appropriately attributed to access consciousness, leaving phenomenal consciousness devoid of any functional significance it can properly call its own [Chalmers, 1997]. The source of this unhappy consequence is the picture of phenomenal and access consciousness as two separate properties sitting side by side at the same theoretical level. But if, as I have argued, phenomenal consciousness (or part of it) is the categorical basis of access consciousness, then the latter can be readily construed as the functional role *of* phenomenal consciousness. That is, phenomenal consciousness is the *occupant* of a functional role, part of the specification of which is given by access consciousness (namely, the part concerned with the poise for free use in reasoning and action control). Here, the relation between phenomenal and access consciousness is construed as the relation of an occupant to its role: phenomenal consciousness *occupies*, or *plays*, access consciousness (if you will). Thus once we construe phenomenal consciousness as the categorical basis of access consciousness, and access consciousness as the functional role of phenomenal consciousness, we can again attribute certain functions to phenomenal consciousness: the functions are construed as *part of* access consciousness and as *performed by* phenomenal consciousness. This appears to avoid potential conceptual confusions caused by fully divorcing phenomenal from access consciousness.

It is, of course, open to someone like Block to claim that phenomenal consciousness is in fact devoid of any functional significance (as [Velmans, 1992] has done) or has very limited functional significance (as [Libet, 1985] suggests). But I take it that such epiphenomenalism, hard or soft, is a *liability* on a theory of phenomenal consciousness, one better avoided when possible.

A big part of Block's argument for the full divorce between phenomenal and access consciousness is his claim that there are possible circumstances in which

one can occur in the absence of the other. In particular, access consciousness could occur in the absence of phenomenal consciousness in what Block calls super-blindsight, and the converse may occur in certain cases of perception of habituated stimuli.

These particular cases will be examined in some detail in the next section. But it is worth noting that, in general, the relationship between a disposition and its categorical basis is not *supposed* to hold with metaphysical necessity.

Consider the fragility of the wine glass. The wine glass is fragile in that it is disposed to break in relatively undemanding circumstances. But this is so not only due to the physico-chemical constitution of the glass, but also in part due to the actual force of gravity on Earth. If gravity was a thousand times weaker, the wine glass would be a thousand times less likely to break in any given circumstance, and so would be a thousand times less fragile, if you will. That is, the glass would not *be* fragile — in that its disposition to break would be very limited. Yet the physico-chemical constitution that is the disposition's categorical basis in the actual world would remain the same. Likewise, in a possible world in which the laws of psychology are radically different from those of the actual world, mental states with subjective character might well not display the disposition to be freely used in reasoning and action control. That is, where the laws of nature are sufficiently different, the categorical basis of access consciousness could certainly occur in the absence of access consciousness. (At the same time, we must maintain a nomologically necessary relation between categorical bases and the dispositions for which they are bases. Thus, the relevant microphysical structure brings about fragility in all possible worlds in which the same laws of nature hold. This is necessary to license scientific inference from dispositions (and their manifestations) to the underlying categorical properties.)

Conversely, some objects are fragile that have a physico-chemical constitution very different from the wine glass'. A vase of completely different constitution can be equally fragile. Thus similar dispositions can have dissimilar categorical bases. Moreover, the functional role occupied by the categorical basis is, like other functional roles, multiply realizable: it allows different occupants to play the exact same role. In similar fashion, access consciousness could readily occur in the absence of its actual categorical basis — if some other categorical properties served as its basis.

So in summary, the fact that there are metaphysically possible circumstances in which phenomenal and access consciousness occur in the absence of one another does not tell against the thesis that a component of the former is the categorical basis of the latter. There is thus an intimate conceptual connection between the two even if it is not a metaphysically necessary connection.

Block may present another objection to the view defended here, namely, that subjective character, or for-me-ness, is more appropriately considered an element of self-consciousness, not phenomenal consciousness. Self-consciousness, Block writes, involves "the possession of the concept of the self and the ability to use this concept in thinking about oneself." [Block, 1995, 389] This, again, is a more cog-

nitive and less phenomenal notion, so self-consciousness should be distinguished from phenomenal consciousness, as Block [1995, 389-390] indeed does. According to this objection, then, although subjective character is commonly to be found in phenomenally conscious states, it is not a constituent, but rather a contaminant, of phenomenal consciousness.

In the previous section, I drew a distinction between transitive and intransitive self-consciousness — between being self-conscious *of* an experience and self-consciously experiencing. While I acquiesce in the need to distinguish phenomenal consciousness from *transitive* self-consciousness, I have argued at length elsewhere (mainly [Kriegel, 2004]) that no consciousness can occur in the absence of *intransitive* self-consciousness. It is impossible to experience something consciously without experiencing it self-consciously. This is, as I admitted at the opening of the last section, not an uncontroversial claim. But as I also underlined, it should not affect the argument of the present paper, because Block himself accepts that there is an element of what he calls "me-ishness" in phenomenally conscious experiences. So the *existence* of this phenomenon, and even its typical *presence* in phenomenally conscious states, is not something Block wishes to deny. What Block does wish to deny is that the this somehow vindicates the scientific community's focus on it.

In other words, Block would deny my claim that this me-ishness, of for-me-ness, is the categorical basis of access consciousness. More specifically, he explicitly denies that the existence of such me-ishness somehow vindicates scientific practice. He writes [Block, 1995, 390]:

> P-conscious [i.e., phenomenally conscious] states often seem to have a "me-ishness" about them, the phenomenal content often represents the state as a state of me. But this fact does not at all suggest that we can reduce P-consciousness to self-consciousness, since such "me-ishness" is the same in states whose P-conscious content is different. For example, the experience as of red is the same as the experience as of green in self-orientation, but the two states are different in phenomenal feel.

Assuming (plausibly) that Block's notion of me-ishness is more or less the same as my notion of subjective character, or for-me-ness, we may interpret his objection as follows. A bluish experience of the sky and a whitish experience of a wall have different phenomenal characters: the one is bluish while the other is whitish. Yet the for-me-ness involved in both experiences is the same. That is, the qualitative character of conscious experiences varies independently of their subjective character. In fact, it may well be that the subjective character of conscious states is a standing feature of them that is always the same — even though conscious states vary widely with respect to their phenomenal character. Therefore, research into consciousness that focused entirely on subjective character, to the exclusion of qualitative character, would miss out on the central component of phenomenal consciousness, the component that accounts for phenomenal differences among different conscious states.

There are two rejoinders to this objection one could explore. The one I will *not* consider is that subjective character does vary, in some subtle and hardly noticeable way, in phenomenally different experiences. The rejoinder I will pursue may appear initially more surprising: that from the fact that subjective character remains unvaried across phenomenal experiences it would not follow that to focus on it would be to miss out on something essential to phenomenal consciousness. That is, it is a fallacy to infer that since a component of phenomenal consciousness is the same in phenomenally different experiences, the study of that component is insufficient for the understanding of phenomenal consciousness.

To see why this is so, consider the possibility that a conscious state's qualitative character is what makes it have the *specific* phenomenal character it has, but its having subjective character is what makes it have phenomenal character *at all*. That is, qualitative character is what makes a conscious state the conscious state it is (rather than a different conscious state), but it is its subjective character that makes it a conscious state *at all* (rather than a non-conscious state). Thus, the bluishness of my experience determines *what* it is like for me to have my experience, but it is its for-me-ness that guarantees that there is *anything* it is like for me to have it.

This possibility is fully consistent with the notion that subjective character remains the same in all phenomenally conscious states, while qualitative character varies in phenomenally different states.

There are several ways the view could be couched in more technical or theoretical terms. Thus, we may appeal to a distinction between *determinates* and *determinables*, and say that qualitative character is a determinate of phenomenal consciousness where subjective character is the determinable. Or we may draw a distinction between *existence conditions* and *identity conditions*, and claim that while qualitative character determines the identity condition of a phenomenal experience, subjective character is what determines its existence condition. Whatever the jargon, the substantial claim is that subjective character ensures that there is *something* it is like for the subject to have her conscious experiences, whereas qualitative character determines which particular *way* it is like for her to have her experience.

Now, the "hard problem" of consciousness is not so much the problem of why some phenomenally conscious experiences differ phenomenally the way they do, as the problem of why some mental states are phenomenally conscious *in the first place*. To bridge the "explanatory gap," we must come to an understanding not of how come phenomenal consciousness varies the way it does, but of how come there *is* such a thing as phenomenal consciousness *at all*. So in order to solve the hard problem, or bridge the explanatory gap, cognitive scientists ought to target the determinable, or the existence condition, of phenomenal consciousness. They must target what makes a mental state phenomenally conscious *in the first place*, not what makes it the specific phenomenal experience it is. If that is subjective character, then scientists are fully justified in focusing on subjective character. Therefore it is fallacious to infer that subjective character cannot be the proper

focus of scientific research into phenomenal consciousness if it remains unvaried across phenomenally different experiences.

4 CASE STUDIES FROM THE COGNITIVE SCIENCES

In this section, I will examine a couple of case studies from consciousness research in the cognitive sciences and analyze them in the light of the framework set forth in the last section regarding the interrelations among phenomenal consciousness, access consciousness, and subjective character.

The cognitive sciences are several in number, but they fall, crudely, into two main groups: the branches that study the psychological level and the branches that study the neurophysiological level. In sub-sections 4.1 and 4.2, I will examine key phenomena studied in each under the heading of consciousness studies, and argue that the study of such phenomena cannot be taken to shed light on the nature of qualitative character, but only on the nature of subjective character, or intransitive self-consciousness. This is, in fact, a dual claim. The negative claim is:

(NC) In studying what they do, cognitive scientists are *not* targeting the phenomenon of qualitative character.

And the positive claim is:

(PC) In studying what they do, cognitive scientists *are* targeting the phenomenon of subjective character (intransitive self-consciousness).

If so, it is presumably the latter phenomenon that cognitive scientists have in mind when they set out to study consciousness.

In a sense, then, what I will argue for is in line with Block's view on current scientific research: Block is right that cognitive scientists do not focus on qualitative character or the component of phenomenal consciousness that accounts for phenomenal *differences* among different conscious experiences. But in light of the argument of the previous section, I will not take this to constitute a critique of scientific practice, but rather an elucidation of its conceptual foundations.

4.1 Subliminal and Habituated Perception

At the *psychological* level, under the heading of consciousness research scientists focus on comparing conscious and unconscious execution of the same cognitive functions. In this way, they hope to isolate the singular contribution of consciousness to the execution of different cognitive functions, thereby treating consciousness as a scientific variable. This methodology has been articulated and expounded chiefly by Bernard Baars, who calls it *contrastive phenomenology* (see [Baars, 1994]). But it is practiced, more implicitly, by many others.

Among the phenomena Baars suggests cognitive scientists focus on are the contrasts between conscious and unconscious perception, imagery, attention, memory,

and problem-solving. One paradigmatic phenomenon for contrastive phenomenology is *subliminal perception* (see [Dixon, 1971] for the *locus classicus*). In subliminal perception, a perceptual system in some modality executes the function of feature discrimination, but does so unconsciously. This function is often executed consciously, however, so by comparing and contrasting subliminal (unconscious) discrimination and conscious discrimination of the same feature, cognitive science can investigate the singular contribution of consciousness to feature discrimination.

Phenomena of subliminal visual perception have been studied for over a century now, with Sidis [1898] and Dunlap [1900] conducting the first regimented experiments. Still one of the most interesting findings in research on subliminal vision is Dunlap's discovery that subliminally perceived stimuli can have an immediate effect on conscious perceptual experience [Dunlap, 1900, 436]. Thus, Dunlap succeeded in reproducing a conscious Müller-Lyer effect using angular lines that were only subliminally perceived. Subjects who were presented with two lines of equal length accompanied by masked angular lines that could not themselves be consciously perceived reported that one of the lines appeared longer than the other (see [Merikle and Daneman, 2000; Merikle *et al.*, 2001] for recent discussion).

A similar phenomenon in the same category is the *perception of habituated stimuli*, such as the auditory perception of the noise produced by the refrigerator pump.[8] When we get habituated to the humming of the refrigerator, our perception of it stops being conscious. At least this is what cognitive scientists assume when they compare it to conscious auditory perception of the refrigerator pump in an attempt to learn about the singular contribution of consciousness to the execution of auditory perception.

Our task here is not to comment on the plausibility of this methodological approach, but to expose its more theoretical presuppositions. That is, we must consider whether the contrast between conscious perception and, say, subliminal or habituated perception of the same stimulus is a contrast between a qualitative perception and non-qualitative perception or a contrast between an intransitively self-conscious perception and a perception that is *not* intransitively self-conscious. When a subject x has a normal conscious perception of the sky, both of the following are true:

3. x perceives the sky qualitatively.

[8] Block's example is the auditory perception of a drill outside one's window as one is engrossed in a conversation and is thus inattentive to the noise of the drill. He writes [1995, 386-7; italics original]: "Suppose that you are engaged in intense conversation when suddenly at noon you realize that right outside your window, there is — and has been for some time — a pneumatic drill digging up the street. You were aware of the noise all along, one might say, but only at noon are you *consciously aware* of it. That is, you were P-conscious [phenomenally conscious] of the noise all along, but at noon you are both P-conscious *and* A-conscious [access-conscious] of it." That is, Block takes this to be a case of phenomenal consciousness in the absence of access consciousness. Clearly, this is not how cognitive scientists treat this case. Indeed, that is the basis for Block's accusation that cognitive scientists are focusing on the wrong phenomenon. But as I will argue in the text, this is precisely a case in which the scientists are focusing on subjective character rather than qualitative character.

4. x perceives the sky self-consciously.

The question we are confronted with is which one of the two becomes false when x's perception is not conscious. More precisely, we must ask which of the two can be *established* to be false with a decent degree of scientific evidence. We would then be in a position to conclude that the contrast between conscious and unconscious perception is supposed to isolate the property of perception denoted in *that* kind of report.

In the cases of subliminal and habituated perception, we must therefore assume that the following report cases of non-conscious perception:

5. $x$has a subliminal visual perception of a blue patch.

6. x has a habituated auditory perception of a refrigerator pump.

And consider whether what x's subliminal and habituated perceptions lack is qualitative character or subjective character (intransitive self-consciousness).

Consider first qualitative character. If qualitative character is what x's perceptions reported in (5) and (6) lack, then the following should be false:

7. $x$perceives the blue patch bluishly.

8. x perceives the refrigerator pump hummingly.

Where "bluishly" and "hummingly" are supposed to be special cases of "qualitatively": the cases where the relevant qualitative character is 'bluish' or 'humming'.

Are (7) and (8) clearly false, in a way that can be established with a scientifically acceptable degree of evidence? My own intuition is that they may not be false at all. It may very well be that x's perceptions here are in fact bluish and humming, but x is simply completely unaware of their being so.

This is particularly plausible with respect to habituated perception. When the refrigerator pump goes off, x immediately notices this, but more pertinently, she can also remember what the qualitative character of her (just terminated) habituated perception had been like. She can reliably tell that it was a humming sound, rather than the sound of a trumpet. She can even compare that humming to her habituated perception of her laptop's hum, as being, say, louder or more acute. And if x can remember what the qualitative character of her perception had been like, then her perception must had had a qualitative character for x to remember. To be sure, this is not *guaranteed* to be the case: the possibility that x *constructs* these memories after the fact is not precluded by the considerations adduced thus far. But I take it that the simpler, more straightforward explanation is that (typically) x remembers her actual experiences.

I have suggested that qualitative character may well be present in perceptions of habituated stimuli. But this claim is inessential to the main point I would like to make. The main point is that perception of habituated stimuli is not a clear or obvious case of absence of qualitative character, so contrasting it with conscious

perception of the same stimuli cannot be taken to shed light on qualitative character. If so, in studying the contrast between conscious and habituated perception, in an attempt to understand consciousness, scientists are not trying to understand qualitative character. This is, in effect, our negative claim with regard to the perception of habituated stimuli:

(NCH) In studying the perception of habituated stimuli, cognitive psychologists are *not* targeting the phenomenon of qualitative character.

That is, the study of habituated perception cannot be taken to shed light on the phenomenon of qualitative character.

With subliminal perception, things are less clear than with perception of habituated stimuli. We really cannot tell whether (i) the subliminal perception of blue has a bluish character of which the subject is simply unaware or (ii) the perception really does have no bluish character. But again, this already suggests that subliminal perception does not constitute clear evidence for the absence of qualitative character, and therefore its contrast with conscious perception cannot be taken to shed scientific light on the nature of qualitative character. So the following negative claim is also warranted:

(NCS) In studying subliminal perception, cognitive psychologists are *not* targeting the phenomenon of qualitative character.

Together, (NCH) and (NCS) provide strong evidence for (NC).

Consider by contrast subjective character, or intransitive self-consciousness. If this is what x's subliminal and habituated perceptions lack, then the following should be false:

9. x perceives the blue patch self-consciously.

10. x perceives the refrigerator pump self-consciously.

This seems right: x is in fact completely unaware of her visual perception of the blue patch and her auditory perception of the refrigerator pump. (Once the pump goes off, she becomes aware that she *had been* hearing it, and may self-consciously remember having heard it. But as long as the stimulus is habituated, she does not hear it self-consciously.) That is, there is no intransitive self-consciousness involved in such perceptions. The contrast between conscious and subliminal or habituated perception therefore *is* a contrast between perception with and without intransitive self-consciousness (and hence subjective character). This establishes the positive claim with respect to the study of subliminal and habituated perception:

(PCS) In studying subliminal perception, cognitive psychologists *are* targeting the phenomenon of subjective character (intransitive self-consciousness).

(PCH) In studying the perception of habituated stimuli, cognitive psychologists *are* targeting the phenomenon of subjective character (intransitive self-consciousness).

That is, the study of subliminal and habituated perception *can* be taken to shed light on the nature of subjective character. This provides strong evidence in favor of (PC). A similar analysis may well apply to the study of other phenomena of unconscious execution of cognitive functions, though I will not pursue such an analysis here. At least as far as subliminal and habituated perception are concerned, then, research into consciousness at the psychological level does not target qualitative character, but rather subjective character.

4.2 Blindsight

Let us move on, then, to the *neurophysiological* level. At this level, the phenomena studied in consciousness research are mainly neurological syndromes in which a subject performs certain cognitive functions, but is unaware of doing so. A quite exhaustive and still relevant survey can be found in [Farah, 1995]. The phenomena she covers are blindsight, neglect, extinction, covert face recognition in prosopagnosia, and covert reading in pure alexia.

The paradigmatic phenomenon in this category is no doubt *blindsight*. Blindsight is a condition caused by lesion to the higher brain. When a blindsighted patient is asked to report what she perceives, she reports that she perceives nothing; but when she is asked to *guess* what she perceives, her guesses are correct well above chance. (The *locus classicus* here is [Weiskrantz, 1986]; see also [Weiskrantz, 1997].) Cognitive scientists conclude that the blindsighted patient does perceive her surroundings, but her perceptions are non-conscious. The disorder is not primarily at the level of perceptual functioning, then, but at the level of conscious awareness. This is what makes it relevant to study of consciousness, as a means for isolating the singular contribution of consciousness to the execution of the functions the blindsighted executes unconsciously.[9]

Again, our task is to consider what concept of consciousness is operative in cognitive scientists' work when they suppose that blindsight involves loss of consciousness. That is, we must assume, with the scientist, that something like the following reports a non-conscious perception:

11. x blindsees a snow ball.

And consider whether what x's blindseeing lacks is qualitative character or intransitive self-consciousness (subjective character).

[9] In an attempt to present a case of access consciousness in the absence of phenomenal consciousness, Block [1995, 385] appeals to what he calls "super-blindsight." For what is still an unclear reason, blindsighted person do not seem capable of spontaneously prompt themselves to "guess" what they blindsee. Super-blindsight is an imaginary condition in which the blindsighted patient has no difficulty telling herself spontaneously to "guess" what she perceives, knowing full well the theory of blindsight and the fact that chances are her "guess" is correct, and so navigates her way around the world with little difficulty. Block argues that this person's perceptual state would be access-conscious but not phenomenally conscious. This seems quite right, but as I argued in §3, does not threaten the main thesis of the present paper, the thesis that phenomenal consciousness, or rather a component of it, is the categorical basis of access consciousness.

This time let us consider first intransitive self-consciousness. If this is what x's blindseeing lacks, then the following should be false:

12. x blindsees the snow ball self-consciously.

This indeed appears to be false. It is part of the very concept of blindsight that the subject is unaware (not even peripherally aware) of perceiving what she does; that is, that she does not perceive the stimulus self-consciously. This establishes our positive claim with regard to blindsight:

(PCB) In studying blindsight, cognitive neuroscientists *are* targeting the phenomenon of subjective character (intransitive self-consciousness).

That is, the study of blindsight (and perhaps similar disorders) can, and in fact should, be taken to shed light on intransitive self-consciousness, or subjective character. This constitutes further evidence in support of (PC).

What about qualitative character? If that is what blindsight lacks, then the following is false:

13. x blindsees the snow ball whitely.

But this is, again, something we cannot determine with any confidence. The problem is again that there are two *prima facie* equally plausible possibilities. It can be that (i) x perceives the snow ball in a non-qualitative manner; or it can be that (ii) x perceives the snow ball qualitatively (more specifically, whitely), but is unaware of this. (She is unaware of this, presumably, precisely because her perception is not intransitively self-conscious.)

Ordinarily, we rely on subjects' first-person reports to determine whether their mental states exhibit qualitative character. But in the case of blindsight, we cannot rely on the subject's report on whether her perceptual state is qualitative, since the subject is unable to report that she has the perceptual state in the first place. The question, then, is whether we can decide the issue on the basis of third-person findings. The answer to this question is unclear at the moment, but the little third-person evidence we have is favorable to (ii), that is, to the possibility that x perceives the snow ball qualitatively but is unaware of the qualitative character of her perception.

The evidence in question is that blindsighted patients can apparently discriminate colors (see especially [Stoerig and Cowey, 1992]; also [Weiskrantz, 1997]). This evidence is most straightforwardly taken to suggest that perceptual states involved in blindsight do exhibit color qualities. This evidence is clearly inconclusive, though. After all, it is also possible that the blindsighted is capable of detection and discrimination of wavelengths in a completely non-qualitative manner, without ever harboring qualitative states that represent genuinely *colorful* objects. Both hypotheses accommodate the evidence. My claim here is not that the third-person evidence we have of color discrimination strongly supports the notion that blindsight can be qualitative, but rather that it supports it *weakly*.

This is because the notion that blindsight is qualitative accommodates the evidence more simply or straightforwardly than the notion that blindsight involves a non-qualitative form of wavelength discrimination at least in one sense, namely, that it constitutes a lesser departure from what happens in normal conscious color perception.

What I have said thus far is supposed to suggest, if somewhat weakly, that blindsighted perception does exhibit qualitative character. But again, more importantly for our present purposes is the fact that it appears that we cannot really know with any confidence whether or not blindseeing is qualitative. This point is sufficient, in itself, to establish our negative claim with regard to blindsight:

(NCB) In studying blindsight, cognitive neuroscientists are *not* targeting the phenomenon of qualitative character.

For the study of blindsight cannot be taken to shed light on a phenomenon nobody knows is relevant to it. If we do not *know* whether blindseeing lacks qualitative character, studying blindseeing cannot augment our *knowledge* (or further our understanding) of qualitative character. Thus (NCB) provides further evidence for (NC). The same may go, *mutatis mutandis*, for the other similar disorders noted by Farah [1995] and studied under the umbrella of consciousness studies in cognitive neuroscience. But establishing that would require a more extensive survey and analysis of the phenomena in question.

(There is a possible, though not particularly central, line of objection to the argument of this subsection, one that may be worth discussion at some length. It might be argued that quite a lot of work in neuroscientific research into consciousness targets a phenomenon that is very different from intransitive self-consciousness, namely, the phenomenon of *binding*. Thus, Crick and Koch's [1990] celebrated neurobiological theory of consciousness conceives of binding as the mark of the conscious. Let me say a little about what the phenomenon of binding is, then why Crick and Koch's work poses a *prima facie* challenge to (PC), and finally why this challenge is ultimately ineffective.

The phenomenon of binding is the fact that the various aspects of a stimulus are experienced by the subject as belonging to one and the same object, even though they are processed in different parts of the brain. For instance, when x perceives a snow ball, the whiteness of the snow ball is represented in one part of the brain, whereas the roundness of it is represented elsewhere. Yet x has a unified, cohesive conscious experience of a round white snow ball. This means that x must have a mechanism by which she *binds* the information about the roundness and the information about the whiteness of the snow ball. (The currently accepted model of binding is in terms of neural synchronization. This model was originally proposed by von der Malsburg [1981]. It still has its detractors (see, e.g., [Shadlen and Movshon, 1999]), but is generally thought to be on the right tracks.)

Crick and Koch [1990] develop a theory of the mechanism in question, and then offer this as a theory of consciousness. That is, their theory of consciousness presupposes that consciousness is essentially a phenomenon of binding, since it

is a theory of binding. So for them "x consciously perceives the snow ball" is equivalent to:

14. x perceives the snow ball bound-ly.

That is, the notion of consciousness they work with takes binding to be the mark of the conscious.

Now, there is no reason to suppose that x cannot perceive a snow ball boundly without perceiving it self-consciously. If so, work on consciousness that focuses on the binding phenomenon — in particular, work by the Singer group and by the Logothetis group — cannot be taken to shed light on the phenomenon of intransitive self-consciousness. It thus undermines (PC).

This would be a serious challenge to (PC), if anyone in the neuroscientific community accepted Crick and Koch's presuppositions regarding consciousness. In their original presentation of their theory, Crick and Koch [1990] explain at great length why they think conscious states involve binding. But nowhere do they indicate why they think non-conscious ones do not. There is in fact no reason to think that when x has a subliminal perception of a snow ball, her perception is incohesive and disunified, such that the roundness and the whiteness are each represented individually, but not as belonging to the one and the same object. This problem was quickly noticed by Crick and Koch's colleagues. Early on, Singer [1994] suggested that binding may be a *necessary condition* for consciousness, but not a *sufficient condition* (see also [Revonsuo, 1999]). This was recently admitted by Crick and Koch [2003] — "We no longer think that synchronized firing [i.e., the mechanism implementing binding], such as the so-called 40 Hz oscillations, is a sufficient condition for the NCC [the neural correlate of consciousness]," they write [Crick and Koch, 2003, 123] — and is currently the consensus in neuroscientific work on binding and consciousness: all conscious states are bound, but not *only* conscious states are; some bound states are non-conscious. Work on binding is conceived of as work on a necessary condition for consciousness, not as work on consciousness *per se* (see [Engel *et al.*, 1999] for a recent survey of the field).

What exactly is missing from unconscious bound mental states, that scientists deem them unconscious? One straightforward answer is that intransitive self-consciousness is what is missing. To sustain an argument against PC on the basis of neuroscientific research into binding, the objector would have to exclude this possibility. The focus on binding itself is no argument against PC, since nobody today takes binding to be sufficient for consciousness. That is, nobody implicitly works with a notion of consciousness that sustains an equivalence between "x perceives consciously…" and "x perceives boundly… ."

It might be objected that, at the end of the day, the position I defend is only superficially different from Block's. My own distinction between subjective and qualitative character basically parallels Block's between access and phenomenal consciousness, and my own analysis of current work in consciousness studies only reinforces his claim that it targets the former instead of the latter.

The objection fails in two ways. Firstly, the notion of subjective character does not parallel that of access consciousness, inasmuch as the latter is dispositional while the former is not. Secondly, subjective character is an aspect of phenomenal consciousness, and so cannot importantly parallel a non-phenomenal notion of consciousness. This is not merely an expedient verbal decision on my part. Labels aside, what is essential to phenomenal consciousness in the present context is that it is the property which generates the mystery of consciousness. It is a substantive claim that subjective character is an aspect of the property that generates the mystery of consciousness, not an unmysterious accompaniment thereto.

5 CONCLUSION: IN DEFENSE OF SCIENTIFIC PRACTICE

This completes the main argument of this paper. I have set to approach the problem of consciousness from the angle of the philosophy of science, by treating consciousness as a scientific concept pregnant with certain philosophical presuppositions, presuppositions the exposition of which is the philosopher's main contribution to research. I have argued that when we approach consciousness from this philosophical-of-science angle, it appears that consciousness is fundamentally a matter of subjective character, not qualitative character. That is, what makes a phenomenally conscious state what it is (i.e., the existence condition of phenomenal consciousness) is the elusive and unusual awareness we have of our concurrent conscious states, not the sensuous quality typically present in them.

The claim that by studying what they do, cognitive scientists are *not* targeting the phenomenon of qualitative character, and *are* targeting the phenomenon of subjective character (construed in terms of intransitive self-consciousness), has been exemplified through a couple of key phenomena, from both psychological and neurophysiological empirical research.

Perhaps more fundamentally, however, I have argued that current scientific practice in consciousness studies is fundamentally sound. The argument may be schematized as follows: 1) Access Consciousness is a dispositional property, the categorical basis of which is subjective character, which is a component of phenomenal consciousness; 2) Science often tries to understand a categorical property by studying the dispositional properties for which it is the basis; therefore, 3) Cognitive Scientists may be trying to understand phenomenal consciousness (or at least a component thereof, namely, subjective character) by studying access consciousness.

Contrary to Block, then, cognitive scientists have not shied away from phenomenal consciousness in order to focus on a closely related but inherently different property. Rather, they have attacked the problem of phenomenal consciousness by examining a dispositional property (access consciousness) whose categorical basis is a crucial component of phenomenal consciousness, the component that *makes* mental states phenomenally conscious.

ACKNOWLEDGEMENTS

For comments on earlier drafts of this paper and/or relevant discussions, I would like to thank George Graham, Richard Healey, Paul Thagard, Cybele Tom, and especially David Chalmers. I have also benefited from presenting this paper at CREA and the University of Arizona's Cognitive Science Program. I would like to thank the audiences there, in particular Steven Biggs, Alexei Grinbaum, Keith Lehrer, Farid Masrour, Jean Petitot, and Orlin Vakarelov.

BIBLIOGRAPHY

[Armstrong, 1968] D. M. Armstrong. *A Materialist Theory of the Mind*. New York: Humanities Press, 1968.

[Baars, 1994] B. J. Baars. Contrastive Phenomenology: A Thoroughly Empirical Approach to Consciousness, 1994. Second version in Block *et al.*, 1997.

[Block, 1995] N. J. Block. On a Confusion About the Function of Consciousness. *Behavioral and Brain Sciences* 18: 227-247, 1995. Reprinted in Block *et al.*, 1997.

[Block, 1997] N. J. Block, O. Flanagan, and G. Guzeldere, eds. *The Nature of Consciousness: Philosophical Debates*. Cambridge MA: MIT, 1997.

[Carruthers, 1989] P. Carruthers. Brute Experience. *Journal of Philosophy* 85: 258-269, 1989.

[Carruthers, 2000] P. Carruthers. *Phenomenal Consciousness*. Cambridge: Cambridge UP, 2000.

[Chalmers, 1995] D. J. Chalmers. Facing Up to the Problem of Consciousness. *Journal of Consciousness Studies* 2: 200-219, 1995.

[Chalmers, 1997] D. J. Chalmers. Availability: The Cognitive Basis of Consciousness? *Behavioral and Brain Sciences* 20: 148-149, 1997. Reprinted in Block *et al.*, 1997.

[Clark, 2000] A. Clark. A Case Where Access Implies Qualia? *Analysis* 60: 30-38, 2000.

[Crick and Koch, 1997] F. Crick and C. Koch. Towards a Neurobiological Theory of Consciousness. *Seminars in the Neurosciences* 2: 263-275, 197. Reprinted in Block *et al.*, 1997.

[Crick and Koch, 2003] F. Crick and C. Koch. A Framework for Consciousness. *Nature Neuroscience* 6: 119-126, 2003.

[Dennett, 1995] D. C. Dennett. The Path Not Taken. *Behavioral and Brain Sciences* 18: 252-253, 1995. Reprinted in Block *et al.*, 1997.

[Dixon, 1971] N. F. Dixon. *Subliminal Perception: The Nature of a Controversy*. London: McGraw-Hill, 1971.

[Dretske, 1995] F. I. Dretske. *Naturalizing the Mind*. Cambridge MA: MIT Press, 1995.

[Dunlap, 1900] K. Dunlap. The Effect of Imperceptible Shadows on the Judgment of Distance. *Psychological Review* 7: 435-453, 1900.

[Engel *et al.*, 1999] A. K. Engel, P. Fries, P. Konig, M. Brecht, and W. Singer. Temporal Binding, Binocular Rivalry, and Consciousness. *Consciousness and Cognition* 8: 128-151, 1999.

[Farrah, 1995] M. Farah. Visual Perception and Visual Awareness after Brain Damage: A Tutorial Overview. In C. Umilta and M. Moscovitch (eds.), *Attention and Performance XV: Conscious and Nonconscious Information Processing*, 1995. Reprinted in Block *et al.*, 1997.

[Fodor, 1968] J. A. Fodor. *Psychological Explanation*. New York: Random House, 1968.

[Fodor, 1975] J. A. Fodor. *The Language of Thought*. New York: Thomas Y. Crowell, 1975.

[Fodor, 1983] J. A. Fodor. *The Modularity of Mind*. Cambridge MA: MIT Press, 1983.

[Huntington ..., 1993] Huntington's Disease Collaborative Research Group. A novel gene containing a trinucleotide repeat that is expanded and unstable on Huntington's disease chromosomes. *Cell* 72: 971-983, 1993.

[Kriegel, 2003a] U. Kriegel. Consciousness, Higher-Order Content, and the Individuation of Vehicles. *Synthese* 134: 477-504, 2003.

[Kriegel, 2003b] U. Kriegel. Consciousness as Intransitive Self-Consciousness: Two Views and an Argument. *Canadian Journal of Philosophy* 33: 103-132, 2003.

[Kriegel, 2004] U. Kriegel. Consciousness and Self-Consciousness. *The Monist* 87: 185-209, 2004.

[Kriegel, 2006] U. Kriegel. The Same-Order Monitoring Theory of Consciousness. In U. Kriegel and K. Williford (eds.), *Consciousness and Self-Reference*. Cambridge, MA: MIT Press, 2006.
[Kriegel, forthcoming] U. Kriegel. Naturalizing Subjective Character. *Philosophy and Phenomenological Research*, forthcoming.
[Levine, 2001] J. Levine. *Purple Haze: The Puzzle of Consciousness*. Oxford and New York: Oxford UP, 2001.
[Libet, 1985] B. Libet. Unconscious Cerebral Initiative and the Role of Conscious Will in Voluntary Action. *Behavioral and Brain Sciences* 8: 529-566, 1985.
[Lycan, 1990] W. G. Lycan. Consciousness as Internal Monitoring. *Philosophical Perspectives* 9: 1-14, 1990. Reprinted in Block *et al.*, 1997.
[Malsburg von der, 1981] C. von der Malsburg. The Correlation Theory of Brain Function. Technical Report 81-2, Max-Planck-Institute for Biophysical Chemistry, Gottingen, 1981.
[Merikle and Daneman, 2000] P. M. Merikle and M. Daneman. Conscious Vs. Unconscious Perception. In M. S. Gazzaniga (ed.), *The New Cognitive Neurosciences* (2^{nd} edition). Cambridge MA: MIT Press, 2000.
[Merikle *et al.*, 2001] P. M. Merikle, D. Smilek, and J. D. Eastwood. Perception without Awareness: Perspectives from Cognitive Psychology. *Cognition* 79: 115-134, 2001.
[Nagel, 1974] T. Nagel. What Is It Like to Be a Bat?" *Philosophical Review* 83: 435-450, 1974. Reprinted in Block *et al.*, 1997.
[Revonsuo, 1999] A. Revonsuo. Binding and the Phenomenal Unity of Consciousness. *Consciousness and Cognition* 8: 173-185, 1999.
[Rosenthal, 1991] D. M. Rosenthal. The Independence of Consciousness and Sensory Quality. *Philosophical Issues* 1: 15-36, 1991.
[Rosenthal, 1997] D. M. Rosenthal. A Theory of Consciousness, 1997. In Block *et al.*, 1997.
[Sidis, 1898] B. Sidis. *The Psychology of Suggestion*. New York: D. Appleton and Company, 1898.
[Shadlen and Movshon, 1999] M. N. Shadlen and J. A. Movshon. Synchrony Unbound: A Critical Evaluation of the Temporal Binding Hypothesis. *Neuron* 24: 67-77, 1999.
[Shoemaker, 1979] S. Shoemaker. Identity, Properties, and Causality. *Midwest Studies in Philosophy* 4: 321-342, 1979.
[Singer, 1994] W. Singer. The organization of sensory motor representations in the neocortex: A hypothesis based on temporal coding. In C. Umilta and M. Moscovitch (eds.), *Attention and Performance XV: Conscious and Nonconscious Information Processing*. Cambridge MA: MIT Press, 1994.
[Stoerig and Cowey, 1992] P. Stoerig and A. Cowey. Wavelength sensitivity in blindsight. *Brain* 115: 425-444, 1992.
[Strawson, 1994] G. Strawson. *Mental Reality*. Cambridge MA: MIT Press, 1994.
[Tye, 1995] M. Tye. *Ten Problems of Consciousness*. Cambridge MA: MIT Press, 1995.
[Tye, 2000] M. Tye. *Consciousness, Color, and Content*. Cambridge MA: MIT Press, 2000.
[Velmans, 1992] M. Velmans. Is Human Information Processing Conscious? *Behavioral and Brain Sciences* 14: 651-669, 1992.
[Weiskrantz, 1986] L. Weiskrantz. *Blindsight: A Case Study and Implication*. Oxford: Oxford UP, 1986.
[Weiskrantz, 1987] L. Weiskrantz. *Consciousness Lost and Found*. Oxford and New York: Oxford UP, 1987.

ON RESTRICTING THE EVIDENCE BASE FOR LINGUISTICS

C. Iten, R. Stainton, and C. Wearing

1 INTRODUCTION

This chapter will investigate two questions about the appropriate evidence base for linguistics. The first question is: *should the evidence base for linguistics be restricted a priori and in principle in any fashion?* This isn't the question of whether one should make reasonable guesses about where to look for evidence. Of course one should. The first question, rather, is whether we can determine in advance of inquiry, from the 'armchair' as it were, what evidence could not possibly play a part. The second question is: *should evidence from the cognitive sciences be excluded?* In particular, we will be interested in the status of evidence from mentalistic psychology and neuroscience. This second question is in part a special case of our first question: should the evidence base for linguistics be restricted *a priori* and in principle to exclude evidence from mentalistic psychology and neuroscience? At the same time, it includes the more practical question of whether, even if admitted in principle, there is good reason to think that such evidence will not actually play a part in linguistics. (An analogy may clarify the difference between the 'in principle' and the 'in practice' questions. In the novel *Dirk Gently's Holistic Detective Agency*, the eponymous detective insists that everything is connected, so that evidence about, say, where a missing cat has gone could come from anywhere. Thus Dirk would reject any in principle restriction on the evidence base for his detections. He adds, however, that because of the fundamental interconnectedness of all things, he begins each investigation with a client-financed fortnight on the beach in Bermuda — even if the investigation is about a lost cat in Surrey. Here he (implausibly) suggests that *in practice* all evidence is on equal footing.)

We will pursue these two questions by considering three *pro-restriction* answers to them — answers that favour restricting the evidence base for linguistic enquiry, specifically setting aside neuropsychological evidence. In particular, we focus on positions inspired by Jerrold Katz, W.V.O. Quine and Ludwig Wittgenstein. Our goal in discussing these negative answers will be to clarify two things: first, the views of language (and language study) underlying each answer, and second, how each of these views leads to the restriction of the evidence base. Our aim is not exegetical. Instead, we want to make clear to linguists who *are* prone to use

Handbook of the Philosophy of Science. Philosophy of Psychology and Cognitive Science
Volume editor: Paul Thagard
General editors: Dov M. Gabbay, Paul Thagard and John Woods
© 2007 Elsevier B.V. All rights reserved

evidence from the cognitive sciences why anyone would favour restrictions. As will emerge, restricting the evidence base may ultimately be a mistake; but we want to make clear why it isn't an obvious mistake. That's the burden of §§2-4. In §5, we will respond to the pro-restriction arguments, showing why they are less compelling than they initially appear. Given the problems with restricting the evidence base, we will conclude that such restrictions are not ultimately justified. We end with one quite unexpected way in which evidence from mentalistic psychology and neuroscience has recently illuminated English morphology, syntax and compositional semantics. The overall conclusion is that, though the alternative can be tempting, linguistics should proceed in the same way as other scientific endeavours, namely, by allowing that evidence from any quarter, including specifically from the domains of psychology and cognitive science, might in principle prove to be illuminating. What's more, though it may initially seem that neuropsychological evidence won't turn out in practice to be of great use, this impression is ultimately mistaken as well.

2 JERROLD KATZ: AN ONTOLOGICAL ARGUMENT FOR RESTRICTION

As a preliminary to our discussion of Katz, consider that there are at least three kinds of objects: concrete physical objects, mental objects, and abstract objects. Some examples of concrete objects include snow, cats, buses, the Empire State Building, and the black marks on this page. Mental objects include such things as pains, beliefs, memories, hopes, and dreams. By contrast, each of Beethoven's 5^{th} *Symphony*, Tolstoy's *War and Peace*, the American Constitution, and the number 8 is an abstract object — none is identical with any particular physical instantiation of it (a specific performance of certain sounds, a specific copy of the book, or piece of parchment, or squiggle on a piece of paper).[1]

One set of reasons for restricting the evidence base when studying a thing X arises from certain conceptions of what sort of object X is. In particular, if X is *not* the kind of object that certain types of evidence could (in principle) tell us anything about, then there is no need to bring those types of evidence to bear on the study of X. More than that, to do so just generates confusion. A specific example may help. Simplifying for expository purposes, Gottlob Frege [1884] insisted, *contra* a dominant attitude of the day, that numbers, numerical facts, and mathematical systems were not mental entities. That the square root of 698 is larger than 25 is a fact about numbers, not about human psychology. And no psychological experiments could refute a mathematical proof of this fact. Indeed, Frege's arguments have convinced most contemporary philosophers that numbers are acausal, atemporal entities. He further urged that, in light of these

[1]This does not exclude the possibility that something may belong to more than one of these categories. In particular, mental objects such as pains may also be physical objects (e.g., brain states). Nor do we wish to exclude concrete properties, mental properties, etc. Given present purposes, however, such metaphysical niceties need not detain us.

ontological considerations, only confusion was generated by reflecting on mental things when doing mathematics. Thus he insisted that mathematics be insulated, both in principle and in practice, from work in psychology: "... psychology should not suppose that it can contribute anything at all to the foundation of arithmetic. To the mathematician as such, these mental images, their origin and change are irrelevant" (87).

For Frege's line of thought to be plausible, it's crucial to make a distinction. We must distinguish between studying something and studying our psychological access to that thing. For example, studying the weather involves gathering information about such physical phenomena as winds, temperatures, and tides. But studying our psychological grasp of weather-facts involves looking at us — specifically, at how we acquire information about the weather, as well as how we store that information and how we put it to use. The latter is an investigation into our psychology, but the former is not. The same distinction holds for abstract objects: studying abstract objects such as numbers is a different enterprise from studying our psychological ability to grasp those objects. Studying numbers, which is what mathematicians do, involves discovering and proving mathematical theorems. Investigating how humans develop and use their mathematical knowledge (e.g., how children learn numbers), by contrast, is a job for psychologists. So, studying numbers and studying our knowledge of numbers are different kinds of enquiry. Moreover, if numbers are in fact abstract objects, then it looks as though the sort of enquiry that psychologists undertake into our knowledge of mathematics is not the right sort of enquiry to yield mathematical truths about numbers. In other words, psychological evidence about our knowledge of mathematics, though of terrific interest in itself, is not relevant to studying mathematics proper. (Knowing, for example, how children learn to add does not help us prove that 4 plus 3 equals 7.)

Now consider languages like English and Swahili. Suppose that such languages are also abstract objects, just as the number 8 is. One might infer, following Frege's lead, that psychology should have no place in the study of language either. This is precisely Katz's take on things. On this sort of view, the restriction of the evidence base follows from one's conception of what language is — the restriction derives from *ontological* considerations about the nature of language. That's the first sort of argument we will consider for restricting the evidence base in linguistics.[2]

The first step in the argument to the effect that languages are the wrong kind of thing to be illuminated by psychology is to show that languages are not concrete items. That is, they are not simple physical objects, consisting of non-mentally-specified physical marks and sounds, paired with the objective concrete objects which trigger the production of those sounds. Such a wholly concrete view of

[2]David Lewis [1975] and Scott Soames [1984; 1985] defend views that are similar to Katz's both in their treatment of languages as abstract formal objects and in drawing consequences from this for the evidence base appropriate for linguistics. As will emerge below, however, not everyone who takes the point that languages are abstract infers thereupon that psychological evidence is irrelevant to linguistics. For intermediate positions and discussions of them, see Barber [2003], Devitt [2003], George [1989] and Higginbotham [1983].

language was purportedly held by structuralist linguists such as Bloomfield before the 1950s; however, it was shown by Chomsky to be radically inadequate. It would take us too far afield to rehearse the key points; and we are especially uninterested in exegesis here. But notice, for example, that things which count as the same 'linguistic sound' can be physically very different. Thus, in (one dialect of) Spanish, "bee" and "vee" are different ways of pronouncing the very same linguistic sound. To take another example, think of the different acoustics of the word "vitamin" when said by a small child, a woman, an elderly man, someone with a bad cold, someone yelling, a tracheotomy patient, and so forth. Here again, what counts linguistically as 'the same sound' is acoustically diverse. These examples show that even the most 'concrete' aspect of language isn't entirely concrete or physical: what holds them all together as 'the same sound' is something more abstract. The same holds for marks. Consider all the different ways of writing a word: printed by hand (in countless different handwritings), drawn in the sand, written in cursive, typed in bold, italics, all-caps, etc. In fact, what all these ways of writing have in common is something about how language users process them. Similarly for the soundings of "bee" and "vee", or "vitamin". So treating language as a concrete physical object precludes the possibility of making the right generalizations by preventing abstraction away from the physical variation exhibited across different instances. Equally problematic for taking a language to be a simple concrete, non-mental thing, is that the same acoustic item can realize different 'linguistic sounds'. Thus we English speakers 'hear different sounds' in the middle of "writer" and "rider". But in fact, acoustically, the middle sound is the same: it's just that the surrounding material causes us to process this acoustic item differently in the two cases.

As noted, the first step in the argument is to show that languages are not concreta. Just about everyone who works in linguistics has granted this by now. More plausibly, languages might be something in the mind — e.g., whatever it is in virtue of which someone is able to produce the marks and sounds that we recognize as linguistic. This would capture the fact that 'linguistic sounds' aren't really acoustically individuated. It would also capture the absolutely essential fact that any given human language contains many, many more sentences than ever have been, or ever will be, produced. But, to come to the second step in Katz's argument, he has urged that mental items are still too concrete to be the sort of thing that a language is. Says Katz, words and sentences, like numbers, are atemporal and acausal entities. They are types, not tokens. As such, he adds, they are ontologically independent both of concrete matter and of human minds. The implications for our questions should be clear: if language is a mental object, then evidence from mentalistic psychology and neuroscience will be important to its study, for such evidence aims to reveal how the mind works; but if language is an abstract object, then it does not follow in the same immediate way that information about the minds or brains of language speakers tells us anything about that abstract object. Psychology might tell us how we come to know and use languages, but it doesn't seem able to tell us about language itself.

Why should one regard language in this way? What arguments support this view of languages as abstract objects not constrained by psychological considerations? Katz views Platonism as a largely neglected alternative to the mentalistic approach, an alternative which is not touched by Chomsky's arguments against the structuralist linguistics of Bloomfield and others. Those arguments clearly demonstrated the explanatory advantages of a mentalistic view of language over the view that language consists merely in a set of physical marks and sounds, but they were not simultaneously aimed at defending such a mentalistic view of language from a Platonist alternative, for no one really considered such a possibility. As a result, Katz argues, it remains an option until proven otherwise.

Taking off from the perspective of Platonism as a neglected alternative, he pursues two lines of argument. First, Katz argues that the empirical commitments of the mentalistic view of language get in the way of an adequate description and explanation of grammar: "there is a possibility of conflict between a demand that grammars satisfy an extrinsic, ideologically inspired constraint and the traditional demand that grammars meet intrinsic constraints concerning the successful description and explanation of the grammatical structure" [Katz, 1985, 196]. One such conflict, although offered by Katz as a mere possibility, arises if the 'psychologically preferable theory' (grammar) is inferior to others on methodological grounds. For example, if the way of representing grammar that English speakers seem to use proves to be less elegant than other possibilities, then the mentalistic approach must compromise such standards as simplicity. Faced with such a possibility, Katz suggests that "abandoning [this approach] is preferable to committing ourselves to such methodologically perverse choices" [1985, 198].

Second, Katz offers a Frege-inspired argument for the abstractness of language, and hence (allegedly) for the irrelevance of psychological evidence. We typically take truths about numbers to be necessary — they could not be otherwise — and psychology doesn't look capable of yielding truths of this sort. Psychology can only tell us about the limitations on our abilities to conceive of numbers as having certain properties, but it can't tell us about how numbers behave independently of how we are able to think about them. So, an explanation of human mathematical capability is incapable of explaining the necessity of mathematical truths. Empirical investigation just can't uncover such deep and abiding necessities. Katz extends this problem to the case of language by focusing attention on the grammatical (semantic) property of *analyticity*, which is exhibited by such sentences as:

- Nightmares are dreams, and
- Flawed gems are imperfect.

(These examples are from Katz [1985, 199]. For more extended discussion, see Katz [1981, Chapter 5].) Katz takes these sentences to be necessary — they are true come what may. Moreover, he takes their truth to depend on the combination of their meanings and their syntax. As a result of this dependency, their

analyticity must be explained by the grammar of English, for it depends solely on grammatical properties. But, as a result of their necessity, the analyticity of these sentences cannot follow from a psychological explanation of the grammar of English — or so Katz insists — for psychology can only deliver "what human beings are psychologically or biologically forced to conceive to be true no matter what. But this is a far cry from what *is* true no matter what" [Katz, 1985, 200]. So, Katz effectively presents a dilemma: either one seeks a grammar that explains the necessity of the truth of these sentences, or a grammar that is psychological. One cannot accomplish both tasks at once, for the analyticity of these sentences is only intelligible if the languages are regarded as abstract entities, rather than mental entities.

Let's sum up so far. Katz argues on two different grounds that languages are abstract objects, on a par with numbers or algebras: first, if we don't treat languages as abstract, we risk not meeting certain intrinsic demands on a theory of language; second, if we don't treat languages as abstract, we won't account for certain necessities. He also urges that if languages are abstract, then neuropsychological evidence cannot bear on their nature. Thus his negative answer to both of our questions: the evidence base for linguistics should be restricted *a priori* and in principle; in particular, evidence from mentalistic psychology and neuroscience should be excluded not only in practice, but in principle. The key rationale for this answer is that the nature of an abstract system is one thing, the nature of our grasp of it is quite another. Psychology can of course tell us about the latter. But it simply cannot be relevant to the former. (See [George, 1989] and [Soames, 1984; 1985] for related points. Devitt [2003], in contrast, urges that the move from ontological status to evidence-restriction is unwarranted.) We turn now to a second, quite different, way of arriving at similar conclusions.

3 QUINE: A RADICAL EMPIRICIST ARGUMENT FOR RESTRICTION

Quine's reasons for restricting the evidence base are, unlike Katz's, more learning-theoretic than ontological. It is not because he views languages as abstract objects that he restricts the evidence base to "what is to be gleaned from overt behaviour in observable circumstances" [Quine, 1987, 5]. Instead, an anti-nativist commitment underwrites Quine's view of the evidence relevant to linguistic enquiry. Or so it seems to us. (Again, we are not interested in exegesis. The aim, to repeat, is to review initially attractive arguments for restricting the evidence base for linguistics.) To foreshadow the second line of argument as a whole: we read Quine as supporting anti-nativism on the basis of the success of the natural sciences; anti-nativism then has implications for what fixes the nature of the language spoken around one; and this, finally, entails that certain kinds of evidence should be avoided.

The starting point for understanding Quine's anti-nativism is his strong commitment to empiricism, i.e. to the view that our knowledge comes from experience. Insofar as our knowledge of language is simply one branch of our total knowledge,

it too must derive from our experiences. One might wonder why this empiricist commitment constitutes a good starting point — why is it plausible that our knowledge comes from experience? One consideration, which constitutes a point in its favour if not an answer, is the fantastic success of the sciences that take empiricism as their foundation. The domains of physics, chemistry, and so forth, which take themselves to be based on and responsible to what we can observe (even indirectly) about the world, achieve ever greater predictive and explanatory power. Clearly, this sort of instrumentalist reasoning in favour of empiricism ('it works, so it's right') doesn't immediately establish its correctness, but it does suggest that the empiricism is a fruitful methodology to adopt.

Now for Quine, what goes for the scientific community goes for the individual as well. (See, for example, [Quine, 1969b].) Both start in the same place, and both use the same tools of investigation and confirmation. That's one key idea. The other is that empiricism doesn't, for Quine, just say that our knowledge requires experience. It says that the *only* evidential source for our knowledge is our observations of the world; we do not have any innately endowed ideas. This is not to insist that we don't make any contribution at all. Like earlier radical empiricists such as Hume, Quine allows for innate general-purpose mechanisms that operate on experience: innate quality spaces, innate mechanisms for forming associations, etc. What he and other radical empiricists reject is the claim that humans have innate information, evidence, beliefs, contents, reasons, ideas, etc. The only information from which knowledge derives, whether for the scientific community or for the infant, is information that comes from experience. In short, we see a move from views about the nature of scientific inquiry, and its success, to anti-nativism about individuals.

Given this general Quinean view about learning, let us consider the situation of the child who sets out to acquire a language. (This will support Quine's behaviorism.) The child, Quine holds, is in the same position as the empirical scientist: she must figure out how language works solely on the basis of the evidence presented to her. As he puts it, "each of us learns his language by observing other people's verbal behaviour and having his own faltering verbal behaviour observed and reinforced or corrected by others" [Quine, 1987, 5]. In this way, Quine takes the child to learn her language strictly and solely on the basis of the evidence presented to her (the sentences she hears others speak) and the corrections of her own attempts to utter appropriate sentences in one situation or another:

> Language is a social art which we all acquire on the evidence solely of other people's overt behavior under publicly recognizable circumstances. Meanings, therefore, those very models of mental entities, end up as grist for the behaviorist's mill [1969a, 26].

> ... even in the complex and obscure parts of language learning, the learner has no data to work with but the overt behavior or other speakers [1969a, 28].

The learner could not, then, be required to use information about what is going on in the minds or brains of other language users, because she has no access to this information. More importantly, she has no *need* of this sort of information, for all she needs to learn is which utterances it is appropriate to produce in response to which stimuli — and how exactly other speakers do this is not something she needs to know. No other evidence, such as the evidence of mentalistic psychology or the other cognitive sciences, could be relevant here, for there are no additional facts for this evidence to uncover. More than that, looking at purported evidence of this kind can only lead to confusion and error, since it may (wrongly) suggest to the investigator that there are more facts than there genuinely are.

In sum, reflecting on the learning situation of the infant, Quine infers that the situation is (and must be) the same for the practicing linguist who aims to discover what we know when we know a language: "in psychology, one may or may not be a behaviourist, but in linguistics one has no choice" [Quine, 1987, 5]. Just as behavioural evidence is sufficient for the child, it must be sufficient for the linguist, because all that is required of our linguistic knowledge is that it "fits all external checkpoints", i.e. conforms to the regularities of stimulus-response pairings in the language community. As Davidson puts it, "The semantic features of language are public features. What no one can, in the nature of the case, figure out from the totality of the relevant evidence cannot be part of meaning" [1979, 235]. Thus we again have negative answers to our two questions: if Quine is right then we know already, prior to any empirical investigation, that the evidence base for linguistics should be restricted in principle; in particular, we know that neuropsychological evidence should be excluded not just in practice but in principle as well. (Note: It is well known that Quine is a behaviorist. But, it's worth stressing, his conclusion here goes far beyond the restriction of evidence to observable behavior: the only kind of behavior the linguist can make use of, if Quine's argument works, is behavior that an ordinary infant has access to. This excludes behavioral evidence from clinical cases, comparison with other languages, etc., as well as other physical evidence from autopsies, brain scans, and so on.)

4 WITTGENSTEIN: A 'NORMATIVE COMMUNITY PRACTICE' ARGUMENT FOR RESTRICTION

We have presented two views on language each of which leads to restrictions in the evidence base for linguistics. We have not tried to show that these views are ultimately correct, nor have we been much concerned with the question of whether the views discussed fit perfectly with the names we have associated with them. Our aim, rather, is to explain why imposing restrictions is neither perverse nor merely contrarian. We turn now to a third example. It is widely agreed that Ludwig Wittgenstein [1953] suggests that language is a kind of social practice. The view has been defended more recently by Michael Dummett. In this view of language we can discern three central features: first, language is community-based; second, it is a practice; and third, it has an irreducibly normative dimension. Weaving

together these central features, the Wittgensteinian argues that linguistic meaning is constituted by *use*: there is nothing to a language beyond the way it is used in a given culture. In defending this view of language, the Wittgensteinian rejects the idea that words have meanings *beyond* how they are collectively used — human-independent meanings that are capable of autonomously adjudicating whether a particular use of those words is correct or incorrect. All that we have are the norm-governed *uses* to which words are put *by us*. (To be clear, by "use" Wittgenstein *et al.* do not mean the kind of thing that a scientistic behaviorist might: e.g., bodily motions of various types. They have in mind culturally constructed actions, norm-bound moves in a collective game. This is already one sense in which Wittgensteinians differs markedly from Quine.) Thus Michael Dummett writes:

> The meaning of [any][3] statement determines and is exhaustively determined by its *use*. The meaning of such a statement cannot be, or contain as an ingredient, anything which is not manifest in the use made of it, lying solely in the mind of the individual who apprehends that meaning: if two individuals agree completely about the use to be made of the statement, then they agree about its meaning. The reason is that the meaning of a statement consists solely in its rôle as an instrument of communication between individuals, just as the powers of a chess-piece consist solely in its rôle in the game according to the rules [1973, 216].

As we shall see, this view of language as *nothing more than* a normative community practice leads Wittgenstein *et al.* to abjure psychological evidence in studying language. (Or better, one can construct on the basis of certain of Wittgenstein's aphoristic remarks some *prima facie* good reasons for doing this. An old joke has it that, if you put three learned rabbis in a room, you'll get four interpretations of the Talmud. Say we, if you put one philosopher in a room, you'll get four interpretations of Wittgenstein!)

Let's examine how language might be an irreducibly normative, community-based practice by drawing an extended analogy with the waltz. First, the waltz is not something that you can do alone. It requires two dancers, and perhaps an orchestra to play appropriate accompanying music. In this way, it is an activity that is only possible with the participation and coordination of multiple people. At the same time, doing the waltz consists in moving in particular conventional ways in coordination with one's partner. It is distinguished by doing certain specific steps in time with specific kinds of music. Which steps constitute the waltz is a matter of convention, which reflects a second, even deeper, sense in which 'others' are required: a convention is the property of a group. Similarly, speaking a language, as Wittgenstein famously argues, is not something one can

[3]In the paper in question, Dummett is discussing mathematical statements. Thus instead of "any statement", he actually says "a mathematical statement". But what he endorses in this paper, he takes to apply to linguistic expressions generally.

do alone – language only makes sense as something that people share in a culture. As he says, "it is not possible to obey a rule 'privately': otherwise thinking one was obeying a rule would be the same thing as obeying it" [1953, §202]. So language, like the waltz, involves, and indeed requires, a community.

Third, we can see the essential normativity of the practice of waltzing in the fact that only certain performances of the waltz count as successful. There is a 'right way' to move your feet which will have you waltzing, while if you move your feet a different way, you may find yourself doing the fox-trot or simply shuffling around the dance floor looking slightly ridiculous. So there is a standard of correctness associated with waltzing: you have to do *the right steps*. Language works the same way: if you don't use the words of a language in the 'right way', you fail to count as a participant in the language-game. Imagine a person who consistently uses English nouns in ways that differ from other English speakers, say, they call dogs "horses", and snowballs "cars", and trees "turnips", and so forth. We would judge that this person is doing something *wrong* — she is failing to do what she is supposed to in order to count as speaking English. So, like the practice of waltzing, the practice of speaking a language comes with a normative standard — only some performances count as speakings of that language.

Finally, coming to the idea that the action is exhausted by these three features, notice that nothing more is involved in doing the waltz than following the norm-governed publicly observable steps. If we want to decide whether someone is doing the waltz, we need to look at their feet to see which steps they are performing. Suppose that they are indeed performing the steps of the waltz (i.e. the same steps that everyone else performs, the steps that are taken to be the steps of the waltz). How exactly their mind and body get their feet to follow those steps, rather than moving in some other way, is *irrelevant* to whether they are doing the waltz. Each dancer could be using a different part of the brain; indeed, something without a human brain could presumably dance the waltz. (Recall the Beatles' song *Being for the Benefit of Mr. Kite*: "And of course Henry the Horse dances the waltz".) All that matters is that the dancer's feet perform the right steps. (Says Wittgenstein: "a wheel that can be turned though nothing else moves with it, is not part of the mechanism" [1953, §271]. See also the invisible beetle in the box analogy, at §293.)

In a similar way, goes the idea, speaking a language involves nothing more than using the words of that language in just those ways that other speakers of the language do. For Wittgenstein, Dummett, *et al.*, what's going on in someone's head in virtue of which they manage to participate successfully in their community's language-game is irrelevant to their participation. All that matters is that their actions conform to the community's linguistic norms. Put otherwise, which words are to be used when is a matter of convention, and (so Wittgensteinians argue) there is nothing more to speaking the language than coordinating one's behaviour with that convention. What goes on 'inside' is just beside the point. (Other examples of cultural practices where what's going on inside seems irrelevant

include voting for the Democratic candidate in your riding, depositing $10.00 in your checking account, and driving over the speed limit.)

If this view of language as an irreducibly normative, community-based practice is correct, then the implications for the evidence base for linguistics seem clear: evidence from such fields as psychology and neuroscience are largely irrelevant, for there is nothing in the heads of individual speakers that is particularly relevant to understanding something whose nature is *exhausted by* being a norm-governed community practice. More precisely, there may be facts about what people are capable of learning or remembering, and facts about how their psychology affects their social coordination, that are relevant to *how* one manages to speak. But there is an important sense in which what is going on in speakers' heads is quite irrelevant to explaining *whether* they are participating in the language-game correctly. Their correct participation amounts to nothing more than their conforming successfully to the norms of the game, and we can assess their success best simply by observing their 'outside'.

So it seems to follow from the Wittgensteinian's view of language that the evidence base appropriate to linguistics is restricted to observations of "moves in the language game". Information about the psychology of speakers, while it might seem capable of illuminating how that behaviour is produced, is at the very least always trumped by, if not plainly irrelevant to, what we learn from observation of speech.

To conclude our discussion of these pro-restriction answers, let's briefly contrast the ways that Katz's, Quine's, and Wittgenstein's views lead to restrictions of the evidence base for linguistics, and in particular, how evidence from the cognitive sciences is excluded both in practice and in principle. For Katz, such evidence is irrelevant because the nature of language is such that language doesn't 'come from us' — languages are abstract objects independent of mind and matter. As a result, evidence about our psychology cannot illuminate the nature of language. For Quine and Wittgenstein, by contrast, language does 'come from us', but the evidence of psychology and neuroscience is nonetheless irrelevant to its study because of *the way* in which each of them thinks language comes from us. For Quine, the radically restrictive learning conditions of the infant lead to constraints on what evidence could be relevant to studying language. For Wittgenstein and his followers, such constraints emerge from the view of language as nothing more than a norm-governed cultural practice — from meaning being exhausted by use.

5 AGAINST RESTRICTING THE EVIDENCE BASE

Our central questions are (i) whether we should impose *a priori* and in principle restrictions on the evidence-base for linguistics, (ii) whether in particular such restrictions should be imposed on evidence from the cognitive sciences, and (iii) whether, even if allowed in as a matter of principle, there is reason to think that neuropsychological evidence won't in fact turn out to be relevant. (Recall Dirk Gently: he may be right that all evidence is *in principle* relevant to finding a lost

cat; but he is surely wrong that a two-week trip to Bermuda is as good a way as any to begin the search.) The views of language and linguistics discussed in §§2-4 of this chapter all give a positive answer to our questions, and on all three counts.

In this section, we argue for the alternative view that no *a priori* restrictions should be imposed, and hence that evidence from the cognitive sciences — because evidence from any area of inquiry — is potentially relevant to linguistic inquiry. More than that, we will argue that such evidence has been and will continue to be relevant in practice — in a way that, say, the number of planets orbiting Alpha Centauri is unlikely to be. The game plan for the rest of the chapter is as follows. In this section, we respond to the pro-restriction arguments introduced above, and provide some novel anti-restriction arguments. In the next section, we offer a detailed example of how neuropsychological evidence has recently provided very surprising evidence about morphology, syntax and compositional semantics. (Many of the ideas below owe their origin to lectures by, and conversations with, Noam Chomsky. They do not, however, represent his views in this area. For his seminal writings on these topics, see especially [Chomsky, 1969, 1986, 2000, 2002].)

5.1 Replies to Wittgenstein

Recall the analogy, on the Wittgensteinian view, between a language and the waltz. Both are nothing more than social practices with conventions that govern their correct execution. From this viewpoint, psychological evidence about the internal states of participants in the practice looks irrelevant to studying the practice itself because what occurs internally has no bearing on one's successful participation in the practice — all that matters is what is outwardly displayed. So the only evidence that will be illuminating must come from observation of the practice.

There are two points to make about this view. First, a note about what practicing linguists do. Just as a matter of fact, linguists do not study human linguistic activity for its own sake, cataloguing patterns; instead, they are interested in explaining the *causes* of linguistic behaviour. An important cluster of these causes is taken to be internal to the humans producing them, and specifically, to be part of their cognitive apparatus, so that the linguist is seeking to explain something quite different from (and causally responsible for) the Wittgensteinian language-games. Hence, there is already a worry that these Wittgensteinian considerations don't speak to linguists.

Just because linguists do this, of course, does not show that what they are doing is right. Is this search for causes a mistake? It's not, if linguistics is any kind of science at all. That's because it cannot just be assumed that the theorist knows which phenomena *are genuinely linguistic*. (See [Fodor, 1981] on this.) To take a simple example, is the fact that people are unlikely to ever make a move in the language game with "The cat the dog chased died" a linguistic fact, or not? One school, inspired by Chomsky, would say that such avoidance is not a properly linguistic fact, because the cause is a curious and specific limitation of human short-term memory. Or again, other discoveries about human short-term memory

constraints may well help to explain why our linguistic practices have the general form they do — say, why the sentences we utter are typically not hundreds of words long. Thus the use of certain specific constructions, and our use of language more broadly, can be informed by studying human memory — so as to tell us which features of the 'game' are really due *to language*. This alone suggests that the restriction of evidence to what we can observe about linguistic interaction is unnecessarily strict.

Here is a second argument in reply to the Wittgensteinian stance. It again assumes that linguistics is some kind of science. It is bad scientific practice in general to assume *a priori* that certain kinds of evidence are not relevant to one's enquiry. For example, assuming that only terrestrial data will be relevant to understanding the ocean's tides might seem plausible. After all, the tides are on the earth, so it seems *prima facie* reasonable that they should be affected only by terrestrial interactions. As it turns out, of course, the tides are caused by the gravitational attraction exerted by the moon, so evidence about the moon turns out to be vital for explaining the tides. But this was not obvious at the outset of the enquiry. In a similar way, it is bad practice to prejudge the question of what evidence is relevant to helping us understand language — even if we suppose that Wittgenstein is correct to characterize language as a social practice.

The fundamental insight behind both of the foregoing replies to Wittgenstein, Dummett, *et al.* (as we are reading them) is this: the job of the scientist is to find out how things are connected. Since we don't know this in advance — not least because we don't ourselves fix how they are connected — it's not a good idea to rule out *a priori* that certain connections hold, sealing us off forever from looking in such-and-such locales. More than that, what the evidence base for linguistics itself turns out to be depends on how things end up being connected. (Put into a slogan: What the evidence base is turns out to be an empirical question.)

True enough, as investigation advances we may make enough progress in finding the causes of linguistic performances that we reasonably hypothesize that certain *a priori* possible evidential linkages do not obtain — and this can lead us to prioritize our search in practice. This is what makes it possible to answer "No" to "Should there be any *a priori* and in principle restrictions on linguistic evidence ?" while answering "Yes" to "Is there good reason to think that evidence of such-and-such kind will not actually play a part?" Given this, the Wittgensteinian characterization of language may, for all that has been argued above, privilege the data gleaned from the observation of linguistic moves in the sense that such data is sought first, and arguably trumps information from other sources. But the view is not thereby exempt from the more general maxim against excluding neuropsychological data *a priori*.

We have been drawing lessons from the practice of ordinary science to respond to Wittgensteinian arguments: scientists look for causes, they don't presume to know at the outset which phenomena properly fall within the purview of their domain of inquiry and, related to this, they do not stipulate beforehand that certain connections simply cannot hold. But what if one denies that the study of

language should be scientific? Some precisely read Wittgenstein as suggesting that the scientific approach is inappropriate here: it would be like a 'science of poetry' or a 'science of fashion journalism'. The grounds for restricting the evidence base could then be couched in terms of a slightly different analogy: it would be foolhardy indeed, when trying to arrive at a proper literary interpretation of Kurt Vonnegut's novel *Slaughterhouse-Five*, to do CT scans on monkeys looking at its notoriously peculiar title page. More than that, it's hard to see how literary criticism, or fashion journalism for that matter, could gain anything from neuropsychological findings about brains. And, taking this stance, one won't be at all moved by comparisons with other sciences — since the whole idea is that the study of cultural linguistic practices shouldn't be scientific at all. Truth be told, we are not sure how one can reply to a theorist who simply refuses to look at language scientifically. One can't even say: "Look how successful the science of linguistics has been so far", since to their eyes the whole discipline will seems confused. On a happier note, such an anti-theory Wittgensteinian also cannot make claims about how one ought ideally to restrict the evidence base for linguistics, since she would have us scrap the whole enterprise, in favour of another topic entirely. Put another way, as we suggested above, our questions are about linguistics in the sense in which it exists as a discipline — i.e., they are questions about what evidence those who publish in linguistics journals, teach in linguistics departments, get degrees in the area, etc., should appeal to. Specifically, the issue is whether practitioners of that discipline are misguided if they try to employ evidence from the cognitive sciences. It isn't really an answer to *this* question to say that one should simply abandon the science of linguistics altogether.

5.2 Replies to Quine

Like the Wittgensteinians, Quine is motivated by a commitment to emergentism, in the sense of taking languages to be human products — things that come from us, rather than being independent of us. At the same time, Quine endorses empiricism: he is very strongly committed to respecting both the results and the methods of the empirical sciences. As we read him, the former commitment underwrites his view that there cannot be more to a given language than what the child acquires, while the latter commitment underwrites his anti-nativism about language acquisition. From these two commitments, we suggested, Quine derives his view of how language should be studied: there can be nothing more to what she has to learn than what can be gleaned from external stimuli (there are no further linguistic facts); consequently, the linguist must confine herself to what can be learned from the evidence available to the child. Worse, appealing to other data is apt to mislead the linguist into thinking that such data decides questions about language that simply cannot be decided.

Before introducing the first difficulty for the argument, we need to make two distinctions about what "empiricism" means. First, there is a necessary versus necessary-and-sufficient issue. Some read "empiricism" as saying that sensory

experience is necessary for knowledge. Others, whom we have called "radical empiricists", think that — aside from innate physiological structure and general-purpose mechanisms such as forming associations — sensory experience is both necessary and sufficient for knowledge. Second, some read "empiricism" as a label for a doctrine in philosophical epistemology — a doctrine about which kinds of data support which kinds of theories. Others read it as a label for a view about how human infants develop psychologically. Since these distinctions cross-cut, we have four possibilities:

SENSES OF "EMPRICISM"	Experience is Necessary	Experience is Necessary and Sufficient
An Epistemological Doctrine about Justification of Theories	A	B
A Psychological Doctrine about Infant Development	C	D

With these distinctions in place, note that Quine adopts empiricism as a thesis about both epistemic justification for theories *and* infant's psychological development. At the same time, he takes both the developing infant and the theorizing scientist to be subject to the 'necessary and sufficient' demand – their knowledge does not merely *require* experience, but, general learning-theoretic machinery aside, comes *entirely* from experience. In sum, Quine endorses both B and D from the table above. More than that, if we read him correctly, then Quine understands B and D to be part and parcel of one doctrine, empiricism.[4]

With this terminological clarification in hand, let's revisit the argument from §3. The fact that there is a good reason to be an empiricist about epistemic justification, namely, the success of the empirical sciences, does not give us a reason to be an empiricist about human psychological development. We cannot assume that empiricism is the correct approach to one just because it looks like the correct approach to the other. Thus we cannot infer from the truth of A or B to the truth of C or D. The second problem is that the success of science supports only the 'necessary' side of the first distinction — the view that knowledge requires

[4]Again, our aim in this chapter is to explain what might be tempting about restricting the evidence base for linguistics. It's not exegesis. Nevertheless, since readers may wonder why we read Quine this way, we should at least say this: we can find in Quine no genuine effort to defend, on the basis of empirical evidence, anti-nativism about human development. This suggests that he takes empiricism, which is non-negotiable for him, to immediately license, as a sub-doctrine, anti-nativism in psychology: "Two cardinal tenets of empiricism remain unassailable, however... One is that whatever evidence there *is* for science *is* sensory evidence. The other... is that all inculcation of meanings of words must rest ultimately on sensory evidence" [Quine, 1969, 75].

observation or experience, rather than the view that nothing beyond observation or experience is required. That is, what observations of successful sciences support is A, not B. As a result, Quine cannot draw any reasonable inference from the success of the sciences to the extreme view that the child is an empiricist-learner whose only source of information is experience.

The first point, then, is that Quine hasn't supported one of his key premises. A further problem with this aspect of Quine's account is that his (empirical) claim that the child learns her language *solely* on the basis of her observations of others and their corrections of her behaviour is probably false. Poverty of stimulus arguments have been offered by Chomsky and many others to show that such data cannot possibly suffice to explain how the child comes to speak a language. What the child *acquires* (e.g. how to form yes-no questions properly, or how to group a wide range of sounds into instances of a single word) is not something that she could plausibly infer from the very limited range of data that she typically has. Instead, it seems that she herself contributes some information which combines with the data that she receives so that she ends up knowing how to speak a language. (See Laurence and Margolis [2001] and Stainton [2006] for recent discussions of the Poverty of Stimulus arguments, including philosophical challengers.)

However, even if anti-nativism were true, and the child contributed nothing of substance to her own language-learning, it still would not follow that the practicing linguist is in the same situation as the child with respect to how to learn about language. This constitutes a final problem for Quine's restriction of the evidence base to observable behaviour in the language group: it simply is not the case that how the studied object develops (in this case, how the child comes to learn her language) limits how a group should study that object. Consider carrots. Suppose, counterfactually, that carrots developed solely on the basis of nutrients that they receive from outside themselves. This still would not mean that the scientist who wishes to understand how carrots develop should study only their nutrients. He should study the carrots themselves, too, their effects on various things, how they rot, and perhaps even other root vegetables. Applying this to the case of language, even if the child were an empiricist-learner, this is not a reason why the scientist studying language should be restricted to the data and methods that are available to the child.

To sum up, Quine's constraints on the evidence base available to the linguist are rendered problematic for several reasons. First, radical empiricism about human development is not supported by reflections upon the methods of the successful sciences. Second, it is unlikely that children are empiricist-learners of language. Third, even if they were, this would not require that the scientist should study language using only the resources available to the child. As a result, Quine's view does not ultimately give us a principled reason to restrict the evidence base available for linguistics.

5.3 Replies to Katz

We end this section of rebuttals by recalling Katz's reasons for thinking that the evidence base for linguistics should be restricted. As discussed in §2, Katz takes human languages to be abstract objects, knowable by us and yet independent of mind and matter. As such, he claims, their nature will not be revealed by psychological evidence, for their nature is not dependent on our psychology.

Notice first off that Katz systematically takes mathematical objects as his point of comparison. However, his argument for regarding languages as abstract objects don't establish that they are of the mathematical kind. As discussed in §2, the range of abstract objects also includes such entities as Tolstoy's *War and Peace*. Unlike numbers, this abstract object seems to depend crucially on the activities of a particular human — there is only as much to *War and Peace* as Tolstoy put into it. In systematically comparing languages with numbers and algebras rather than abstract objects like *War and Peace* or the American Constitution, Katz exaggerates the extent to which abstract objects are divorced from the activities of people. This is especially important because of how abstract *creations*, unlike abstract mathematical objects, can be illuminated by psychology: accepting that novels such as *Slaughterhouse-Five* and *Moby Dick* are types does not turn American Literature into a branch of mathematics; and certainly some psychological evidence will bear on the nature of these novels, e.g., the mental states of Vonnegut and Melville when they wrote them. More broadly, insofar as these abstract things emerge from us, we can psychologically investigate *the thing known* and thereby find out about it, the abstract thing. (See [Higginbotham, 1983] for more on this.)

Katz will object that there is a good reason to assimilate natural languages to mathematical systems, viz. that both give rise to abiding necessary truths. The argument, recall, rests on the claim that only by seeing languages as mathematical can one account for properties such as analyticity, which Katz claims are exhibited by some linguistic items. Such properties, he argues, simply cannot be explained by regarding language as a mental object along the lines defended by the cognitivist. Our reply takes the form of a dilemma. Either there are language-based necessities or there are not. If, as Quine [1953a; 1953b] has suggested, there really aren't any, then there is nothing for Katz-style Platonism about language to account for. If, in contrast, there are language-based necessities, then it's plain enough how the emergentist can account for them: they are language-based! The reason why it's necessary that flawed gems are imperfect, for example, is because of something we English speakers did: we made "flawed" semantically related in the right way to "imperfect". (Of course if such necessities exist *and* humans cannot bring them into existence with their conventions, then a problem remains. However, Katz hasn't given us reasons to think that. See [Boghossian, 1997] for contemporary discussion of the larger issues surrounding analyticity.)

The first argument against treating languages as mathematical entities is that this misses their dependence upon us. A second argument against Katz's Platonist

view is this: while conceptually possible, it is not interesting. Fodor [1981, 158–9], for example, describes Platonism as, on the one hand, "unassailable", but on the other hand, as something that "nobody is remotely interested in". Fodor notes that Platonism about language does not *preclude* the pursuit of psycholinguistics, but simply remains indifferent to its findings. Insofar as Platonism dedicates itself to finding out which grammar(s) is consistent with the intuitions of speaker/hearers (which everyone grants to be relevant), it may run in parallel, up to a point, with psycholinguistics. At the point where they diverge, where the psycholinguist rejects certain candidate grammars on the basis of data *other than* intuitions accepted by the Platonist, Fodor complains that the Platonist's project ceases to be worthy of pursuit. The reason is that the Platonist view of language rests on *stipulation* — what it chooses to study is not determined by how the world is, but by the investigator. Obviously, the investigator is at liberty to choose which *language* she investigates; but more importantly, on Katz's view she seemingly also has considerable liberty in determining which data, and even which intuitions, are properly linguistic, and hence will count as the facts to which her grammar must be responsible. Unlike the psycholinguist, whose restrictions (if any) must be (eventually) empirically grounded, the Platonist linguist is not so constrained because what she aims to explain is not assumed to have an empirical foundation. Thus, she may study whichever language-intuitions pair she chooses.

Katz [1985, 177] has complained that the objection begs the question against the Platonist, to the extent that it forgets that the Platonist is not interested in 'speaker/hearer capacities' at all. If, as the Platonist insists, languages are abstract, but nonetheless real, objects, then ignoring language-intuition pairs that are empirically grounded in experimentation in favour of others that are not is no failing: "it can no more be to the discredit of Platonism that it doesn't pay attention to psychological capacities than it can be to the discredit of Fodor's psychologism that it doesn't pay attention to abstract objects" [1985, 177]. However, Katz underestimates the extent of stipulation that is involved in delineating the scope of the Platonist's enquiry into any particular case. Specifically, recalling a point made when discussing the Wittgensteinian view, Katz overlooks how much stipulation will in fact be required to isolate the 'purely linguistic facts' that he takes the Platonist linguist to be concerned with. Consider, for instance, the following list of facts:

- English speakers produce "Ouch!" instead of "Ay!" when injured
- We seldom encounter sentences like "The cat the dog chased died"
- One can request the salt by saying "Can you reach the salt?"
- When one says "Stella and Wayne just got married" one usually means that they got married to one another
- Certain linguistic greetings are more polite than others
- In writing, a sentence starts with a capital and ends with a period

- Every utterance of "I cannot speak a word of English" is false
- The sentence "Juan's mother and father like themselves" cannot mean that Juan's mother, his father and Juan all like themselves
- "The if two Uruguayan" is less grammatical than "The man seems sleeping", though both are ungrammatical to some degree
- "I lifted up him" sounds odd but "I lifted up John's dog" does not

The question is: which of these facts are properly linguistic? Katz's Platonist must account for all the *linguistic* facts about a given language... but first he must identify them. It is not clear how the Platonist is able to do this — there don't seem to be any constraints. One simply stipulates.

This point also speaks to Katz' complaint that mentalistic approaches may impose extrinsic demands on grammars. First off, it is not clear that Katz can draw a sharp line around the 'purely linguistic' properties, so that the 'conflicts' Katz points to may be quite difficult to cash out in a principled way.[5] More specifically, he takes it to be a point against such views that, for instance, what they term the correct grammar (the one speakers actually know/use) might be messier than other descriptively adequate competitors — say, for reasons to do with the evolutionary development of linguistic abilities through the co-opting of pre-existing non-linguistic mechanisms. But if we do not assume that the 'purely linguistic' facts can be found except by looking at psychological facts, then grammatical messiness may simply be a reality.

The final objection to Katz takes the form of a differential certainty argument. This is a kind of argument where one must choose between two propositions, typically one supported by abstruse philosophical reasoning, and one that seems immediately obvious. The strategy is to say, "Though we aren't sure what it is, there must be *something* wrong with the abstruse argument, because what it seemingly supports is far less plausible than other things we know". (An example: following Zeno, one might argue that because any distance can be divided in half, one cannot really walk across a room. Long before knowing what was wrong with this line of thought, a differential certainty argument can show that there has to be *something* wrong — for people cross rooms all the time.) Here is how the argument form applies here. Katz suggests that if languages are ontologically abstract, then psychological evidence is not relevant; he then affirms the antecedent of this conditional. His conclusion is that psychological evidence should be avoided. But this conclusion conflicts with the patent relevance of certain psychological facts to the nature of language. Take two short examples. (A much more detailed

[5]Louise Antony [2003] pursues this line of argument against Soames [1984], who, like Katz, holds that one can distinguish the 'purely linguistic' properties *a priori*. However, Soames differs from Katz in taking linguistics to be an empirical discipline, such that his arguments (Antony suggests) are more vulnerable to her criticisms. In any case, Katz would need to defend a principled distinction between linguistic and psychological properties, without *assuming* Platonism, in order for the 'conflicts' to constitute problems for the mentalistic approach.

one appears in the next section.) We'll begin with a very obvious real-world example. Even those who say that they aren't interested in psychology don't genuinely ignore it in practice. For instance, *everyone* who undertakes to describe a language takes on board that the language must be finitely specifiable, because otherwise it would be unlearnable. Everyone pays heed to learnability precisely because it's obvious that a grammar's being learnable is relevant to whether it's the grammar we use/know. But, of course, being learnable is a psychological requirement. Here is a more exotic case. There is a rather rare anomia that attacks the ability to use proper names, but leaves the ability to use quantifier phrases (e.g., "every dog", "several computers", "all hamburgers") more or less intact. Given this, consider a thought-experiment. Suppose we found out that this anomia leaves definite descriptions such as "The king of France" intact as well: this phrase is spared in just the way "Some king of France" is. This data-point clearly would not prove that definite descriptions are more like quantifier phrases and less like names. But surely a theorist who supposed, with Frege [1892], that definite descriptions are name-like would have to at least explain away the data. As a linguist, she could not just say, "That is irrelevant to my topic". Given such examples, differential certainty now comes in. Which is more secure, our sense that the anomia and learnability are relevant to linguistics, or our confidence in Katz' abstruse reasoning? The answer is clear. (Of course, this argument doesn't tell us exactly where Katz went wrong in his reasoning. On that issue, however, see above.)

Before moving on, it's worth making a point about the relationship between what one takes the ontology of languages to be and one's view on methodology. Katz, for example, sometimes writes as if the very fact that languages are abstract entities means that psychological evidence will not be apposite. Certain Wittgensteinians equally seem to suppose that just because a language is, ontologically speaking, a normative cultural practice, it follows that psychological evidence should be avoided. Ontological reflections can also *support* an ecumenical methodology, of course: if languages just are mental items, then of course mentalistic evidence will be relevant both in principle and in practice.[6] This recognition can then encourage the idea that it's *only if* languages are mental items that mentalist evidence will be so relevant. But, say we, it's not necessary, for evidence from all sources to be permitted, that one follow Chomsky and take languages to be internal psychological entities. To the contrary, the relationship

[6]By the way, taking languages to be mental items is not, we think, ultimately in conflict with taking languages to also be abstract or social in certain regards. Suppose, for example, that what determines which linguistic facts obtain is something directly about the minds of its native speakers. Still, those 'internal' facts will fix a cluster of 'external' facts — about which sentences mean what, which sentences are well-formed and to what degree, etc. (Put crudely, the intensionally correct mental grammar settles which other grammars are extensionally correct. See [Laurence, 2003] for more.) What's more, even supposing this is how the facts are settled, there is no bar whatever to writing grammars that don't track the psychology of native speakers, to capture the resulting external facts, for whatever purpose — and to doing it in a way that distinguishes correct versus incorrect grammars. (For instance, language instructors may want a correct grammar for English that is easy to teach to Malagasy speakers.)

between ontology and methodology is far less tight than this. What matters is not what kind of thing a language is — concrete, abstract, mental, etc. — as much as what languages are connected with and (especially) what they are constituted by. Specifically, if the nature of our languages emerges from us — if as the metaphysicians say it 'supervenes' on us — then even if languages themselves are abstract things, or social things, rather than mental things, mental evidence can easily be relevant. (Just think of the waltz, or *Slaughterhouse-Five*. Why exactly shouldn't evidence about the psychology of dancers, readers and authors bear on the nature of these?)

To summarize so far, we began by introducing three families of arguments in favour of restricting the evidence base for linguistics. The aim, recall, was to make clear why anyone would feel tempted to do this; it was not to provide a detailed exegesis of the theorists discussed. Having provided arguments in favour of restriction, we then responded to them. We hope to have shown that, though imposing restrictions is not absurd or fool-hardy, it is ultimately incorrect. (Further arguments along the same lines may be found in [Antony, 2003], [Laurence, 2003] and [Stainton, 2001].) Most of our rebuttals have been directed to the issue of restrictions *a priori* and in principle. It remains to consider in detail whether results from the cognitive sciences are likely, in practice, to afford important evidence. We address that issue next, by means of a detailed example.

6 A DETAILED EXAMPLE

The previous section gives reasons why evidence of a psychological kind ought not be excluded in principle. Here we provide a specific example of how data from the cognitive sciences in fact play a useful role.

A standard view about the morphology of regular verbs like "start" and "link" is that their compound forms are built from the root plus a suffix. For instance, "started" is built from two elements: "start" [the root] and "-ed" [the suffix]. Put in terms familiar to philosophers and logicians, this amounts to something like: English contains syntactic axioms for "start" and for "-ed", and a formation rule that says how to put them together. There is then a semantic axiom for "start" and one for "-ed", and a compositional semantic axiom that captures the meaning-effect of combining these two items. The standard view says something different about irregular verbs like "swim", "hit" and "fall". Morphologically speaking, they apparently don't have a suffix; and in particular, their past tense forms ("swam", "hit", "fell") pretty clearly don't contain the suffix "-ed". What the standard view says about these is that they are atomic. Again, putting the point in terms of axioms and rules, the idea is that there is a special *axiom* that says what the past form of "hit" is, and a corresponding semantic axiom with its meaning.

This standard view isn't the only option. An alternative is that this supposed difference between regular and irregular verbs is an illusion. Not because "swam", "hit" and "fell" are compounds built from the "-ed" suffix; rather, because "started" *et al.* are equally atomic, both in terms of how they are formally

built up and in terms of how their meaning gets fixed in the language. (Note: the issue isn't whether the present and past tense forms of a regular verb are *semantically* related. Of course they are. The issue is what kind of rules account for their semantic relatedness: Are they semantically related because they share a morphosyntactic element — e.g., because the very same item "start" is brought to bear in deriving the meaning of both "start" and "started"? Or are they semantically related for some other reason?)

In short, we have two hypotheses about English. One, corresponding to the standard view, says that there is one set of rules for regular forms, and a different set of rules for irregular forms. The other says that there is one set of rules that applies to all verbs — so that the supposed regular/irregular distinction isn't linguistically important. Having introduced this disagreement about English, we will now explain some recent experiments in (neuro)psychology. The point, given our questions, will be that the results of these neuropsychological experiments pertain evidentially to the properly linguistic question.

Numerous psychologists have investigated how we process and represent past tense forms. They have looked in particular at whether native speakers of English, to arrive at the meaning of a sentence like (2), process their language in terms of (1):

1. *past tense = verb stem + -ed*

2. John walked past the Eiffel Tower.

There were two psychological hypotheses to consider. Hypothesis 1 has it that regular past tense forms are processed in accordance with the rule in (1), while irregular forms are stored separately in their entirety. An alternative hypothesis, Hypothesis 2, is that all past tense forms, regular and irregular, are stored in the lexicon individually and retrieved whole, without a rule like (1) being represented anywhere. A good way of testing this is via priming experiments. The idea behind such experiments is that recognizing a word is easier and therefore quicker after one has just been exposed to the same word or a closely related one. For instance, you are quicker to recognize "cat" as a word if you've just read "cat" a moment before. If Hypothesis 1 is right and regular past tenses really are processed as *verb stem + -ed*, one would expect the past tense to prime the stem to the same extent as the stem itself, because processing the past form necessitates accessing the stem form. To give an example, if you process "walked" as "walk + -ed", reading "walk" immediately after "walked" should make you quicker to recognize "walk" as a word, much the same way as reading "walk" immediately after "walk" would. With irregular past tenses, on the other hand, one would not expect the same amount of priming if the first hypothesis is correct. By contrast, Hypothesis 2 would predict no difference in priming effect across 'regular' and 'irregular' verbs: according to that account, processing "walked" does not require you to access "walk" any more than processing "went" requires you to access the actual word "go" — and thus there should be no difference in how fast you recognize "walk" after "walked" and "go" after "went".

Stanners *et al.* [1979] conducted such a priming experiment using reaction times (i.e., how long it takes a subject to recognize a letter or sound sequence as a word) and the results they found support Hypothesis 1: regular past forms prime the stem to a similar extent as the stem itself, while there is less priming in the case of irregular past forms. Applied to an example, subjects recognize "walk" as a word more quickly after reading "walked" than they recognize "go" after reading "went". Subsequent studies replicated these results for regular forms, i.e. they all agree that regular past forms prime their stems. However, Münte *et al.* [1999] point out that the results for irregular forms are less clear. For instance, Kempley & Morton [1982] found no priming in the irregular cases, while Fowler *et al.* [1985] and Forster *et al.* [1987] report full priming. According to Münte *et al.*, this highlights that priming studies *that crudely test reaction times* are problematic and may not be getting at the right phenomenon. It could be, for instance, that the reason for priming in regular past tense forms is that the past forms (e.g. "walked") contain the stems (e.g. "walk") phonologically and orthographically, though the former is still genuinely atomic in terms of how it is stored. In contrast, the equally atomic irregular past forms (e.g. "went") don't resemble their stems (e.g. "go") phonologically and orthographically to the same extent.

Because of these problems, Münte *et al.* set out to do a priming study that doesn't measure reaction times. Instead, they made use of technology that measures electric brain activity. It is known that people's event-related brain potentials (ERPs) pattern in particular ways when they're exposed to a word, a picture or a face for the second time (e.g. [Rugg, 1985]). In other words, measuring the electric activity in people's brains gives an indication of whether they've processed a particular form (linguistic or otherwise) before. Applying this to the question of past tense morphology, Hypothesis 1 about processing would be supported if verb stems following regular past tense forms led to ERPs characteristic of repetition, while verb stems following irregular past forms didn't. Hypothesis 2 would be supported if neither caused the brain activity typical of repetition.

A main goal of Münte *et al.*'s study was to avoid the possibility that the priming effect in regular past forms came down to the formal resemblance between regular past forms and verb stems. They wanted to make sure that the differences in observed effect, if any, was due to a difference in how morphology, syntax and semantics are processed — not because of very superficial differences between regular and irregular forms. For this reason, they did not expose subjects to the prime (e.g. "walked") and the target word (e.g. "walk") in immediate succession. Instead, they presented them with lists that always had between five and nine words between the prime and the target word — because previous evidence suggested that formal priming was a short-term effect. They also made sure to present the prime and the target word in different case letters, e.g., if the prime was written "walked" the target word would be written "WALK". Finally, they introduced two control conditions. In a phonological control condition the target word was preceded by a word that did or didn't share its initial phonemes (e.g. "sincere"–"sin" vs. "board"–"sin"). The nonce (invented) word control included made up words

that either followed a regular or an irregular pattern (e.g. "renited"−−"renite" vs. "ploke"−"plike"). The rationale behind this was the following. If the priming effect in the regular past tense case rests on the superficial formal resemblance between the past form and the verb stem, the same effect should also appear in the case of phonological similarity (e.g. "sincere"−"sin") and regular nonce words (e.g. "renited"−"renite"). If the priming effect were due to purely semantic relatedness (e.g., if it were due to the fact that "start" and "started" both pertain to *beginning*, while "leave" and "left" both pertain to *departing*), it should be produced by both regular and irregular past forms. Should it turn out, however, that it's only regular past tenses that prime the verb stem, Hypothesis 1 would be supported: processing "walked", for example, would, indeed, involve accessing the actual morphosyntactic item "walk" and the suffix "-ed", and deriving the meaning of the compound from these items.

What Münte *et al.* found was the following. Regular past tense forms resulted in ERPs typical of repetition when the subjects were presented with the verb stems. Irregular past forms did not have this effect. Interestingly, the phonological control condition where a target word was preceded by a word sharing its first few phonemes also produced no priming effect. Similar results were obtained in the nonce word condition: the characteristic patterns of electric brain activity associated with repetition did not appear, irrespective of whether the invented words followed a regular or irregular pattern. In terms of our example, reading "walk" after "walked" produced the typical repetition ERP pattern, while reading "go" after "went", "sin" after "sincere", and "renite" after "renited" didn't. This suggests that the priming effect in the regular past tense cases cannot be due to a superficial resemblance between the past form and the verb stem (if that were the case at least one of the resemblances in the control conditions should have led to the same effect) and it cannot be due to semantic relatedness either (otherwise irregular forms should also have a priming effect). Instead, it seems clear that subjects do access the verb stem in processing the regular past form, but not in processing the irregular form. This supports Hypothesis 1, whereas Hypothesis 2 is not compatible with the evidence found by Münte *et al.*.

Those are the neuropsychological results. Now consider what these experiments suggest about English. If the standard view about the contrasting morphology, syntax and compositional semantics for regulars versus irregulars is correct, this psychological effect is immediately accounted for. If English contains a syntactic axiom for "start" and another for "-ed", and it contains a semantic axiom for "start" and another for "-ed", then someone who knows English knows all of these rules, and she would naturally employ them to understand "started". Since she uses the axioms for "start" to understand "started", the various psychological and neurological priming effects are independently predicted to occur. Making the point in a general way, claims about the inner workings of the thing-known have implications; in particular, such claims have implications for psychology and neuroscience. Now, the thing-known in the case at hand is a language. So, claims about the structure of a language will yield testable predictions in the cognitive

sciences. To be clear, to affirm that there will be testable predications is not to suggest that, e.g., the results presented above out-and-out *prove* that English has this structure. There may be better explanations of the observed effects, or the effects may be shown to be illusory, and so on. The question, applied to our detailed example, is this: could a reasonable person inquiring into the morphology, syntax and semantics of English, who is not wholly blinded by ideology, react to these results by saying "I don't care. My view makes no predictions about processing"? We think not.

One last thing. It may be objected that the problem isn't so much using psychology to find out about the rules of English, but the very search for "the right linguistic rules" in the first place. Stich [1972] suggests, for instance, that from a mathematical point of view there will be more than one 'descriptively adequate' grammar for any given language. In other words, there will be more than one grammar that succeeds in generating all the well-formed expressions of a language from its primitive parts. To borrow a pair of examples from that paper, if A were a descriptively adequate grammar for English, then we could construct another descriptively adequate grammar, B, by adding a few superfluous rules to A, or by replacing one of A's rules with several distinct rules that collectively cover all the cases covered by the original rule.[7] Building on such reflections, one might insist that the only job a linguist should take on is describing the set of well-formed sentences: taking seriously the question of which axioms actually make up the language is already conceding too much. Continues the objection, given that the foregoing considerations about ERPs and the like don't speak to the nature of English *understood simply as such a set of sentences*, they aren't pertinent to the nature of the language after all. This is a familiar objection, but, following James Higginbotham [1985], we think it actually carries little weight. First, the notion of a set of well-formed formulae that we are trying to describe comes to us from the study of invented logical languages. In those languages, a formula is either fully well-formed or it is fully ill-formed, and ill-formed formulae lack meaning entirely. What's more, with invented logical languages there isn't an empirical issue about which formulae are well-formed: this is stipulated, not discovered. But human languages just aren't like that. Natural language expressions that are somewhat ill-formed can be meaningful, hence must be accounted for; natural language expressions come in degrees of well-formedness, and this too must be accounted for; and reasonable people can disagree, on empirical grounds, about which category an expression belongs to. So, the 'pre-established set of sentences' picture just does not fit well with natural, spoken, languages. Second, in this chapter we have been considering whether those working in the discipline of linguistics should pursue their task while eschewing certain sorts of evidence. Put otherwise, the word

[7] As Stich points out, less trivial variations are also possible, but they are vastly more difficult to construct. The background issue here, by the way, is 'indeterminacy': Stich [1972] exhibits a sort of Quinean [1960] pessimism about the possibility that there is a single descriptively adequate grammar internalized by speakers of a language, seeing room instead for an indeterminacy of grammars. For a reply to such worries, see [Chomsky and Katz, 1974].

"linguistics" in our title is to mean... *the actual discipline of linguistics*. Now, even linguists who are suspicious of taking linguistics to be a psychological enterprise would emphatically reject the idea that the only thing they are trying to do is find some way, any way, of characterizing the well-formed expressions of various languages. So, even beyond its other weaknesses, this objection simply doesn't speak to the issues actually at play.

We take this example about priming effects to afford several lessons. Before reviewing them, however, let us stress again what the lesson is *not*: it's not that the distinction between irregular and regular verb forms is linguistically important. It doesn't matter at all for our purposes who ends up being right about that dispute. Instead, what one should take away from the example is, first, the cogency in this particular case of the differential certainty point: it's far more clear that reaction times and ERP data are relevant to the nature of English verbs than it is that, say, English is an abstract object, or a cultural practice such that... In particular, second, it couldn't be more clear both that infants do not have access to the results of priming experiments, and that it would be well-nigh irrational to simply ignore such evidence. Third, the structure of the system known has a causal impact on our minds/brains. Because of this, we can make inferences to the best explanation from neuropsychological effects to the structure of the thing-known. In the present case, for instance, we can infer from contrasting processing effects for regular and irregular verbs to contrasting rule systems in the language known: such a difference in rule systems would account for the observed effects. Finally, this case illustrates nicely that evidence from domains never imagined, not even fifty years ago, can end up bearing on a linguistic issue. As hinted, this is possible precisely because we are trying to find out how things are connected — and so we can end up being quite surprised by what connections actually obtain. (Recall the tides and "terrestrial evidence" example. Another very nice illustration appears in [Antony, 2003]: drawings on Amerindian ceramic pots turned out to be relevant to establishing the timing of a supernova.)

Our goal in this chapter has been to introduce the debate about whether there should be restrictions on the evidence base for linguistics. Specifically, we took up the question whether the evidence base should be restricted *a priori* and in principle, and whether it should exclude, in principle or in practice, evidence from mentalistic psychology and neuroscience. We described three views — whose origins lie in some sense with Katz, Quine, and Wittgenstein — that lead to placing such restrictions on the evidence base for language study. Our aim was to make clear why, on such views, there are *prima facie* plausible grounds for restricting the evidence base for linguistics. In the final two sections, however, we responded to the three pro-restriction views, arguing that the restrictions they seek to impose ultimately should be avoided.

ACKNOWLEDGEMENTS

We are grateful to Steven Davis, Ray Elugardo and Dave Matheson for comments on an earlier draft. Portions of this paper were presented at the Workshop on Semantics and Psychological Evidence at Washington University in St. Louis, April 15^{th}, 2005. We thank the organizer Sam Scott and the audience members for helpful feedback. Especially useful were comments from participants José Bermúdez, Philip Robbins and Ken Taylor.

BIBLIOGRAPHY

[Antony, 2003] L. Antony. Rabbit-pots and supernovas: On the relevance of psychological data to linguistic theory. In A. Barber (ed.), *Epistemology of Language*. Oxford: Oxford University Press, 47–68, 2003.

[Barber, 2003] A. Barber. Introduction. In A. Barber (ed.), *Epistemology of Language*. Oxford: Oxford University Press, 1–43, 2003.

[Boghossian, 1997] P. Boghossian. Analyticity. In B. Hale and C. Wright (eds.), *A Companion to the Philosophy of Language*. Oxford: Blackwell, 331–368, 1997.

[Chomsky, 1969] N. Chomsky. Quine's empirical assumptions. In D. Davidson and J. Hintikka (eds.), *Words and Objections*. Dordrecht: D. Reidel, 53–68, 1969.

[Chomsky, 1986] N. Chomsky. *Knowledge of Language: Its Nature, Origin, and Use*. New York: Praeger, 1986.

[Chomsky, 2000] N. Chomsky. *New Horizons in the Study of Language and Mind*. Cambridge, UK: Cambridge University Press, 2000.

[Chomsky, 2002] N. Chomsky. *On Nature and Language*. Cambridge: Cambridge University Press, 2002.

[Chomsky and Katz, 1974] N. Chomsky and J. Katz. What the linguist is talking about. *Journal of Philosophy*, 71, 347–367, 1974. Reprinted in N. Block (ed.), *Readings in the Philosophy of Psychology*, Vol. 2. Cambridge, MA: Harvard University Press, 223–237, 1981.

[Davidson, 1979] D. Davidson. The inscrutability of reference. *The Southwestern Journal of Philosophy*, 10: 7–19, 1979. Reprinted in his *Inquiries into Truth and Interpretation*. Oxford: Oxford University Press, 227–241, 1984.

[Devitt, 2003] M. Devitt. Linguistics is not psychology. In A. Barber (ed.), *Epistemology of Language*. Oxford: Oxford University Press, 107–139, 2003.

[Dummett, 1973] M. Dummett. The philosophical basis of intuitionistic logic. In his *Truth and Other Enigmas*. Cambridge: Harvard University Press, 215–247, 1973.

[Fodor, 1981] J. Fodor. Some notes on what linguistics is about. In N. Block (ed.), *Readings in the Philosophy of Psychology*, Vol. 2. Cambridge, MA: Harvard University Press, 197–201, 1981. Reprinted in J. Katz (ed.), *The Philosophy of Linguistics*. Oxford: Oxford University Press, 146–160, 1981.

[Fodor, 1983] J. Fodor. *The Modularity of Mind*. Cambridge, MA: MIT Press, 1983.

[Forster *et al.*, 1987] K. I. Forster, C. Davis, C. Schoknecht and R. Carter. Masked priming with graphemically related forms: repetition or partial activation? *Quarterly Journal of Experimental Psychology*, 39: 211–251, 1987.

[Fowler *et al.*, 1985] C. Fowler, S. Napps and L. B. Feldman. Relations among regular and irregular morphologically related words in the lexicon as revealed by repetition priming. *Memory and Cognition*, 13: 241–255, 1985.

[Frege, 1884] G. Frege. *Foundations of Arithmetic*, 1884. Reprinted in M. Beaney (ed.), *The Frege Reader*. Oxford: Blackwell, 84–129, 1997.

[Frege, 1892] G. Frege. On sense and reference. In M. Beaney (ed.), *The Frege Reader*. Oxford: Blackwell, 151–171, 197.

[George, 1989] A. George. How not to become confused about linguistics. In A. George (ed.), *Reflections on Chomsky*. Oxford: Blackwell, 90–110, 1989.

[Higginbotham, 1983] J. Higginbotham. Is grammar psychological? In L. Cauman *et al.* (eds.), *How Many Questions?* Indianapolis, IN: Hackett, 1983.

[Higginbotham, 1985] J. Higginbotham. On semantics. *Linguistic Inquiry*, 16(4): 547–593, 1985.
[Katz, 1981] J. Katz. *Language and other Abstract Objects*. Totowa, NJ: Rowman and Littlefield, 1981.
[Katz, 1985] J. Katz. An outline of Platonist grammar. In J. Katz (ed.), *The Philosophy of Linguistics*. Oxford: Oxford University Press, 172–203, 1985.
[Katz, 1990] J. Katz. *The Metaphysics of Meaning*. Cambridge, MA: MIT Press, 1990.
[Kempley and Morton, 1982] S. Kempley and J. Morton. The effects of priming with regularly and irregularly related words in auditory word recognition. *British Journal of Psychology*, 73: 441–454, 1982.
[Laurence, 2003] S. Laurence. Is linguistics a branch of psychology? In A. Barber (ed.), *Epistemology of Language*. Oxford: Oxford University Press, 69–106, 2003.
[Laurence and Margolis, 2001] S. Laurence and E. Margolis. The poverty of the stimulus argument. *British Journal for the Philosophy of Science*, 52, 2001.
[Lewis, 1975] D. Lewis. Languages and language. In K. Gunderstone (ed.), *Minnesota Studies in the Philosophy of Science*, Vol. VII, 1975. Reprinted in D. Lewis, *Philosophical Papers*, Vol. I. Oxford: Oxford University Press, 1983.
[Münte *et al.*, 1999] T. F. Münte, T. Say, H. Clahsen, K. Schiltz and M. Kutas. Decomposition of morphologically complex words in English: Evidence from event-related brain potentials. *Cognitive Brain Research*, 7: 241–253, 1999.
[Quine, 1953a] W. V. O. Quine. The problem of meaning in linguistics. In his *From A Logical Point of View*. Cambridge, MA: Harvard University Press, 47–64, 1953a.
[Quine, 1953b] W. V. O. Quine. Two dogmas of empiricism. In his *From A Logical Point of View*. Cambridge, MA: Harvard University Press, 20–46, 1953b.
[Quine, 1960] W. V. O. Quine. *Word and Object*. Cambridge, MA: MIT Press, 1960.
[Quine, 1969a] W. V. O. Quine. Ontological relativity. In his *Ontological Relativity and Other Essays*. New York: Columbia University Press, 26–68, 1969a.
[Quine, 1969b] W. V. O. Quine. Epistemology naturalized. In his *Ontological Relativity and Other Essays*. New York: Columbia University Press, 69–90, 1969b.
[Quine, 1987] W. V. O. Quine. Indeterminacy of translation again. *Journal of Philosophy*, 84, 1987.
[Rugg, 1985] M. D. Rugg. The effects of semantic priming and word repetition on event-related potentials. *Psychophysiology*, 22: 642–647, 1985.
[Soames, 1984] S. Soames. Linguistics and psychology. *Linguistics and Philosophy*, 7: 155–179, 1984.
[Soames, 1985] S. Soames. Semantics and psychology. In J. Katz (ed), *The Philosophy of Linguistics*. Oxford: Oxford University Press, 204–226, 1985.
[Stainton, 2001] R. Stainton. Communicative events as evidence in linguistics. In J. de Villiers and R. Stainton (eds.), *Communication in Linguistics*, Vol. 1: *Papers in Honour of Michael Gregory*. Toronto: Editions du GREF, 329–346, 2001.
[Stainton, 2006] R. Stainton (ed.). *Debates in Cognitive Science*. Oxford: Blackwell, 2006.
[Stanners *et al.*, 1979] R. Stanners, J. Neiser, W. Hernon and R. Hall. Memory representation for morphologically related words. *Journal of Verbal Learning and Verbal Behaviour*, 18: 399–412, 1979.
[Stich, 1972] S. Stich. Grammar, psychology, and indeterminacy. *Journal of Philosophy*, 79(22): 799–818, 1972. Reprinted in J. Katz (ed.), *The Philosophy of Linguistics*. Oxford: Oxford University Press, 126–145, 1981.
[Wittgenstein, 1953] L. Wittgenstein. *Philosophical Investigations*. Oxford: Blackwell, 1953.

EMOTION: COMPETING THEORIES AND PHILOSOPHICAL ISSUES

Jesse J. Prinz

1 INTRODUCTION

Everyone agrees that emotions play important roles in psychology. There is also some broad agreement about the kind of roles that emotions play. They are important for decision making, for evaluative judgment, for social behavior, and for motivating action. Emotions may also play important roles in perception and memory. Despite all this consensus on *where* emotions have an impact, there is very little consensus on *what* emotions are. Diametrically opposed theories of the emotions have been proposed, and their adherents continue to disagree. In this chapter, I will survey some of the disagreements. I will begin by considering a debate between two camps: those who regard emotions as cognitive and those who regard emotions as noncognitive. Within these camps, very different theories of the emotions have been proposed. I will then turn to three philosophical questions about emotions. These are issues that have an empirical dimension, but where philosophers may be especially well poised to make a contribution.

My survey of emotion research will not be neutral. This is an opinionated introduction. But few of the arguments below are decisive, and all of the controversies that I consider are enduring.

2 THEORIES OF EMOTION

2.1 Noncognitive Theories: James-Lange and Beyond

In ordinary talk, the word "emotion" is often used interchangeably with "feeling." Feelings are conscious episodes of experience. They are traditionally contrasted with thoughts. Twinges, twitches, and pangs are examples. To identify emotions with feelings is to contrast emotion with thinking, and this equation is known, therefore, as a noncognitive theory. Other emotion researchers insist that thoughts are essential to emotions. They defend cognitive theories. The debate between defenders of cognitive and noncognitive theories is the most important division in emotion research. I will survey some of the major contenders on either side of this division. I begin with the feeling theory because it comes closest to unreflective commonsense.

Handbook of the Philosophy of Science. Philosophy of Psychology and Cognitive Science
Volume editor: Paul Thagard
General editors: Dov M. Gabbay, Paul Thagard and John Woods
© 2007 Elsevier B.V. All rights reserved

If emotions are feelings, what kind of feelings are they? There are two possible answers to this questions: reductive and nonreductive. Reductive feeling theories identify emotions with a class of feelings that can be characterized without reference to emotions. Nonreductive feeling theories say that emotions are *sui generis*; e.g., the distinctive feeling of anger cannot be characterized in any other way — it is an angry feeling. The nonreductive view may be closest to the commonsensical view, but it has had surprisingly few adherents in emotion research. Spinoza may have conceived his two primary affects, pain and pleasure, as *sui generis* feelings, but he resisted explaining other emotions in this way.

The *sui generis* feeling view was subjected in an influential critique in the late 19^{th} century. William James [1884] said that commonsense is committed to the following account of how emotions take place. First, there is an eliciting event, say, an encounter with a bear; the bear is perceived, and the danger of the situation is recognized; this causes a distinctive feeling, which we call fear; the fear causes our bodies to react to the bear by taking flight. James argued that this sequence gets things backwards. Emotions do not precede bodily response. When we recognize that we are in danger, our bodies immediately react and fear is the feeling of the bodily changes underlying those reactions. Fear is the feeling of our hearts racing, our vessels constricting, our muscles tensing, our lungs inhaling, our eyes widening, and so fourth. Fear is a feeling of a patterned change in the body, not a *sui generis* feeling. James's theory is reductive because it identifies emotions with feelings that can be characterized (as I just did) without reference to emotions. But James's theory is still a feeling theory. Carl Lange [1885] proposed a very similar theory independently of James at around the same time, though he emphasized changes in blood vasculature, rather than more global and varied bodily changes. In some passages, Lange implies that the emotion is the bodily change itself rather than the feeling of the change. In other passages, he identifies emotions with bodily feelings. Related views have been recently defended by philosophers [Prinz, 2004; Robinson, 2004] and neuroscientists [Damasio, 1994].

Both Lange and James defended their theory by appeal to introspection. They asked their readers to imagine an intense emotion and then to subtract away every bodily symptom. If we try to imagine fear with no racing heart, with relaxed muscles, with steady breathing, and so on, there will be nothing left, no residual non-bodily feeling, that we would want to call fear. The James-Lange theory has an advantage over the *sui generis* feeling theory. It explains how emotion feelings arise, as does so without multiplying feelings beyond those that everyone is committed to, in so far as every one agrees that we can consciously experience certain bodily changes. James and Lange win on parsimony, and their introspective argument is very compelling. While emotional feeling may seem *sui generis* at first, we can usually associate them with particular bodily changes if we try to. Next time you have a strong emotion, see if you can't identify some of its bodily components. In some cases, bodily states are difficult to describe (visceral perturbations may be hard to localize precisely through introspection), but some somatic conditions are easy to discern: we have no difficulty noticing constrained

breathing, a racing heart, a lump in the throat, butterflies in the stomach, tensed muscles, laughter, tears, gasps, and grimaces.

The problem with this strategy for defending the James-Lange view is that it relies on introspection, and introspection can mislead. If we want to establish the necessity of felt bodily change, it would be better if we could find other kinds of empirical support. It occurred to James that a good way to test his theory would be to find people with spinal cord injuries who have no capacity to feel their bodies. If these people no longer feel emotions, the James-Lange theory would find support. When James wrote his article, he knew about the emotional consequences of spinal cord injury only through anecdotal reports. A systematic study of these individuals was first attempted by Hohmann [1966] decades after James wrote his seminar paper. Hohmann asked 25 spinal patients to report on how their emotions had changed since injury. Patients reported a reduction in all emotions, except for sadness and sentimentality. It is possible that these two emotions have bodily symptoms centered around the neck and head, which can still be felt after spinal injury. Hohmann's most striking finding was that the degree of emotional reduction increased with the height of the injury on the spinal cord. Higher injuries cause a greater reduction in bodily perception. This is just what the James-Lange theory would predict.

There are several ways to challenge the James-Lange theory. I want to consider five objections here. First, one might question the evidence. Hohmann's study has been challenged by Chwalisz *et al.* [1988]. They complain that Hohmann failed to compare his subjects' responses to control groups', and they claim that his findings have alternative explanations. In the 1960s, spinal patients were encouraged to resign themselves to inactive lives, and this may have lead to the reduction in emotional responses. In addition, Hohmann's subjects were relatively old, and their emotions might have reduced through aging. Chwalisz *et al.* examined younger spinal patients who were living more active lives. These subjects did not report that their emotions had diminished, and their self-reported levels of emotional intensity were actually higher than wheelchair bound control subjects.

The Chwalisz *et al.* study casts doubts on Hohmann's result and erodes support for the James-Lange theory. But the study may suffer from problems of its own. For example, Chwalisz *et al.* reply on subjects' recollections of how their emotions were before injury. The average time since injury for these subjects was nine years, and many of the subjects had sustained their injuries in childhood. Moreover, both this study and Hohmann's may be based on a false premise about the James-Lange Theory. If emotions are feelings of bodily change, it doesn't follow that the body must actually change every time an emotion is felt. A bodily feeling can be a feeling of the kind that one would have if one's body were to change, even if the change has not actually taken place. James makes this point in a footnote, and the point ahs been vigorously pursued by Damasio [1999]. Damasio defends James by arguing that bodily feeling can bypass the body via a pathway in the brain that activates regions associated with bodily perception without circuiting through the body. Of course, one might expect that this "as-if" loop would lead to bodily

feelings that are less intense than feelings brought on by actual bodily changes. That prediction is actually borne out by Chwalisz et al. Like Hohmann, they found that emotional intensity varies as a function of the height of spinal injuries. People with no capacity to perceive their bodies still experience emotions, but these are less intense than emotions that derive from the body.

The claim that emotions involve feelings of bodily states, even in the absence of the capacity to perceive the body, may seem desperate. The Chwalisz *et al.* study is consistent with the James-Lange theory, but only by adding an ad hoc mechanism. This concern can be answered, however, because there is independent evidence for the claim that emotions involve brain regions associated with the body. Every neuroimaging study of emotional response seems to show activation in somatic areas. Anterior cingulate cortex, primary and secondary somatosensory cortices, and insular cortex are common players. Each of these areas has been independently associated with perception and regulation of the body. The amygdala, which has been frequently implicated in the initiation of certain emotions (especially fear) is essentially a body-control center. It sends signals to structures such as the hypothalamus, the dorsal motor nucleus of the vagus, and central gray, all of which orchestrate patterns of bodily change. The functional anatomy of emotion is perhaps the best evidence for the claim that emotions involve perceptions of the body.

This still does not mean that James and Lange were right. Several other objections are available. A second objection focuses on the claim that emotions must be felt. The brain regions associated with emotion are involved in body perception, but we have no reason to think that activity in those regions always corresponds to conscious experience. If we can perceive patterned bodily changes unconsciously, perhaps we should say there can be unconscious emotions. This conjecture is the major distinction between Damasio's [1994] theory of the emotions and the James-Lange theory on which it is based. In other words, Damasio defends a somatic theory of the emotions, but not a feeling theory. Deciding between Damasio and the James-Lange theory is not an easy task, and I will return to the issue in section 3.1.

A third objection to the James-Lange theory involves the relationship between emotion and action. Intuitively, emotions make us act. We flee because we are afraid. We aggress because we are angry. We mope because we are sad. Behaviorists have sometimes been inclined to identify emotions with dispositions to behave in these ways, especially behaviorists in the philosophical tradition (cf. [Wittgenstein, 1953]). The James-Lange theory certainly relates to action, because the body patterns we are alleged to perceive are patterns that prepare us for behavioral response. The heart races in fear in order to send blood to the extremities in preparation for flight. So there is a sense in which action disposition theory and the James-Lange theory are sides of the same coin. There is, however, an important difference. For James and Lange, emotions are on the input side. They are perceptions of bodily change. For action disposition theorists, emotions are on the output side: they are commands for action. Without introducing commands

of some sort, the defender of James and Lange cannot account for the fact that emotions seem to cause action, rather than merely registering action. Of course, this is exactly what James wanted to challenge when he said flight causes fear, rather than the other way around. James expressly denies that emotions can contribute to psychological explanations in the way that has been presumed by folk psychology. Should we accept this claim?

Accepting James's claim that emotions are effects of action rather than causes is methodologically problematic. "Emotion" is a theoretical term borrowed from folk psychology. We postulate emotions to explain certain things. It is essential to folk explanation that emotions contribute to action. To deny that emotions play this role is like denying that water is clear and wet. Such reference fixing features are, arguably, non-negotiable. It is not essential that every emotion be related to some specific action on every occasion. Perhaps folk psychology doesn't go that far. But an adequate theory of emotions has to make sense of sentences such as, "He avoided heights because he was afraid of them." If the James-Lange theory cannot do that, then it may depart too radically from folk psychology. It may be an eliminativist theory, rather than a reductionist theory.

The defender of James and Lange can resolve problem in two compatible ways. First, it is possible to challenge the distinction between perceptions of action dispositions and issuing action commands. We ordinarily think of inputs and outputs as distinct systems in the mind brain, but this may be a false dichotomy. It is possible that the representations involved in perceiving the body also play a role in instructing the body to act. If so, bodily perceptions are action commands, and the James-Lange theory is equivalent to an action disposition theory. There is little direct evidence for this possibility. We have reason to believe that the very same neurons can be involved in both action and perception of action (e.g., mirror neurons), and we know that brain regions associated with emotion are implicated in both perception of the body and bodily regulation. Anterior cingulated cortex, second somatosensory cortex, and the insula are all noteworthy in this regard [Critchley *et al.* 2000; Adolphs *et al.* 2000, Kuniecki *et al.* 2003].

There is a second response available to the defender of James and Lange. Rather than insisting that emotions are nothing but perceptions of bodily changes, defenders of this approach can argue that emotions are *valenced* bodily perceptions. Valence can be understood as an internal motivational command that affixes itself to other mental states [Prinz, 2004]. Valence causes one to try to decrease (in the case of native valence) or promote (in the case of positive valence) the mental states to which they are affixed. A valenced bodily perception is a sensation of the body coupled with a command that says "More of this!" or "Less of this!" If emotions contain valence commands, they are motivating. This is a supplementation of the original James-Lange theory, but it preserves the core insight that emotions are perceptions of bodily changes. They are valenced perceptions.

The James-Lange theory may need to be supplemented in ways that depart more dramatically from the spirit of the original account. Evidence from neuroanatomy suggests that perceptions of bodily changes are necessary for emotions (though see

more on this below), but are they really sufficient? If emotions contain valence commands, then the answer is negative. Critics of James and Lange argue that other additions are needed as well. Before considering their recommendations, it's worth asking why James and Lange believed that perceived bodily changes were sufficient. Both of them argue for this conclusion.

Lange noted that emotions could be altered by drinking alcohol. He presumed that alcohol had this effect by modifying blood flow in the body. If alcohol can induce emotions by mere bodily change, then perceptions of the body are sufficient for emotions. James noted that one could affect ones emotions by changing one's bodily state. To cope with his own depression, he would adopt an upright posture and a joyful countenance, resisting every urge to slouch and frown. By changing his body, he changed his affective state. This effect has been confirmed through contemporary research on facial feedback. Emotional experiences can be altered by changing our facial expressions even if we do not realize we are doing so. Strack *et al.* [1988] had subjects fill out a questionnaire, without using their hands, by holding a pen that was held in their mouths. Some subjects were instructed to hold the pen between puckered lips, resulting in an inadvertent grimace, while others were told to hold the pen between their teeth, forcing an inadvertent smile. Among other things, they had to rate the humor of cartoons in the questionnaire, and subjects in the smiling condition gave higher ratings. Levenson *et al.* [1990] found that making facial expressions had global bodily effects. Different emotional expressions cause different patterned changes in the autonomic nervous system. Facial expressions evidently trigger more global bodily changes, and these affect one's emotional state.

These lines of evidence are suggestive but they do not establish the sufficiency of perceived bodily responses. Alcohol and facial feedback certainly affect the body, but they may have other effects as well. Those other effects may be essential components of our emotions. What sort of effects? Let me consider two. These can be regarded as two more objections to the James-Lange theory.

First, there is considerable evidence that emotions influence the way we process information. When we are afraid, we pay attention to threatening things in our environment [Mogg and Bradley, 1999]. When we are sad, our self-assessments switch from inflated to realistic [Alloy and Abramson, 1988]. When we are happy, our capacity to think creatively increases [Isen, 2000]. As far as I know, no one has attempted to induce these kinds of effects in a facial feedback paradigm, but I suspect that would be possible. Information processing effects are robust; they seem to be inextricably bound to emotions. The James-Lange story leaves them out. Oatley and Johnson-Laird [1987] go so far as to identify emotions with processing modes. I think this goes too far. Their theory does not give the body its due. It would be much more reasonably to propose a hybrid, according to which emotions involve both bodily perceptions *and* modes of information processing. This would be a departure from the pure James-Lange theory. One could try to preserve a pure James-Lange by arguing that processing modes are contingent effects of emotions rather than constitutive parts [Prinz, 2004]. If that argument does not

succeed, the fan of James and Lange may be forced to settle for the hybrid.

There is one final line of objection to consider. Many researchers believe that emotions essentially involve thoughts. Emotions may involve perceptions of bodily changes, but the objection goes, these do not suffice. Without an accompanying thought, bodily changes are mere twinges and pangs. They are not full-fledged emotions. I will explore this idea when I discuss cognitive theories in the next section. The evidence for cognitive theories is the most pressing objection to a pure James-Lange account.

In this section, I have focused on the James-Lange theory of emotions, which is, to my mind, the most tenable non-cognitive theory. Along the way, other non-cognitive theories have been considered. We encountered *sui generis* feeling theories, action disposition theories, and processing mode theories. I also noted that there is a close cousin of the James-Lange theory that identifies emotions with perceptions of the body that can be unconscious, rather than insisting that emotions are feelings. The James-Lange theory cannot be reconcile with the *sui generis* feeling theory, but it can, I argued, be reconciled with a version of the behavioral disposition view. Departing from the original formulation, defenders of the James-Lange theory may be tempted to treat emotions as hybrids that combine bodily perceptions with modes of information processing. This would by a hybrid theory, rather than a pure James-Lange theory. The hybrid would preserve an important feature of the original account, however. It would preserve the assumption that emotions do not have thoughts as component parts. Processing modes affect cognition, but they are not themselves cognitive in this sense. Thus, the hybrid theory is a noncognitive theory. The defender of James and Lange would abandon the noncognitive spirit of the original account if they were forced to concede that emotions have thoughts as components. That is the challenge I turn to now.

2.2 Cognitive Theories: Appraising Appraisals

The idea that emotions are noncognitive is unpopular in contemporary and psychology and philosophy. The majority of prominent emotion researchers insist that emotions have an essential cognitive component. Within psychology the major tide-change came when Schachter and Singer [1962] published their influential critique of the James-Lange theory. Schachter and Singer agreed that emotions had a necessary bodily component, but they sought to prove that this component is insufficient on its own. To do that, they conducted one of the most famous experiments in the history of psychology.

Subjects in Schachter and Singer's experiment were injected with adrenalin and told that it was a drug designed to improve vision. Some subjects were told that the injection would cause a state of arousal, and others were not. All subjects were then asked to wait in a room for a short time to let the drug take hold. In that room, the experimenters planted a stooge, who pretended to be a subject in the same study. Some of the subjects were placed in a room with a stooge who engaged

in silly antics, like throwing a paper airplane a playing with hula hoops. Others subjects were accompanied by a disgruntled stooge. In that condition, subjects were given an insulting questionnaire, and the stooge they were with griped about the questions, and ultimately tore up the questionnaire and stormed out. Subjects who had been informed about the affects of the injection, did not show strong emotional responses; subjects with the nasty stooge, acted as if they were angry; subjects with the silly stooge acted as if they were happy. The experimenters concluded that distinct emotions have a common underlying physiological basis (genetic arousal). What distinguishes emotions from each other and from non-emotional states of arousal is an accompanying cognitive attribution. We label arousal as anger in contexts of insult and as euphoria when ridiculous things occur.

The Schachter and Singer study looks like strong evidence for the necessity of a cognitive component in emotions. The study, however, is flawed. First, the fact that subjects in the two stooge conditions appeared angry and happy does not entail that they actually experienced those emotions. They may have just acted that way to be polite. In exit interviews, subjects in both conditions reported mild happiness, rather than rage or delight. There is also independent evidence that emotions can be physiologically differentiated [Levenson *et al.*, 1990; Philippot *et al.* 2002]; though see also [Cacioppo *et al.*, 2000]. Second, if the subjects did experience different emotions, it does not prove that emotions have a cognitive component. The generic state of arousal caused by the adrenalin may have transformed into distinctive patterns of physiological change in the two conditions. The experimenters did not measure physiological response during the experiment. The most striking result in the study is that the subjects who had been informed about the effects of the injection did not react strongly to the two stooges. On the face of it, this is actually quite odd, since the stooges antics were decidedly unsubtle. Perhaps the subjects who were informed about the effects of adrenalin were preoccupied with monitoring the effects of the drug. Or perhaps these subjects experienced emotions, but mistook their emotions for effects of the drug. This is the opposite of the interpretation offered by Schacter and Singer: cognitive attributions are not necessary for emotions; but they can lead us to think we are not having emotions when we actually are. Finally, if these objections are not sufficient, it also turns out that the experiment has not held up on attempted replication [Reisenzein, 1983].

Despite the shortcomings of the Schachter and Singer experiments, there are still good reasons for taking cognitive theories seriously. Here are three simple arguments. First, emotions are influenced by thoughts. We get angry when we believe that we have been insulted, and we become afraid when we believe that we are in danger. In experimental settings, changing a person's beliefs can change her emotions. It's not clear how a thoughts could cause emotions if emotions did not contain thoughts. Second, emotions seem to be meaningful. Fear is not just a feeling, it seems, but a signal, indicating that I am in danger. The most natural explanation of this is that emotions contain thoughts. Thoughts after all are meaningful. Third, emotions can have intentional objects. We can be afraid

that there is a prowler in the house, but we cannot have a pang or a twinge that there is a prowler in the house [Pitcher, 1965]. Thoughts can have intentional objects.

Arguments of this kind have led to promotion and proliferation of cognitive theories. The Schachter and Singer theory emphases an attribution process whereby why try to make sense of our inner feelings. Other cognitive theories emphasize, not our evaluation of inner feelings, but our evaluations of objects, situations, and events in the world around us. Fear, for example, might be said to involve the judgment that I am in danger. To be afraid of a prowler is to judge that I am in danger due to the prowler. Aristotle is credited with having one of the first theories of this kind. In the *Rhetoric* he says that anger involves the judgment that I have been slighted. For Spinoza [1677], an emotion is a judgment plus a feeling of pain or pleasure. For example, he defines pity as pain accompanied by the thought that evil has befallen someone that we take to be like ourselves. Contemporary philosophers have proposed similar views. Greenspan's [1988] theory of emotion echoes Spinoza's. Solomon follows Sartre in characterizing emotions as ways of projecting meaning onto events in the world around us [Solomon, 1976]. Armon-Jones [1990] says emotions are ways of conceptualizing things (e.g., conceiving of something as dangerous). In a similar vein, Nussbaum [2001] talks of value-laden appearances.

Psychologists have developed similar proposals. According to the most dominant tradition in psychology today, emotions always involve appraisal judgments. These are defined as judgments about organism-environment relations that bear on well-being [Lazarus, 1991]. Psychologists develop the idea by suggesting that there is a small set of appraisal questions that we ask ourselves when we consider an object or situation. Each emotion corresponds to a different set of answers to these questions [Arnold, 1960; Smith and Lazarus, 1993; Roseman, *et al.* 1990; Scherer, 1997]. To illustrate, consider the theory that has been developed by Klaus Scherer [1997]. He calls the appraisal questions "stimulus evaluation checks." These come along five dimensions. We ask whether the stimulus is novel or familiar, whether it is pleasant, whether it is relevant to our goals, what the potential is for coping with it, and whether it is consistent with out standards. Some of these dimensions have sub-questions. For example, in asking about goal relevance, we assess whether the stimulus is conducive to our goals and whether it demands urgent response. When asking about coping potential, we assess who the agent is responsible for the stimulus and whether we have control over the situation. Scherer uses these questions to distinguish emotions. Anger is the emotion that results when we judge that something that is inconsistent with goals, urgent, caused by another person, within our power to change, and incompatible with our standards. Sadness is the result of assessing a stimulus as unpleasant, not urgent, inconsistent with goals, and beyond our power to change. Other emotions correspond to other judgments, and no state qualifies as an emotion if it is not the result of these stimulus evaluation checks.

Cognitive theorists do not generally deny that bodily responses are involved in

the emotions. They think that these responses occur in response to appraisal judgments. An emotion might be regarded as a compound state or episode involving both judgments and a bodily response. Some cognitive theorists have pushed this farther by denying that physiological responses are necessary. Solomon [1976] argues that we can have an emotion without any bodily perturbation. Long-standing emotions are the best example. One can be in love for years, but one does not sustain a consistent bodily pattern all this time [Solomon, 1976]. This can be regarded as another argument against defenders of noncognitive theories.

The arguments favoring cognitive theories are intuitively compelling, but they are not powerful enough to refute noncognitive theories. Let me consider them in turn. First, consider the fact that emotions can be caused by thoughts. Does this show that emotions are not cognitive? I don't think so. There are two ways to interpret the argument. On one reading, the argument alleges that thoughts can cause only other thoughts. This principle cannot be maintained. Cognitive theorists must admit that thoughts can cause bodily perturbations. By their own account, thinking about dangers can cause one to shudder. On a second reading, one might take the argument to imply that the relationship between emotions and thoughts is systematic: the same thought always causes the same emotion. If thoughts about danger reliably cause fear, mustn't such thoughts be essential to fear? This question implies a non sequitur. Suppose that thinking about rancid food always causes me to wrinkle my nose. Must such thoughts be essential to nose wrinkling? Obviously not. For one thing, the thoughts that lead to nose wrinkling are not component parts of the facial expression. Cognitions might be essential causes of emotions without being emotions. For another thing, nose wrinkling can be caused by many things, cognitive and noncognitive. Smelling rancid food would do the trick.

There is good evidence that emotions, like nose wrinkling, can be caused without mediating thoughts. Zajonc [1984] gives the example of emotions induced by drugs, by direct brain stimulation, and by facial feedback. These cases may be exotic, but there are many more mundane instances of noncognitive emotion induction. Emotions are often triggered by perceptions rather than thoughts. Seeing a snake can trigger a fear response before any judgments have time to form — indeed, fear can be perceptually triggered via pathways from the optic nerve to the amygdala, bypassing the neocortex, which is the presumed seat of judgment [LeDoux, 1996]. Cognitive theorists might argue that the state caused by seeing a snake cannot count as fear unless there is an accompanying judgment. This, however, begs the question. The state induced by seeing a snake includes a physiological pattern that is just like episodes of fear that have cognitive causes. We would experience this state as fear. It is a paradigm case. And there is no evidence that judgments are necessarily required in inducing the state or once the state has been induced.

What about the argument that emotions are meaningful? Does this show that emotions must be cognitive? Here too I am skeptical. A mental state is meaningful if it represents something. According to one of the leading theories of representation, a mental state represents something if it is reliably caused by that thing, and

if it has the function of being so caused [Dretske, 1986]. The bodily response associated with fear is reliably caused by dangers, and it is adapted to that purpose. Our perception of that state is also a reliable danger detector. When we perceive that our body has entered the pattern characteristic of fear, we thereby pick up the information that we are in danger. Our capacity to perceive patterned bodily responses is, plausibly, evolved for the purpose of picking up on such information. If this is right, then perceptions of patterned changes in the body represent whatever external conditions reliably cause those changes to occur. The perception of the bodily state associated with fear represents danger. One might say that such a perception *registers* a bodily change, but *represents* a danger, because our capacity to perceive such bodily patterns was attained because of the role such perceptions play in danger detection. Perceptions of other bodily patterns represent losses, offenses, goal attainment, and so on. They represent the same relational properties that appraisal judgments represent, but they are not cognitive. I call such meaningful body perceptions embodied appraisals [Prinz, 2004].

Defenders of cognitive theories will resist this embodied alternative. They will argue that, even if bodily perceptions can be meaningful, they lack the rich informational structure that empirical work on the emotions has brought out. In particular, they cannot carry the kind of information emphasized by cognitive appraisal theories. Recall that these psychologists associate emotions with several dimensions of appraisal, rather than simple categories such as loss, danger, or insult. A perception of a bodily response could be a danger detector, but it could not plausibly represent the complex fact that there has been a novel, unexpected, urgent, goal obstructing, power depriving event caused by another person or some other external factor. This very complex set of judgments is essential to fear on Scherer's theory and theories like it. A highly structured cognitive representation would be needed to capture this complex multidimensional assessment. Bodily perceptions won't do. If appraisal theorists are right to say that emotions depend on representations of this complexity, then noncognitive theories are hopeless. Psychologists try to support their dimensional appraisal theories empirically. They have subjects imagine different emotions, and then they ask them to answer questions like those in Scherer's stimulus evaluation checks. The theories correctly predict which emotions will be correlated with the various combinations of answers to these questions. This method is problematic. It reveals something about how people think about their emotions, but it does not necessarily get at the underlying processes by which emotions are elicited. It is certainly true that fear occurs in response to things that are incongruent with our goals, unexpected, power depriving, and so on. It is therefore unsurprising that we associate these features with fear. But that does not mean that we make each of these highly sophisticated judgments every time we are afraid. If someone is frightened by a seeing a spider on her shoulder, is it really plausible that she reflects on her powerlessness or her goals before the fear arises? I am skeptical. I think emotions are meaningful, but I am unconvinced by attempts to show that their meaning has the rich structural complexity postulated by prevailing psychological models.

Emotions can represent organism-environment relations, such as a danger or loss, even if they are not cognitive. I just rejected one strategy that cognitive theorists might use to argue that this is not enough. Above, I mentioned another argument that must now be addressed. In addition to representing organism-environment relations, emotions have specific intentional objects. One can be afraid that there is a prowler in the house, sad about the death of a loved one, or furious about a courtroom verdict. If emotions were judgments, one could readily explain how they come to have such intentional objects. Fear about the prowler is a thought that there is a prowler in the house coupled with a judgment that this fact presents a potential danger. I have demonstrated that perceptions of the body can represent dangers, but how can such perceptions take on specific intentional objects? A spine tingle may tell you that you are in danger, but it makes little sense to say, I spine tingle that there is a prowler in the house.

I don't think this is a deep worry. It can be addressed by introducing what might be called emotional attitudes. Emotions themselves are not judgments, or any other kind of propositional attitude, but they be featured in certain propositional attitudes. Propositional attitude can be analyzed in functional terms. They are functional roles played by mental representations of states of affairs [Fodor, 1985]. Consider a propositionally structured representation of the state of affairs that we would express by saying "there is a prowler in the house." For simplicity, we can think of this as a sentence in a language of thought [Fodor, 1975], thought it might also be an image of some kind [Prinz, 2002]. This propositional representation is not yet a propositional attitude. To become an attitude, it must play a role in information processing. A representation p is a belief if one treats p as if it were true. If you believe there is a prowler in the house, you might call the police and assert that there is a prowler in the house. Likewise, to suspect that p is to treat p as if were probably true. By analogy, we can define emotional attitudes as functional roles involving emotions. The fear that p is to suspect that p *and* to have fear as a result. Fearing is a special kind of suspecting that has an embodied appraisal as one of its constitutive effects. To be angry that p is to believe that p and to have anger as a result.

This proposal can be compared be compared to Damasio's [1994] somatic marker hypothesis, according to which we make decisions by simulating the emotional consequences of our actions. Suppose you are gambling and you need to decide whether to draw a card from a deck that has of cards that has not paid off. You fear that if you take another cased, you will lose. On Damasio's view this thought involves an embodied emotional response that arises when you imagine taking another card. Similarly, I would say this thought involves the suspicion that taking a card will lead to defeat couple with an embodied emotional response to that suspicion. This is very sketchy, of course. Developing a full analysis of the functional roles that constitute our emotional attitudes would be a difficult empirical undertaking. We do not even have a good analysis of the functional role that constitutes believing. But I see no reason for pessimism. When we say that someone has a belief that p, we mean that she has a representation of p

playing a certain role. Likewise, for a fear that p. When we say that emotions are about intentional objects, we express the fact that emotions can be caused by representations of those objects in a particular way.

This brings us to the final argument advanced against noncognitive theories: the allegation that emotions can exist without perceived bodily perturbations. I mentioned Solomon's example of long-standing love. Emotions like this pose a direct challenge to noncognitive theories of the James-Lange variety. According to James and Lange, perceived bodily changes are necessary and sufficient for emotions. Does long-standing love prove that they are wrong? Perhaps not. The foregoing account of emotional attitudes may provide the basis of a response. An emotional attitude is a functional role in virtue of which representations of states of affairs have emotional effects. One can also have an emotional attitude towards an object or person, whereby representations of that object or person cause emotional responses. The causal relationship here can be defined in dispositional terms. This is a standard feature of functional analyses. Pain causes us to say ouch, but we don't always say ouch when we are in pain. Beliefs cause us to make assertions, but we don't always assert our beliefs. Likewise, emotional attitudes dispose us to undergo bodily perturbations and to feel those perturbations, but these effects do not always arise. On this analysis, loving someone is being disposed to have an information bearing bodily response when one thinks of that person. One can love someone without the body response, but not without the disposition to have the bodily response. We would doubt the sincerity of a person who claimed to love someone, but never had a pang, tingle, or twinge when thinking of that person. If we distinguish emotions from emotional attitudes, the objection under consideration loses its teeth.

I have surveyed leading cognitive theories of the emotions, and arguments used to defend those theories over their noncognitive rivals. The arguments that I have considered are insufficient. They do not demonstrate that emotions are cognitive. Theories in the James-Lange tradition may have unappreciated resources for explaining how emotions can be meaningful. Of course, my discussion of these issues has not been exhaustive. Defenders of cognitive theories will dig in their heels. My aim here has been to give an overview of prevailing theories and a vivid impression of the core disputes dividing them. Most of these disputes are, as yet, unresolved. In section 3, I review three other disputes that are orthogonal to the debate between defenders of cognitive and noncognitive theories.

3 THREE PHILOSOPHICAL ISSUES

3.1 Unconscious Emotions

In section 2.1, I mentioned a difference between the James-Lange theory, as it is usually interpreted, and recent descendents of that theory, such as the account defended by Damasio. James and Lange sometimes equate emotions with feelings. They imply that emotions must be consciously felt. Damasio rejects this

assumption. He thinks that emotions can be unconscious. To some ears, the term "unconscious emotion" sounds like an oxymoron. Is there any principled way to decide who is right?

To answer this question, it is helpful to consider the debate about unfelt pains. Defenders of unfelt pains argue that unconscious states can behave in ways that are very much like conscious pains (e.g., [Rosenthal, 1991]). Conscious pains register injuries or maladies and they lead to analgesic behavior, such as stroking or clutching the affected body part. A mental state can register damage and promote such behaviors without being conscious. This may happen for example, when people are briefly distracted from their pains. A sudden knock on the door can momentarily distract one from a headache. It may also occur when conscious pains are overridden by other psychological process. An injured athlete may limp and grimace as if in pain while adrenalin prevents her from feeling anything unpleasant. In these cases, it is tempting to say there are unfelt pains. Unfelt states can play the ordinary pain role, minus the feeling. Likewise for emotions. An intense feeling of sadness may briefly vanish as one reads the cooking instructions on a frozen dinner box, only to return in full force a moment later. A student giving her first public address may tremble with anxiety even as her own words take her mind off of the conscious feeling. Emotions can be unconscious, the argument goes, because unconscious states can play functional roles characteristic of emotions. Unconscious states can represent organism-environment relations and register bodily changes. We talk about unconscious perception for much the same reason. One can subliminally *see* because one can register visual information without consciousness. Unfelt seeing is still seeing. If emotion is a form of bodily perception, arguments for unconscious perception help support the claim that emotions can be unconscious.

Arguments for unfelt pain, vision, and emotion find further support in certain theories of consciousness. Rosenthal [1991] defends a higher-order thought theory, according to which mental states become conscious when we have thoughts about them (see also Kriegel, this volume.). Pain becomes conscious when we form the noninferential thought: I am in pain. It order to form this thought there must be a first order state to think about; there must be a pain that we recognize as such. Unconscious emotions must exist in order to form the consciousness-conferring thoughts: I am sad, I am annoyed, and so on. One can also defend unconscious emotions by arguing that consciousness requires attention. When we don't attend to a visual stimulus, we don't consciously experience it, but, nevertheless it is visually processed, or "seen." Likewise, when we don't attend to our bodily states, we may fail to consciously experience the emotions we are having (for further discussion, see [Prinz, 2005]). In the throws of an argument, we may be so focused on the content of what we are saying that we fail to notice the irritation or rage that would be immediately obvious to anyone listening to the tone in our voices.

Opponents of unconscious emotion can defend their position in several different ways. They can argue that emotion words name feelings, rather than naming

the other psychological and behavioral processes associated with emotion. This strategy would not appeal to many emotion researchers today, because most agree that emotions are not mere feeling. Emotions necessarily carry information and dispose us to certain kinds of behavior. A second strategy is to argue that emotions cannot serve these other roles unless they are felt (compare [Clore, 1994]). This is an empirical claim, and it can probably be empirically refuted. Strahan *et al.* [2001] subliminally presented subjects with sad faces and then asked them to make choices about what music to listen to. Subjects who had seen the sad faces chose more upbeat music than control subjects, as if they were compensating for negative emotions induced by the faces. But, when asked to report their feelings, both groups came out alike. This suggests that an unconscious state can register negative information and promote coping strategies.

To argue against unconscious emotions, there is but one plausible strategy. One can insist, on conceptual grounds, that emotion words refer to states that have several essential attributes: they have meaning, they promote behaviors, *and* they are felt. A state that lacks one of these attributes is very much like an emotion, but it would be better to call it a quasi-emotion, or an affective state. It is hard to know how to refute this proposal, because it is advanced as a conceptual claim rather than an empirical claim. But there is a methodological point that one might make in response. Concepts, in general, tend to comprise non-essential features. Concepts tend to have a prototype structure. They encode features that are typical, but not necessary for category membership. No one doubts that emotions are typically felt (or, trivially, that felt emotions are the ones we typically notice). Those who think that it is a conceptual truth that emotions must be felt may be mistaking a contingent feature of their emotion concepts for a necessary one. It is like insisting that birds must have feathers. They usually do. Having feathers is part of our concept of birds. But it would be a mistake to think that you could deprive a canary of its avian status by giving it a thorough plucking.

3.2 Natural Kinds

Philosophers often assume that emotions comprise a heterogeneous category (e.g., [Ryle, 1949; Rorty, 1980; Griffiths, 1997]). The term "emotion," they say, does not represent a natural kind. If this is right, that has important implications for research. First, experimentalists who investigate the emotions must be very careful about choosing items in their studies. It would be a mistake to mix different kinds of emotions into a single analysis. Second, there can be no general theory of the emotions. Rather, there must be separate theories for each subcategory. Third, it follows that "emotion" may not be a useful theoretical construct. Perhaps it should be dropped from scientific vocabulary. This last point has a metaphysical counterpart. If "emotion" is not useful for science, then perhaps we should abandon the folk category as well. Perhaps emotions are ripe for elimination.

Griffiths [1997] has given the not-natural-kinds thesis a spirited defense. He argues that emotions fall into at least three different subcategories. Some are

affect programs: modular, automatic patterns of response that have homologues in many other species. The term affect program is borrowed from Ekman [1972], and it is intended to include the six emotions that Ekman claimed to be cross-culturally universal: anger, fear, joy, sadness, surprise, and disgust. In contrast to affect programs, other emotions are non-modular, cognitively mediated responses that usually serve to coordinate social behavior. Many of these "higher cognitive emotions" may be unique to our species. Griffiths believes that some of these emotions are species typical products of natural selection within the recent hominid line. Guilt and jealousy are examples. But other higher cognitive emotions may be socially constructed. In Japan, for example, there is an emotion called *amae*, described as a feeling of dependency, which may reflect the collectivist orientation of that culture [Doi, 1973]. Griffiths entertains the possibility that some socially informed emotions may work very differently than ordinary higher cognitive emotions. They may be "disclaimed actions." These are responses that people believe to be passive and involuntary, but actually reflect unconscious choices and strategies. If this analysis is right, higher cognitive emotions divide between those that are evolved and relatively automatic and those that are covertly willful and heavily influenced by culture.

If Griffiths is right, then emotions are not a natural kind. Affect programs, evolved cognitive emotions, and socially constructed disclaimed actions involve different mechanisms. Grouping them in a single scientific theory would be a mistake. This argument is not at all implausible, but there may be room for resistance. First of all, by Griffiths own admission, the very same emotion term is used to refer to episodes that fall into each of the subcategories that he described. Anger, for example, can be triggered automatically by a simple cue, such as seeing someone glare or getting cut off on the highway. But anger can also be set off by highly deliberative cognitive process, as when one ruminates on political policies. And anger can also be a disclaimed action. Griffiths mentions the example of staging anger in a lovers' quarrel. In this light, the divergent subcategories of emotion begin to look like different ways in which one and the same emotion can be caused. Emotions are uniform, but they can be elicited in different ways.

For another thing, emotions in all of the alleged subcategories have much in common. They all represent organism-environment relations that bear on well being; they are all felt; they all command attention; they all compel us to act; they are all negatively or positively valenced; and, I would argue, they are all typically embodied by perceived patterns of bodily change. We get pangs of jealousy and sinking feelings of guilt. In neuroimaging studies on higher cognitive emotions and Ekman's size basic emotions, the same brain areas are active. The insula, anterior cingulate, and other structures associated with bodily response are implicated for all emotions that have been investigated. Unless it can be demonstrated that these apparent commonalities across the range of emotions are illusory, there is good reason to resist Griffiths strong disunity thesis. There is good reason to think that emotions all involve the same kinds of mechanisms.

3.3 Emotions and Reasoning

I will consider one more topic of philosophical interest. Emotions are often regarded with suspicion. They are seen as disruptive forces that prevent us from seeing things clearly and making level-headed decision. Emotions are said to be like animal urges, vestiges from earlier species that could not rely on the power of reason. In a word, emotions are said to be irrational, or at least arational. They may lead to appropriate behavior, on some occasions, but they can easily lead us astray and derail reasoning. In cases where decision can be made intellectually, we are better off keeping our passions in check.

This is the picture that Damasio decries as "Descartes' Error" in his book of that title. It may not have been Descartes' error, but it is certainly a prevalent view in philosophy. Aristotle thought that emotions lead to weakness of will, Spinoza thought emotions were confused judgments, and Kant thought they were an affront to moral judgment. Damasio [1994] defends a different picture. He argues that emotions are essential for good reasoning. He bases his conclusion of studies of individuals with injuries to a part of the brain called the ventromedial prefrontal cortex. These individuals suffer from an impairment in their capacity to reason emotionally, and that impairment has an adverse affect on their capacity make good decisions. They get taken in by charlatans, their relationships fall apart, they take imprudent risks, and they behave in socially inappropriate ways. Damasio explains the connection between emotion and decision-making by appeal to the somatic marker hypothesis. We make decisions by imagining future actions. In healthy individuals, risky actions are emotionally marked. We literally experience low levels of negative emotions when we imagine doing something that carries considerable risk. We experience positive emotions when we consider actions that will be beneficial. We select actions in part by picking the ones on with the best emotional markers (Damasio calls them somatic markers, because he defends a Jamesian theory of the emotions). Thagard [2001] proposes a sophisticated model of this process. When making decisions, we do not simply weigh the emotional consequences of a given option in isolation. We consider how that option coheres with other goals and interests that are themselves emotionally marked. Emotions serve as constraints that must be satisfied when deciding what to do. The ultimate weight given to an option depends on how it fits in with affect-laden goals. If this process is impaired, bad decisions follow.

Emotions probably play other valuable roles in reasoning as well. Consider moral cognition. We may be prompted to do the good out of sympathy for others, and we may avoid the bad because we anticipate feeling guilty. Humeans argues that goodness consists in nothing more than being apt to cause approbation in us under certain conditions (e.g., knowing all the relevant facts). Kantians resist this idea, arguing that we can discover the good by pure reason. I will not try to adjudicate between Kant and Hume here. Perhaps we can discover the good without emotion. But, if Hume is right, that isn't possible. Likewise, pure reason may not be able to tell you what paintings are beautiful, what films are entertaining, or

what friends are worth having.

Emotions may also facilitate reasoning in a third way. To make decisions, we often need to sift through a lot of information that may be potentially relevant. Doing an exhaustive search through potentially relevant factors would be very time consuming, and it might even be computationally intractable. To make decisions we need an efficient method of selecting from the relevant information. Emotions may help. De Sousa [1987] argues that emotions serve to make certain things more salient, and thus lead us to attend to the most relevant information.

The point of these examples is not that emotions always lead to better decisions. They may not. In some cases, emotions lead us to do things that go against our best interests. A violent rage is rarely beneficial. Sometimes emotions lead us to do things that are neither rational nor arational. Hurtshouse [1991] gives the example of tearing up the photograph of your ex-lover. Nor is the point that emotions sometimes lead us to do rational things. No one would deny that. The point is rather that creatures like us depend on emotions for good reasoning. Without them we would not be able to act on the basis of future consequences, make evaluative judgments, or solve certain versions of the frame problem. Other imaginable decision-making systems may be able to do some of these things without the aid of emotions, but we, as a matter of fact, cannot. This is a controversial claim, of course. If it is correct, emotions are extremely important for reasoning. Models of both theoretical and practical inference would benefit from taking emotions seriously.

BIBLIOGRAPHY

[Adolphs *et al.*, 2000] R. Adolphs, H. Damasio, D. Tranel, G. Cooper and A. R. Damasio. A role for somatosensory cortices in the visual recognition of emotion as revealed by three-dimensional lesion mapping. *Journal of Neuroscience,* 20, 2683-90, 2000.

[Alloy and Abrahamson, 1988] L. B. Alloy and L. Y. Abrahamson. Depressive realism. In L. B. Alloy (Ed.), *Cognitive processes in depression* (pp. 441-485). New York, NY: Guilford Press, 1988.

[Alloy and Abrahamson, 1979] L. B. Alloy and L. Y. Abramson. Judgment of contingency in depressed and nondepressed students: Sadder but wiser? *Journal of Experimental Psychology*, 108, 441-485, 1979.

[Armon-Jones, 1989] C. Armon-Jones. *Varieties of affect.* Toronto: University of Toronto Press, 1989.

[Arnold, 1960] M. B. Arnold. *Emotion and Personality*, New York, NY: Columbia University Press, 1960.

[Cacioppo *et al.*, 2000] J. Cacioppo, G. Bernston, J. Larson, K. Poehlmann and T. Ito. The psychophysiology of emotion. In M. Lewis and J. Haviland-Jones (Eds.), *Handbook of emotions*, *Second edition* (173-191). New York, NY: Guilford Press, 2000.

[Chwalisz *et al.*, 1988] K. Chwalisz, E. Diener and D. Gallagher. Autonomic arousal feedback and emotional experience: Evidence from the spinal cord injured. *Journal of Personality and Social Psychology, 54,* 820-828, 1988.

[Clore, 1994] G. Clore. Why emotions are never unconscious. In P. Ekman and R. J. Davidson (Eds.) *The nature of emotion: Fundamental questions* (pp. 285-290). New York, NY: Oxford University Press, 1994.

[Critchley *et al.*, 2000] H. D. Critchley, D. R. Cornfield, M. P. Chandler, C. J. Mathias and R. J. Dolan. Cerebral correlates of autonomic cardiovascular arousal: a functional neuroimaging investigation in humans. Journal of Physiology, 523, 259-270, 2000.

[Damasio, 1994] A. R. Damasio. *Descartes' error: Emotion reason and the human brain*. New York, NY: Gossett/Putnam, 1994.
[Damasio, 1999] A. R. Damasio. *The feeling of what happens: Body and emotion in the making of consciousness*. New York, NY: Harcourt Brace and Company, 1999.
[De Sousa, 1987] R. De Sousa. *The rationality of emotions*. Cambridge, MA.: MIT Press, 1987.
[Doi, 1973] T. Doi. *The anatomy of dependence*. Tokyo: Kodansha International, 1973.
[Dretske, 1986] F. Dretske. Misrepresentation. In R. Bogdan (ed) *Belief: Form, content and function*, pp. 17–36. Oxford: Oxford University Press, 1986.
[Ekman, 1972] P. Ekman. *Emotions in the human face*. New York, NY: Pergamon Press, 1972.
[Fodor, 1975] J. A. Fodor. *The Language of Thought*. New York, NY: Cromwell, 1975.
[Fodor, 1985] J. A. Fodor. Fodor's Guide to Mental Representation: The Intelligent Auntie's Vade-Mecum. *Mind*, 94, 76-100, 1985.
[Greenspan, 1988] P. Greenspan. *Emotions and reasons*. New York, NY: Routledge, 1988.
[Griffiths, 1997] P. E. Griffiths. *What emotions really are*. Chicago, IL: University of Chicago Press, 1997.
[Hohmann, 1966] G. W. Hohmann. Some effects of spinal cord lesions on experienced emotional feelings. *Psychophysiology, 3,* 143-156, 1966.
[Hursthouse, 1991] R. Hursthouse. Arational Actions. *Journal of Philosophy*, 88, 57-68, 1991.
[Isen, 2000] A. M. Isen. Positive affect and decision making. In M. Lewis and J. Haviland-Jones (Eds.), *Handbook of emotion, Second edition* (pp. 417-435). New York, NY: Guilford, 2000.
[James, 1884] W. James. What is an Emotion? *Mind*, 9, 188-205, 1884.
[Kuniecki *et al.*, 2003] M. Kuniecki, A. Urbanik, B. Sobiecka, J. Kozub and M. Binder. Central control of heart rate changes during visual affective processing as revealed by fMRI. *Acta Neurobioligiae Experimentalis*, 63, 39-48, 2003.
[Lange, 1885/1922] C. G. Lange. *Om sindsbevaegelser: et psyko-fysiologisk studie*. Kjbenhavn: Jacob Lunds, 1885. Reprinted in *The emotions*, C. G. Lange and W. James (Eds.), I. A. Haupt (trans.) Baltimore, Williams and Wilkins Company 1922.
[Lazarus, 1991] R. S. Lazarus. *Emotion and adaptation.* New York, NY: Oxford University Press, 1991.
[LeDoux, 1996] J. E. LeDoux. *The emotional brain*. New York, NY: Simon and Schuster, 1996.
[Levenson *et al.*, 1990] R. W. Levenson, P. Ekman and W. V. Friesen. Voluntary facial action generates emotion-specific autonomic nervous system activity. *Psychophysiology,* 27, 363-384, 1990.
[Mogg and Bradley, 1999] K. Mogg and B. P. Bradley. Orienting of attention to threatening facial expressions presented under conditions of restricted awareness. *Cognition and Emotion*, 13, 713–40, 1999.
[Nussbaum, 2001] M. C. Nussbaum. *Upheavals of thought: The intelligence of the emotions,* Cambridge: Cambridge University Press, 2001.
[Oatley and Johnson-Laird, 1987] K. Oatley and P. N. Johnson-Laird. Towards a cognitive theory of emotions. *Cognition and Emotion*, 1, 29-50, 1987.
[Philippot *et al.*, 2002] P. Philippot, C. Chapelle and S. Blairy. Respiratory feedback in the generation of emotion. *Cognition & Emotion*, 16, 605-627, 2002.
[Pitcher, 1965] G. Pitcher. Emotion. *Mind,* 74, 324-346, 1965.
[Prinz, 2004] J. J. Prinz. *Gut reactions: A perceptual theory of emotion*. New York, NY: Oxford University Press, 2004.
[Prinz, 2005] J. J. Prinz. Emotions, embodiment, and awareness. In L. Feldman-Barrett, P. Niedenthal, and P. Winkielman (Eds.), *Emotion and consciousness* (pp. 363-383). New York, NY: Guilford, 2005.
[Reisenzein, 1983] R. Reisenzein. The Schachter theory of emotion: Two decades later. *Psychological Bulletin, 94*, 239-264, 1983.
[Rorty, 1980] A. O. Rorty. Introduction. In A. O. Rorty (Ed.), *Explaining emotions* (pp. 1-8). Berkeley, CA: University of California Press, 1980.
[Roseman *et al.*, 1990] I. J. roseman, M. S. Spindel, and P. E. Jose. Appraisals of emotion-eliciting events: Testing a theory of discrete emotions. *Journal of Personality and Social Psychology*, 59, 899-915, 1990.
[Rosenthal, 1991] D. Rosenthal. The independence of consciousness and sensory quality. In E. Villanueva (Ed.), *Consciousness: Philosophical issues*, 1, (pp. 15-36). Atascadero, CA: Ridgeview Publishing Company, 1991.
[Ryle, 1949] G. Ryle. *The concept of mind*. New York, NY: Barnes and Noble Press, 1949.

[Schachter and Singer, 1962] S. Schachter and C. Singer. Cognitive, social, and physiological determinants of emotional state. *Psychological Review*, 69, 379-399, 1962

[Scherer, 1997] K. R. Scherer. Profiles of emotion-antecedent appraisal: testing theoretical predictions across cultures. *Cognition and Emotion*, 11, 113-150, 1997.

[Smith and Lazarus, 1993] C. A. Smith and R. S. Lazarus. Appraisal components, core relational themes, and the emotions. *Cognition and Emotion*, 7, 233-269, 1993.

[Soloman, 1976] R. C. Solomon. *The passions*. New York, NY: Doubleday, 1976.

[Spinoza, 1677/1994] B. de Spinoza. *Ethics*. In *A Spinoza reader*. E. Curley (trans. and ed.). Princeton, NJ: Princeton University Press, 1677/1994.

[Strack *et al.*, 1988] F. Strack, L. L. Martin and S. Stepper. Inhibiting and facilitating conditions of facial expressions: A nonobtrusive test of the facial feedback hypothesis. *Journal of Personality and Social Psychology, 54*, 768-777, 1988.

[Strahan *et al.*, 2002] E. J. Strahan, S. J. Spencer and M. P. Zanna. Subliminal priming and persuasion: Striking while the iron is hot. Journal of Experimental Social Psychology, 38, 556-568, 2002.

[Thagard, 2001] P. Thagard. How to make decisions: Coherence, emotion, and practical inference. In E. Millgram (Ed.), *Varieties of practical inference* (pp. 355-371). Cambridge, MA: MIT Press, 2001.

[Wittgenstein, 1953] L. Wittgenstein. *Philosophical investigations*. G. E. M. Anscombe (Trans.). Oxford: Blackwell, 1953.

[Zajonc, 1984] R. B. Zajonc. On the primacy of affect. *American Psychologist*, 39, 117-123, 1984.

SIMULATION

Alvin Goldman and Kelby Mason

1 INTRODUCTION

Human beings spend a lot of time thinking about thoughts, feelings and other assorted mental states, both their own and those of others. I try to figure out what Albertine really thinks of Oscar, how much a certain gift would please the intended recipient, how I would feel in the morning if I had another drink. We commonly arrive at beliefs about such mental states, and thereby make attributions — retrospective, concurrent or predictive attributions — to the persons in question, the "targets" of attribution. The psychological capacity to perform this kind of task is variously referred to as *mindreading*, *folk psychology*, *theory of mind* or *mentalizing*. This capacity makes an appearance early in life, across cultures and without explicit cultural instruction. Most of our everyday social interaction and communication seems to depend on our beliefs and expectations about the mental states of other people, and therefore on mindreading.

The scientific study of mindreading began with a paper in primatology. Based on the impressive performance of chimpanzees on various tasks, David Premack and Guy Woodruff (1978) argued that they had humanlike mindreading abilities (see Tomasello and Call 1997, Heyes 1998, and Povinelli et al. 2000 for more sceptical reviews of mindreading in non-human primates). The tasks they used, however, could not rule out the possibility that chimpanzees were merely exploiting correlations between situations and targets' behavior (e.g. when food is around, others will eat it), rather than actually thinking about targets' mental states (e.g. that others want the food and will therefore eat it). This prompted several philosophers (Dennett 1978, Harman 1978) to suggest a more rigorous test for mindreading, the *false-belief* task (described below). The false-belief task was introduced to developmental psychology in an influential paper by Heinz Wimmer and Josef Perner (1983), who tested for infants' mindreading abilities. Since then, mindreading has become a major area for research in philosophy and the cognitive sciences, including developmental psychology, social psychology and neuroscience.

One theory of mindreading is that it involves mental *simulation*. Its basic idea is that in predicting the target's mental states, we try to simulate or reproduce in our minds the same state, or sequence of states, that they are in. We put ourselves in their position, and see what we would believe, feel or do. There are various ways to elaborate this basic idea, and there are also alternative accounts of mindreading. To give some idea of the context in which simulation theory emerged, we sketch

Handbook of the Philosophy of Science. Philosophy of Psychology and Cognitive Science
Volume editor: Paul Thagard
General editors: Dov M. Gabbay, Paul Thagard and John Woods
© 2007 Elsevier B.V. All rights reserved

these alternative accounts below. But the main focus of this chapter is simulation, and so we give more attention to versions of simulation theory, detail the evidence for the importance of simulation, and discuss some of the key conceptual issues for simulation theory.

A comprehensive theory of human mindreading should address at least the following six questions:

1. How do we mindread other people? What *processes* do we use to arrive at mental state attributions?

2. How do we mindread ourselves? (Since the processes of first- and third-person attribution may differ, this is a separate question from (1)).

3. The *developmental* question—how and when do we acquire the capacity for mindreading?

4. The *architectural* question—how is this capacity implemented in the brain, and how is it connected with other capacities?

5. The *conceptual* question—what is the content of the mental state concepts used in mindreading?

6. The *comparative/evolutionary* question—what are the mindreading capacities of other species, and how did the human capacity evolve?

Although there are connections between these questions, especially (1) and (2), they are still largely independent. Simulation theory is primarily an answer to (1), and only secondarily to the others. Similarly, the chief alternative to simulation theory, *theory theory*, is in the first instance a theory about the process of third-person attribution. In the next section we outline the basics of the theory theory of mindreading, as well as two other accounts that have been suggested as alternatives to simulation theory, namely *modularism* and *rationality theory*.

2 ALTERNATIVES TO SIMULATION

Theory theory

According to theory theory, our mindreading capacity is implemented by an intuitive theory of mind, a body of knowledge or belief about the causal relations between mental states on the one hand and behavior, environment and other mental states on the other. An intuitive, or folk, theory of mind is supposed to contain a lot of generalizations like the following:

1. If someone wants x, and believes that y is a means to x, then, other things being equal, they will do y.

2. If someone believes p and *if* p *then* q, then, other things being equal, they will believe q.

3. If x is in someone's field of vision and they are attending to the scene, then, other things being equal, they will perceive x.

In mindreading, we use these generalizations and our data about the target's circumstances or behavior to infer some of the target's mental states. Once we have done this, we can then make further predictions or retrodictions about their other mental states. Note that neither the theory of mind nor the process of inference is assumed to be accessible to consciousness. Thus we can be mostly unaware of either its principles or any overt deductive or inductive inferences, while still relying on them at a subpersonal level when we mindread. In addition, the generalizations of theory of mind would need to be much richer and more complicated than the skeletal folk platitudes (1)–(3), but theory-theorists have done limited work on formulating more realistic principles.

In the early days of research on mindreading, theory theory was the default account of mindreading in psychology (and to some extent it still is). This is partly due to the prevalence in philosophy of mind of a position called *functionalism* (Lewis 1972; Block 1978; Shoemaker 1984). According to one brand of functionalism, common sense functionalism, mental-state concepts are understood in terms of their role (or function) in the overall cognitive economy. What we understand by a mental state such as pain or belief, say, are the causal links between that state and behavior, environment and other mental states, links that are specified by a common sense psychological theory. If mental states are understood by their functional roles in a common sense psychological theory, then it is natural to infer that when ordinary people are thinking about mental states—as in mindreading—they apply such a theory.

This functionalist notion of "theory" is one of several reasons for calling the body of our folk-psychological knowledge a theory. Just as in many scientific theories, folk-psychological concepts like belief and desire are interdefined by a cohesive set of causal principles. Moreover, setting aside the controversial status of introspection, mental states are not directly observable, and in other scientific domains such non-observable entities (e.g. atoms, genes) are often called "theoretical". Finally, if mindreading involves making inferences from observable data and general principles, then it resembles scientific processes of prediction. Because it holds that mindreading requires a lot of propositional knowledge like (1)–(3) plus inferences from this knowledge, theory theory can also be called an *information-rich* and *inferentialist* approach.

In its most general form, theory theory is a theory about the process of mindreading. As such, it is compatible with a variety of views about the development of mindreading, its cognitive architecture, and the relation between first- and third-person mindreading. The purest form of theory theory, however, takes a definite stance on all three questions. This is the *child-scientist* theory, advocated by Alison Gopnik, Henry Wellman, and Andrew Meltzoff (Gopnik and

Wellman 1992, 1994; Gopnik and Meltzoff 1997). According to the child-scientist theory, the resemblance between our mindreading capacity and scientific theories is not limited to the theory-like representation of folk-psychological knowledge. The acquisition of that knowledge is also like the acquisition of scientific knowledge, namely a gradual process of theorizing, making generalizations based on the available evidence, revising our theories to handle incongruent evidence etc. Thus the child-scientist theory has an explicit theory of development (which will be assessed in the section on false-belief tasks, below), according to which the child works out a theory of mind just as the scientist works out scientific theories. It also has clear implications for cognitive architecture: the acquisition process uses a domain-general mechanism for theorizing, and the adult mindreading capacity uses domain-general mechanisms for inference. Finally, Gopnik has claimed that an advantage of the child-scientist theory is that it offers a unified account of first- and third-person mindreading, which are supposed to work by the same inferential processes.

Modularism

Another view that has been proposed as an explanatory framework for mindreading is that mindreading is implemented by an innately specified module, which Alan Leslie calls the *Theory of Mind Mechanism* or *ToMM* (Leslie 1994). Modularism seems to share with theory theory an inferentialist account of attribution, so it may be a species of theory theory. Alternatively, it may be taxonomized as a distinct approach. It is not clear whether mindreading is supposed to have all the strictly Fodorian features of modularity (Fodor 1983), but modularists typically claim that mindreading has at least the more important ones.

There are several converging lines of evidence that mindreading is at least somewhat modular. Mindreading follows a typical developmental pattern through typical developmental stages. It is also dissociable from other cognitive capacities, so that people can be impaired in mindreading while functioning normally in other areas, and vice versa. This has led researchers such as Leslie (1987; Scholl and Leslie 1999) and Simon Baron-Cohen (1995) to claim that the development of mindreading proceeds on a more-or-less innately specified pathway.

Modularism is thus incompatible with the child-scientist form of theory theory, which claims that the acquisition of mentalizing itself involves theorizing and that mindreading uses domain-general theoretical processes. But theory theory is not wedded to either this anti-nativism or domain-generality. If theory of mind is conceived of as *core knowledge* (Carey and Spelke 1991; Hirschfeld and Gelman 1994; Botterill and Carruthers 1999), we would expect mindreading to be both modular and information-rich. Since modularism is mainly a theory of architecture and development, it is largely noncommittal on the processing question and therefore cannot decide between theory theory and simulation theory.

Rationality theory

Another theory of mindreading, especially influential in philosophy, is the *rationality theory*, chiefly associated with Donald Davidson (1984) and Daniel Dennett (1987). Following some ideas of W. V. Quine (1960) on *radical translation*, Dennett and Davidson each proposed that there are fundamental constraints on the sorts of interpretations we give to people's utterances and, by extension, on the sorts of mental states we attribute to them. These constraints are seen as deriving from a fundamental *principle of charity* or *assumption of rationality*. For Dennett, we are constrained to interpret people as being more or less *rational*, while for Davidson, we are constrained to interpret their beliefs as being mostly *true*.

Dennett's basic idea is this: when we attribute a mental state to a target, we consider what mental state would be rational for the target to have, given what she is doing, her perceptual surround, what else she believes, and so on. Similarly, when we predict a target's behavior, we consider what behavior would be rational for the target in that situation, with those beliefs and desires, etc. Davidson gives a similar story with "truth" in place of "what would be rational", at least for belief-attribution—that is, what we attribute to a target is mostly beliefs that are true (by our lights). When our interpretations violate these constraints, claim Dennett and Davidson, we conclude that the interpretations must be wrong rather than that the target is massively irrational or deluded.

The rationality theory has not attracted much empirical research with the exception of the work of Csibra, Gergely and their colleagues (Csibra et al 1999, Gergely et al 1995). Their experiments have shown that very young infants (as early as 9 months) interpret simple animated shapes as agents, and expect them to act efficiently. In one study, for instance, infants were shown one figure "jumping over" an obstacle to reach its destination. When shown another scene in which the obstacle was gone, they now expected the figure to go straight to its destination and were surprised if the figure traced the same "jumping" path as before. Such results suggest that mindreaders make some assumptions about efficiency of action, if not precisely rationality.

Like modularism, rationality theory is not a complete alternative to either theory theory or simulation theory, because it does not provide an account of the process of mindreading. The rational thing to do or think does not simply present itself when we mindread; rather, it takes a lot of cognitive work to figure out what it is. Working out the rational thing to do/think must either rely on a theory of what is rational (so rationality theory becomes theory theory) or on simulating what I would do and calling that "rational" (so rationality theory becomes simulation theory). If there are principles of rationality in mindreading, they can only provide constraints on mindreading, rather than an explanation of how mindreading works. Moreover, they are weak constraints at best, since we do not expect people to be perfectly consistent nor to believe everything implied by their beliefs (since any given belief has an infinite number of implications, we had better not expect this!).

3 SIMULATION THEORY

The central idea of simulation was foreshadowed by thinkers like Adam Smith (1759/1976), David Hume (1739/1978), Immanuel Kant (1781/1953), Friedrich Nietzsche (1881/1982) and W. V. Quine (1960, 1990), who discussed in one form or another the use of *empathy*, *imagination* or *projection* in mindreading. Simulation can be considered a kind of empathy, if empathy is understood as feeling one's way into another's position by identifying with them, or imaginatively projecting oneself into their position. Alternatively, simulation can be understood as automatically matching or replicating another person's mental state, without any deliberate imagination or projection. Simulation also has conceptual ties to psychological research on *perspective-* and *role-taking* and *egocentric attribution*.

Simulation was first seriously proposed as a thoroughgoing account of mindreading in the 1980s. Reacting against the orthodoxy of functionalism in philosophy of mind and of theory theory in the emerging psychological research on mindreading, philosophers Jane Heal (1986) and Robert Gordon (1986) independently argued that mindreading might largely depend on simulation rather than on a theory of mind. Simulation has since been further developed by Gordon (1992, 1995, 1996), Heal (1996, 1998, 2003), Gregory Currie (1995, Currie and Ravenscroft 2002), Vittorio Gallese (2001, 2003), Alvin Goldman (1989, 1992a, 1992b, 2000, 2006; Gallese and Goldman 1998; Goldman and Sripada 2005), Paul Harris (1991, 1992), and others.

A large part of the motivation for simulation theory is the suspicion that theory theory is *too* information rich—that it attributes too much knowledge and sophistication to mindreaders, even if the knowledge is only tacit or unconscious. Instead of supposing that mindreaders possess a special information-rich body of folk-psychological knowledge, simulation theory suggests that they are *information-poor* but engage a special skill, namely constructing pretend, simulated or imaginary mental states and engaging in simulated processing. Because mindreaders and their targets are (usually) both members of the same species, their mental states and cognitive systems are similar in many respects. On the simple—and mostly correct—assumption that their targets are "like me" (Meltzoff and Brooks 2001), mindreaders can use themselves as a model or analogue for the target. There is thus no need for the mindreader to store a large body of folk-psychological knowledge.

An analogy is often drawn between this kind of mental simulation and the way we simulate the behavior of aircraft in a wind tunnel. Suppose that we want to know how a new aircraft design will perform under various real-world conditions of turbulence. One way to determine this, which corresponds to mindreading via theory of mind, would be to calculate the behavior from our knowledge of the aircraft specifications and laws of physics. Another way, corresponding to simulation, would be to construct a model of the aircraft and place it in a wind tunnel, which is designed to match real conditions in the relevant aspects. While the first method is *knowledge-driven*, the second is *process-driven* (Goldman 1989).

Of course in real life we would use both theory and simulation to predict how the aircraft will perform. Perhaps "theory", in some broad sense, is needed both before and after we run the simulation; before, when we design the model aircraft, and after, when we use the results of the test-run to predict the behavior of the real aircraft. Similarly, in mindreading there might be theory involved in deciding which pretend inputs to construct, and in transferring our simulated state or behavior from ourselves to the target. Even if it is buffered by theory, simulation still plays a crucial role in such a model. In fact, there is good reason to think that mindreading involves a hybrid of theory and simulation, as we will discuss below.

There is one important difference between using a wind tunnel for aerodynamic prediction and using our own cognitive systems for predicting behavior and inference. In the mindreading case, the modeling systems are the very systems that also drive our own behavior and judgments; so they had better not use simulated inputs the same way they use normal inputs, to produce behavior and judgment. Mindreaders do not carry out the behavior they predict of their targets, nor do they come to share all the mental states they attribute to them. Thus it is crucial that in simulation the systems used to model the target are taken *off-line* and their usual effects inhibited. This is highlighted with the help of the following diagrams (adapted from Goldman, 2006).

Figures 1 and 2 graphically bring out the central idea of offline simulation. In these figures, squares represent desires, ovals represent beliefs, hexagons represent decisions, and the circle represents a decision-making system. Figure 1 represents an actor who desires to bring about a certain outcome g and believes that m *is a means to* g. Here, suppose that g is "I eat some chocolate", and m is "I look in the cupboard". This desire and belief are fed into the decision-making system, which produces a decision to do m, that is, look in the cupboard.

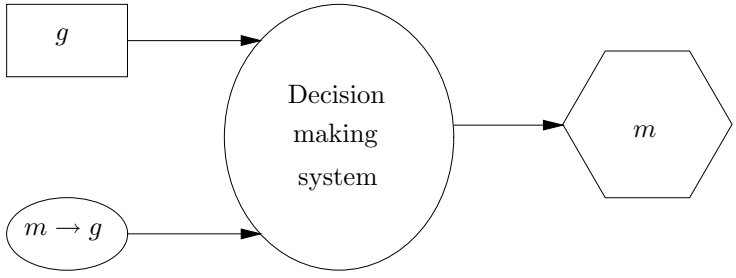

Figure 1. Standard decision to do m

Figure 2 represents a mindreader trying to predict the decision of a target (T) via simulation. The mindreader does not desire g nor believe that m *is a means to* g, but she thinks that T does have this desire and belief. In the first stage of simulation, the mindreader constructs a pretend belief and pretend desire to match the states of the target (pretense is represented by shading, and imputation to T is represented by labeling states with '**T'**). In the second stage, she feeds

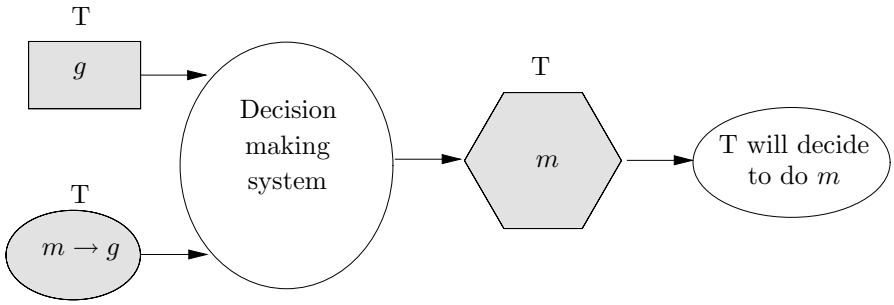

Figure 2. Simulating a decision to do m

these simulated states into her decision-making system and the system produces a simulated decision to do m. Finally, the mindreader transfers this pretend decision to the target and forms a genuine (not merely simulated) belief that *T will decide to do m*.

4 DEVELOPMENT: FALSE-BELIEF TASKS

A large body of developmental research indicates that children's mindreading skills follow a characteristic developmental pattern. There is, however, substantial disagreement about whether this pattern reflects changes in children's mindreading competence, or merely improvements in performance. A classic experimental paradigm is the *location-change* version of the false-belief task, introduced by Wimmer and Perner (1983), and its basic structure is as follows. The child is shown a puppet, Maxi, who places some chocolate in a cupboard. Maxi then leaves the room and, while he is gone, the chocolate is moved to another location. The child is asked where Maxi will look for the chocolate when he returns (where he thinks the chocolate is). The chief finding of Wimmer and Perner was that 5-6 year-olds reliably passed this task, predicting that Maxi would look in the old location, whereas 3-4 year-olds failed, predicting that Maxi would look in the actual location of the chocolate. Later studies have more or less robustly replicated this finding, across a variety of manipulations in the task, demonstrating a significant change in mindreading abilities between ages three and four (reviewed in Wellman et al. 2001).

Other studies involving different tasks also pointed to a similar change between these ages. In the *appearance-reality* task, children are shown an object with a misleading appearance (e.g. a *Hollywood Rock*, a sponge that looks like a rock) and asked what they think it is ("a rock"). They are then told what the object really is and asked what it looks like. Younger children fail this task as well, saying that the sponge actually looks like a sponge (Flavell et al. 1986). Yet another task in this area is the *deceptive-container* task, a different type of false-belief task (Perner

et al. 1987). Children are shown a box of candy and asked what they think is inside ("candy"). After being shown that the box in fact contains pencils, they are asked again what they had originally thought was in the box. Once again, 3-year-olds reliably fail this task, claiming to have correctly thought pencils were inside instead of acknowledging the false thought that candy was inside (Gopnik and Astington 1988).

These and similar findings are taken by many as evidence that three year-olds lack the concept of a *mis*representation, and hence the concept of a false belief. While they do have some understanding of mental states, they don't yet have a *representational* understanding, according to which mental states can represent things otherwise than how they really are (Perner 1991; Gopnik 1993). This group of developmentalists argue that children undergo a process of *conceptual change* between the ages of three and four, during which they acquire a genuinely *representational* theory of mind.

An alternative proposal, first offered by Fodor (1992) and Leslie and German (1995), is that 3-year-olds and 4-year-olds have the same underlying competence. What changes between these ages is rather something about their performance, with 4-year-olds acquiring the processing power to give the right answers. This performance-based explanation is supported by studies showing that various manipulations can help 3-year-olds pass the tests. For instance, when three year-olds are given a memory aid in the false-container task, they can recall their original (false) prediction (Mitchell and Lacohee 1991). They can also give the right false-belief answer when the reality is made less salient, for instance if they are merely told where the chocolate really is, rather than seeing it for themselves (Zaitchik 1991). This suggests that young children have a problem inhibiting what is in another sense the "correct" answer (i.e. Maxi *should* look in the new location if he wants to find the chocolate).

In fact, there is independent evidence against the conceptual change theory. Karen Bartsch and Henry Wellman (1995) performed a detailed analysis of young children's conversations about mental states, based on transcripts in the Child Language Data Exchange System (CHILDES). There and elsewhere (Bartsch and Wellman 1989), they showed that even 3-year-olds have a concept of false belief, although there may be a lag between gaining that concept and deploying it in explanation and prediction.

Better support for a performance-based explanation over a conceptual-deficit explanation of false-belief task skill would be a specification of the nature of the psychological change between three and four that accounts for improved mindreading skill. An excellent candidate has now emerged, viz., growth in *inhibitory control* capacity. Inhibitory control is a specific variant of a wider construct: executive functioning. Executive functioning subsumes processes that monitor and control thought and action, including self-regulation, planning, behavior organization, cognitive flexibility, and response inhibition. Inhibitory control is an executive ability that enables someone to override "prepotent" tendencies, i.e., dominant or habitual tendencies. In the present case, the dominant, habitual, or natural tendency

is to reference reality as one knows, or believes, it to be. A false-belief task requires an attributor to override this natural tendency, to say that Maxi will look for the chocolate somewhere different from where the attributor knows it to be. The proposal is that 3-year-olds have trouble overriding their prepotent tendency because their inhibitory control mechanism has not yet (sufficiently) matured. By four years of age, the inhibitory control mechanism has strengthened and they can do the requisite overriding. Simulation theory might adapt the inhibitory control story by saying that weak inhibitory control capacity produces poor simulation, i.e., failure to imaginatively override one's current belief about reality.

Evidence in support of the inhibitory control explanation of false-belief task changes has been presented by Carlson and Moses (2001). Prior to their work it had been independently established that important developmental changes occur in inhibitory control during just the period that the indicated changes in mindreading take place, between 3 and 6 (Diamond and Taylor 1996; Frye et al. 1995). Also it was known that autistic individuals show impairment on classic executive functioning tasks such as the Wisconsin Card Sort task and Tower of Hanoi task (Hughes and Russell 1993; Ozonoff et al. 1991), as well as on the relevant mentalizing tasks. The specifics of the Carlson and Moses study are as follows.

Children were tested on a battery of four mindreading tasks plus a battery of 10 inhibitory control tasks. The mindreading tasks included the location-change and deceptive-container false belief tasks. The 10 inhibitory control tasks all required children to respond contrary to a prepotent tendency. In their day/night task, for example, the prepotent response was to say "day" for a picture of the sun and "night" for a picture of the moon. Subjects were instead required to say the opposite of what the picture shows. In their grass/snow task, children were required to point to the color that is opposite to its associate, i.e., green for "snow" and white for "grass". Almost all the comparisons in performance between the ten inhibitory control measures and the four mindreading measures showed significant correlations. Moreover, every inhibitory control measure was significantly related to the mindreading battery (i.e., overall) scores; conversely, every mindreading measure was significantly correlated with the inhibitory control battery. The correlation between the two batteries themselves was especially high, viz., 0.66. Although the correlational nature of the Carlson and Moses study does not directly resolve the direction of causality, Hughes (1998) found that executive functioning at age three predicts mindreading performance one year later better than mindreading predicts executive functioning across that age span. This favors a direction of causality from executive functioning to mindreading rather than the converse.

In sum, the inhibitory control approach is a compelling account of false-belief task changes that invokes purely performance factors rather than the conceptual-deficit factor, and thereby undercuts the explanation favored by child-scientist theorists. It is also compatible with simulation theory. Continuing support for this approach derives from studies by Leslie and collaborators (Leslie and Polizzi 1992; Leslie et al., 2005). Contrary to what was widely believed in the early

1990s, therefore, the developmental evidence gives no support to theory theory over simulation as an account of mindreading.

5 SIMULATION AND MIRROR SYSTEMS

There are other domains than mindreading in which simulation plausibly plays a role, and that demonstrate how simulation might work in mindreading. Neuroscientists in the laboratory of Giacomo Rizzolatti discovered an unexpected class of neurons in the premotor cortex of macaque monkeys. This area features neurons that mostly subserve the preparation or planning of actions, with different neurons coding for different types of actions, such as grasping, holding, or tearing an object. Further testing of the monkeys, using single cell recording, revealed the surprising fact that a subclass of these neurons (in Brodmann area 44) fire at close to maximal rates not only when the animal executes its associated type of action but also when it merely *observes* another monkey or an experimenter execute the same action with respect to a goal object. This matching or mirroring of observation and execution was quite unexpected and led the investigators to label the neurons *mirror neurons* (Gallese et al., 1996; Rizzolatti et al., 1996). Although the same group of mirror neurons fire in both the observer and the actor, the observer does not actually execute the action; apparently, the plan or goal that is activated in the observer is neurally inhibited.

By using a variety of methodologies, a somewhat similar action mirroring system has been established in humans (e.g., Fadiga et al., 1995). In the human case, it was also found that observing an action, such as grasping an object or tracing a figure in the air, reliably produces electromyographically detectable activation in the muscle groups appropriate for the observed action, although no overt action is performed. One might say that the plan of action is taken "off-line". Confirmation of this comes from clinical evidence of so-called "imitation behavior". Patients with prefrontal lesions compulsively imitate gestures or complex actions they see (Lhermitte et al., 1986). This behavior is explained as an impairment of the inhibitory control normally governing motor plans. From this it may be inferred that unimpaired humans, when observing others perform actions, regularly generate plans to do the same actions, which are then inhibited.

If mirror-neuron activity is interpreted as described above, it's a matter of a goal state in an actor being replicated, or simulated, by a corresponding goal state in an observer. This led Gallese and Goldman (1998) to see it as a possible component of, or a precursor of, a rudimentary exercise in simulational mindreading. At a minimum, an observer's undergoing a goal-state that mirrors, or resonates with, a corresponding goal-state in a target provides the potential for such mindreading. It is not established that monkeys ever engage in mindreading; nor is it established that human observers who undergo mirror-matching events use those events to attribute corresponding goal-states to their targets. However, mirror-matching activity is appropriate simulational activity that *could* provide the basis for successful simulational mindreading.

Research on mirroring, or resonance, phenomena has produced many other impressive findings of a comparable sort, clear evidence that minds do a lot of automatic simulation of other people's actions (see Gallese 2001, 2003 for reviews). Rizzolatti's group in Parma extended their work with an fMRI study of human subjects who were shown actions made not only with the hand but with the mouth or a foot, e.g., biting an apple or kicking a ball (Buccino et al., 2001). There was matching neuronal activity in each case. In other words, when observing an action, our motor system becomes active *as if* we were executing that very same action being observed. Indeed, action observation determined a somatotopically-organized activation of premotor cortex, a pattern similar to that of the classical motor cortex somatotopy. Such an organization was also found in the posterior parietal lobe (Buccino et al., 2001).

Subsequent experiments show that activation of the premotor cortex (in monkeys) does not depend on actually seeing the final part of the observed action. Even when the final part is hidden, and hence can only be inferred, there is also mirror neuron activation (Umilta et al., 2001). Another study shows that the observation needn't be visual; even sounds related to a specific type of action can generate mirror-matching activation (Kohler et al., 2002).

Another striking domain of mirroring or resonance is touch. When we watch a movie scene in which a tarantula crawls on James Bond's chest, we ourselves shiver, as if the spider were crawling on our own chest. This made Keysers et al. (2004) wonder whether there is such a phenomenon as "tactile empathy". If a mirroring mechanism applies to the sight of touch, then watching someone being touched should activate the same area(s) as getting touched oneself. This was confirmed. The secondary (but not the primary) somatosensory cortex was activated both when participants were touched and when they observed someone else getting touched by objects (Keysers et al. 2004).

There are also some domains of mirroring that are relevant to mindreading. Mirror neurons for pain were initially discovered quite accidentally when a neurological patient was being prepared for cingulotomy. It was observed that a single cell in the patient responded both when he was stimulated with a pin-prick and when he observed the same stimulation applied to the medical examiner (Hutchison et al. 1999). In a functional imaging study, Singer et al. (2004) confirmed the existence of empathy for pain. Volunteers experienced a painful stimulus and also observed a signal indicating that their loved one—present in the same room—was receiving a similar painful stimulus. Overlapping regions were activated under both conditions. Specifically, increased activation of the anterior insula and anterior cingulate cortex was common both to the experience of pain (as contrasted with no-pain) and the "observation" of pain (as contrasted with no-pain) in another.

Finally, an fMRI study of the experience and observation of disgust reveals that there is mirroring, or resonance, in the domain of emotion as well. Wicker et al. (2003) performed an experiment in which participants inhaled disgusting or pleasant odorants through a mask on their nose and mouth and, separately, viewed

movies of individuals smelling the contents of a glass, which were either disgusting, pleasant, or neutral. The same brain regions were preferentially activated during the experience of inhaling disgusting odorants and during the observation of disgust facial expressions.

6 SIMULATION AND IMAGERY

Gregory Currie and Ian Ravenscroft have argued that imagery is a form of simulation (Currie 1995, Currie and Ravenscroft 1997, 2002). Simulation does not require that the simulated state replicate the target state exactly, and so it does not require that the states involved in visually imagining a hammer, for instance, be exactly the same as the visual experience one would have when actually seeing that hammer. All that is required is that one try to replicate, in some significant aspects, the state. When this point is kept in mind, it becomes highly plausible that at least some forms of imagery are simulations. In particular, there is supportive research on visual and motor imagery, which are best understood as simulations of visual perception and motor execution, respectively.

To take visual imagery first, there is extensive evidence that constraints on vision and visual illusions are replicated in visual imagery. For instance, the field of our actual visual experience subtends only a limited angle of the total input from our eyes. Stephen Kosslyn (1978) had one group of subjects actually approach items on a wall until they "overflowed" the visual field and another group imagine doing the same thing. Overflow occurred at almost exactly the same distance in real and imagined space, suggesting that the imagining subjects were simulating visual perception. A more dramatic example of this replication of visual constraints in imagery is the clinical evidence for *unilateral visual neglect*, in which damage to the posterior right parietal lobe can lead patients to fail to see the left half of the scene before their eyes. Bisiach and Luzzatti (1978) found that their two patients with this syndrome neglected not only the left half of their actual visual field but also the left half of their imagined visual field. Asked to imagine a familiar scene from a particular vantage point, they described only the right hand of the imagined scene; when asked to imagine the scene from the opposite vantage point, they described elements of the scene which were now on the right of the imagined scene, but neglected the previously described elements that were now on the left. Similarly, visual illusions like the Müller-Lyer lines and the McCullough Effect appear to have counterparts in visual imagery (Finke 1989).

Neuroscientific evidence also shows that vision and visual imagery activate (at least some of) the same areas of the brain. Both visual perception of faces and visual imagination of faces activate the fusiform gyrus (Kanwisher et al. 1997, O'Craven and Kanwisher 2000), and both these tasks are impaired by lesions to the fusiform gyrus. Kosslyn et al. (1999) studied another visual area of the brain, Area 17, by positron emission tomography (PET) and repetitive transcranial magnetic stimulation (rTMS). Subjects either saw or visualized sets of stripes. PET scans showed that Area 17 is activated in both tasks, and rTMS produced impairments

on both tasks, thus demonstrating that Area 17 is implicated in both vision and visual imagery.

The evidence in the motor domain is even more striking and less well known than in the visual case. For instance, Larry Parsons showed subjects a hand and asked whether it was a left or right hand (Parsons 1987, 1994). Subjects answered by imagining their own hand moving from its current orientation into the stimulus orientation, and compared the two. Moreover, the imagined trajectories were highly constrained by biomechanical constraints on actual hand movements. Parsons et al. (1998) further showed that motor imagery for the two hands is controlled by contralateral hemispheres of the brain, as is actual hand movement. Patients whose hemispheres had been disconnected were accurate at judging the handedness of a hand which was contralateral to the perceiving hemisphere, but at chance when judging the handedness of an ipsilateral hand. Evidently, imagining a hand movement is executed by the same cerebral mechanism that actually executes movement of that hand, which is to say that motor imagery involves neural mimicry, or simulation, of performing the movement. Visual and motor imagery are further discussed in Currie and Ravenscroft (2002) and Goldman (2006).

7 SIMULATIONAL MINDREADING OF EMOTIONS

None of the studies in Sections 5 or 6 speaks directly to the question of mindreading. They all demonstrate automatic mental simulation, or resonance, across individuals, but they leave open the question of whether an observer uses a mirroring event that he experiences (often at a covert level) to make an attribution to the actor, or target. We turn now to evidence that such simulation-based attributions do occur, especially attributions of emotion. Most of the studies to be discussed next involve patients with brain lesions. The principal tests of these patients are tests in which they observe slides or videotapes of facial expressions and are asked to classify the emotions being expressed, where the choices are drawn from a set of about six primary emotions: sadness, happiness, fear, anger, disgust, and surprise. Since these emotions are mental states, the judgments or attributions made by the subjects are attributions of mental states to the depicted targets.

A fundamental finding in this area is that for each of three basic emotions there was a paired deficit in experiencing that emotion and in attributing it, or "recognizing" it, in a facial expression. These deficits were *selective* in the sense that patients impaired specifically in emotion X had no difficulty recognizing emotion Y or Z but only recognizing X. (Sometimes matters were not quite so clear-cut.) In short, face-based emotion recognition, a specific type of low-level mindreading task, had a systematic pattern of paired deficits. Let us review some of these clinical and experimental findings (for fuller reviews, see Goldman and Sripada 2005; Goldman, 2006, chap. 6).

Adolphs and colleagues studied patients to determine whether damage to the amygdala, known to be associated with the production of fear, might also affect face-based recognition of fear (Adolphs et al., 1994). One patient, SM, was a

30-year-old woman with Urbach-Wiethe disease, a rare metabolic disorder that resulted in bilateral destruction of her amygdala. She was described as having a "fearless" nature (Damasio 1999, p. 66), and she reported not feeling afraid when shown film clips that normally elicit fear (e.g., 'The Shining' and 'The Silence of the Lambs'); but she seemed to experience other emotions strongly (Adolphs and Tranel 2000). SM was tested on various face-based emotion recognition (FaBER) tasks. SM was abnormal in face-based recognition of fear. Her ratings of fearful faces correlated less with normal ratings than did those of any of 12 brain-damaged control subjects (with different kinds of brain damage), and fell 2-5 standard deviations below the mean of the controls when the data were converted to a normal distribution. She was also abnormal to a lesser extent in recognizing anger and surprise.

Sprengelmeyer et al. (1999) studied another patient, NM, who also had bilateral amygdala damage. NM was also abnormal in experiencing fear. He liked dangerous activities, such as dangling from a helicopter while hunting deer in Siberia. Like SM, NM also exhibited a severe and selective impairment in fear recognition.

The second emotion for which a paired deficit was found is disgust. The fMRI study of disgust reported earlier was with *normal* subjects; the studies discussed next involve brain-damaged patients. As background, it is well established from both animal and human studies that taste processing is localized in the anterior insula region, and that disgust is an elaboration of a phylogenetically more primitive distaste response (Rozin et al. 2000). Calder et al. (2000) studied a patient NK who suffered from insula and basal ganglia damage. His overall score on a questionnaire for the experience of disgust was significantly lower than the controls, but his scores for anger and fear did not significantly differ from the scores of controls. Thus, his experience of disgust appeared to be selectively impaired. In tests of his ability to recognize emotions in faces, NK similarly showed significant and selective impairment in disgust recognition. That is, only in that emotion was his recognitional ability abnormal.

The third emotion for which a paired deficit has been found is anger. The neurotransmitter dopamine is apparently involved in the processing of aggression in social-agonistic encounters and plays an important role in mediating the experience of anger. Dopamine levels in rats are elevated during social-agonistic encounters, and increased dopamine levels can lead to enhanced appetitive aggression and agonistic dominance. Aware of the paired deficits discovered for fear and disgust, Lawrence and colleagues speculated that lowering the level of dopamine in normal individuals might not only reduce the production or experience of anger but also impair their ability to recognize anger in faces. Thus, Lawrence et al. (2002) administered the dopamine antagonist sulpiride to otherwise normal subjects, producing a temporary impairment in their dopamine system. As predicted, the sulpiride subjects performed significantly worse than controls at recognizing angry faces but no differently in recognizing facial expressions of other emotions.

What accounts for this striking pattern of paired deficits for three emotions? What does it tell us about the *ordinary* performance of FaBER tasks that damage

to the system in which an emotion is produced, or experienced, also damages the individual's capacity to recognize that emotion specifically? An obvious hypothesis is that when a normal person successfully executes a FaBER task for some emotion E, they do so via a simulation heuristic. That is, they use the very same neural system in making a recognition judgment as is used in experiencing emotion E. The classification process somehow uses the target emotion itself, or the neural substrate of that emotion. Naturally, when that neural substrate is damaged, emotion E is not produced and therefore they cannot correctly recognize emotion E in another's face. This germ of a simulationist idea was suggested in several of the original studies reported above, but not developed in detail.

Goldman and Sripada (2005) set out the argument more systematically (also see Goldman, 2006, chap. 6). Let the simulation hypothesis for mindreading be the hypothesis that some of the same events, processes or machinery are used in arriving at an accurate attribution of state M as one uses to token M oneself. This hypothesis predicts that damage to the system used to token M in oneself will also damage one's ability to attribute M accurately to others. Precisely this prediction is what the pattern of paired deficits bears out, in spades. So simulation theory predicts these observed findings. Does theory theory also predict them?

Goldman and Sripada argue to the contrary. They articulate three ways that a theory theorist might try to explain the paired deficit data. One of them, however, is quite ad hoc, and the other two are contravened by further evidence. One possibility is that subjects with FaBER deficits have difficulty with the perceptual processing of faces. In fact many of the studies tested their subjects on measures designed to detect such perceptual deficiencies (such as the Benton Face Matching Task), and the subjects were normal. Another possibility is that subjects with FaBER deficits are selectively impaired in declarative knowledge concerning the emotion in question, for example, its typical elicitors or behavioral effects. But researchers report that their FaBER-impaired subjects have no such declarative knowledge deficit. For example, Calder et al. (2001) reported that patients with disgust recognition impairments are able to provide plausible situations in which a person might feel disgusted and do not show impaired knowledge of the concept of disgust.

If FaBER tasks are indeed executed via simulation, how does the process go (in the case of normals)? The simulation account would be even more compelling if we had a clear picture of how the simulation routine proceeds. Goldman and Sripada propose four possible models for such a simulation routine. We won't review them in detail, but will merely provide one illustration.

One model is called a "reverse simulation" model, the idea being that the attributor runs the appropriate emotional process in the reverse direction. Basic emotions standardly cause a suite of physiological changes, including a characteristic facial expression (Ekman 1992). There is evidence that the process can run in reverse: manipulation of the facial musculature can generate, at least in attenuated form, the corresponding emotional state. Clever techniques have been used to induce smiles or frowns without subjects' awareness that they were smiling or

frowning, e.g., by having them hold a pen in their teeth (which induces a smile) or in their lips (which induces a frown). The selected facial musculature influences their mood or emotional state. Suppose, then, that observers of emotional facial expressions automatically mimic their target's facial expression. Such mimicry is known to occur automatically, rapidly and covertly (Dimberg and Thunberg 1998). These facial exertions would produce traces in the attributor of the same emotion experienced by the target. These traces would be classified for their emotion type and attributed to the target. This would be how normal people do FaBER tasks correctly. Now, if someone has brain damage that prevents him from undergoing a specific emotion, there will be no reliable basis for making correct attributions of *that* emotion. But attribution of other emotions will not be affected. Hence, a paired deficit.

This sounds good, but there is evidence against this model. Calder et al. (2000) found three patients with Mobius syndrome, a congenital syndrome involving complete facial paralysis, who performed normally on FaBER tasks. Keillor (2002) reported a similar finding in which a patient with bilateral facial paralysis performed normally on FaBER tasks. So the simulation process does not seem to depend on literal facial exertions. A permutation of the reverse simulation model is possible to accommodate this finding, but we won't pursue that or other models here (see Goldman and Sripada 2005, or Goldman, 2006, for further discussion).

A different question is whether the foregoing model is genuinely a simulation model. The routine described in this model would not exactly fit the traditional multi-step pattern of simulation-based attribution, involving "pretend" states being fed into appropriate cognitive machinery and the outputs of this machinery being taken "off-line". But on our preferred conception of simulation-based mindreading, it suffices that there be automatic replication (or semi-replication) of the target state, which the attributor uses to classify the state and attribute the classification to the target (see Goldman, 2006, chap. 2). These features are satisfied by the model.

8 HIGH-LEVEL MINDREADING

The examples of simulation discussed in the previous sections were either not cases of mindreading at all (Sections 5 and 6) or cases of only *low-level* mindreading (Section 7). Low-level mindreading involves automatic and pre-conscious mechanisms, responsible for mental ascriptions of relatively simple types, such as emotions or feelings. By contrast, *high-level* mindreading involves more controlled and effortful mechanisms, typically with ascriptions of a more complex type, viz. propositional attitudes (the high/low terminology is from Goldman, 2006). Perhaps surprisingly, there have not been very many attempts to subject the theory-theory/simulation-theory dispute to careful experimental test. At any rate, it has proved difficult to obtain experimental results on high-level mindreading that aren't open to alternative interpretations. Nonetheless, there is interesting evidence that is quite

congenial to the simulationist approach, and a small sample of this evidence will now be presented.

There is a venerable tradition in social psychology that documents people's tendency to project their own thoughts and preferences onto other people (Cronbach 1955; Ross, Greene and House 1977). In recent writing, there are several suggestions that people use their own beliefs and perceptions as a judgmental "anchor" from which they adjust – usually insufficiently – to accommodate differences between themselves and others (Nickerson 1999; Epley, Keysar, Van Boven and Gilovich 2003). As Nickerson (1999) explains it, imputing one's own knowledge or beliefs to a specific other is a default measure; it is what one does in the absence of knowledge, or of a basis for inferring, that the other's knowledge is different from one's own. This is similar to a phenomenon noted by Nichols and Stich's (2003), which they call *default belief attribution*. As mindreaders, we tend to assume that our targets share most of our beliefs, e.g., the beliefs that snow is white, that 2+2=4, etc. According to social psychologists, even when one gets evidence that a target's knowledge is different from one's own, there is still a strong tendency to project or impute one's own knowledge (or preference) to the target.

Here are a few summary conclusions by investigators of the propensity to overestimate the commonality of one's own knowledge. Keysar, Ginzel and Bazerman (1995) write:

> First, when others are more informed than they are themselves, people do not fully take into account others' privileged access to information; they sometimes behave as if the others do not have such extra information. Second, even when people know that others do not have access to their own privileged information, they may behave as if those others had access to this information. (Keysar et al. 1995, p. 283).

Steedman and Johnson-Laird (1980) surmised that speakers in conversations tend to assume that hearers know everything they themselves know about the world and about the conversations. Piaget (1962) had this to say in a similar vein:

> Every beginning instructor discovers sooner or later that his first lectures were incomprehensible because he was talking to himself, so to say, mindful only of his own point of view. He realizes only gradually and with difficulty that it is not easy to place oneself in the shoes of students who do not know what he knows about the subject matter of his course. (p. 5)

Keysar, Lin and Barr (2003) had two people play a communication game, in which the 'director' instructed a participant to move certain objects. The participant had to interpret the director's instruction, and the correct way to proceed, of course, would be to interpret the instructions in terms of what the *director* knew, not what the participant herself knew. But participants displayed a strong tendency to interpret the director's instructions in terms of the participant's own knowledge,

even when the participant was aware that she knew about a hidden object that the director was unaware of. This is an instance of what psychologists call *egocentrism*.

Camerer, Loewenstein and Weber (1989) investigated situations in which people who were relatively well informed with respect to some economic variables were required to predict what other, less informed people would forecast (the forecasts had to do with corporate earnings). It would work to the advantage of these better-informed people to discount completely their own knowledge that the less informed people lacked. But they failed to do so; they discounted it only partially. This difficulty in setting aside one's own knowledge was called *the curse of knowledge*.

Another example, due to Fischhoff (1975), is *hindsight bias*. In hindsight, people are prone to overestimate the degree to which they anticipated a future event before it occurred. This can be seen as a special case of the tendency to overimpute one's own knowledge to others, where the "other", in this case, is oneself at an earlier time. One assumes that one already had knowledge that, in fact, one only recently acquired.

Borrowing the "curse of knowledge" phraseology, Birch and Bloom (2003, 2004) performed an experiment showing that the same phenomenon is found in young children. Children were presented with two toys that contained something inside, one toy described as being familiar to a puppet named Percy and another unfamiliar to Percy. The children were asked to judge whether Percy would know what was inside the toy, where the children themselves sometimes knew the toys' contents and sometimes not. When the children themselves knew the toys' contents, they tended to overestimate what Percy knew, thereby exemplifying the curse of knowledge. This tendency significantly declined from age three to age five. Birch and Bloom speculate that the curse of knowledge might explain young children's failure in false-belief tasks, in both the displaced object version and deceptive container version.

It may strike the reader that all this evidence cuts *against* simulation theory, not in favor of it. Doesn't it show that people are very bad at putting themselves in the shoes of others, or adopting their perspective? To some extent this is right. But it is a mistake to formulate simulation theory as holding that people are very good at perspective taking, or that they even try to make imaginative adjustments for each and every mental state of a target they seek to mindread. Such massive adjustment would be too demanding, so it is more reasonable to assume that people make only *selective* adjustments for the states viewed as most likely to differ as between self and target. Furthermore, differences in knowledge and belief states may be particularly neglected. People experience their beliefs as veridical, unbiased representations of the world rather than subjective constructions; so there is a strong presumption that others will share their beliefs (Gilovich 1990; Griffin and Ross 1991).

These findings are congenial to simulation theory because one of the crucial underlying ideas of the theory is that mindreading involves a projection, imputation, or transference of one's own mental state to the target. In the standard simulationist paradigm, this occurs in the final stage. The usual paradigm features

four steps: the attributor (1) creates some pretend states, (2) feeds them into a cognitive mechanism, (3) lets the mechanism operate on the pretend states to generate an output state (in the self), and (4) transfers or imputes the output state to the target. Step (4), the projection or transference step, is what the foregoing discussion highlights. In fact, simulation theory would say that in default belief attribution the first three steps are omitted. The default belief process, then, should be regarded as a special case, perhaps a "limiting case," of simulation. Nonetheless, the process described is more akin to simulation than to theorizing, because it depicts attributors as transferring one of their own (first-order) states to targets. Such a move has no obvious place within theory theory.

Of course, the full story of belief attribution as told by many of the writers cited above includes a step that allows individuative evidence concerning the target to override the default assumption of belief similarity. This step introduces a 'theorizing' element into the process. This step would be at odds with pure simulationism but not with hybrid simulationism (the view we favor).

Pure theory theory could also devise a hypothesis to accommodate the findings. It might posit two extra beliefs in the head of each belief-ascribing attributor: (i) a second-order belief to the effect that the attributor has the first-order belief in question, and (ii) a belief in the generalization that any of one's own first-order beliefs are likely to be shared by another random individual. These extra beliefs are dubious posits; but it's hard to rule them out, which illustrates why the simulation/theory dispute is so hard to resolve.

Other social psychologists who study mental-state attribution find that the simulation, or perspective-taking, approach helps explain their experimental results. In a study by Van Boven, Dunning and Loewenstein (2000), some participants were given Cornell coffee mugs and asked to indicate the lowest price they would sell their mugs for. Other participants didn't receive mugs but were asked to indicate the highest price they would pay to purchase one. Mug owners then estimated the highest purchase price of the average buyer and buyers estimated the lowest selling price of the average owner. Since prices reflect valuations, these price estimates were, in effect, mental-state estimates. Apparently, being endowed with a mug makes a big difference to one's valuation, although this difference – the so-called endowment effect – is not adequately appreciated. Both owners and buyers underestimated the difference between themselves and their opposite numbers. These differences were called "egocentric empathy gaps". The only manipulation by the investigators that substantially reduced these gaps was the creation of buyers' agents who represented buyers in transactions with owners. Although buyers' agents who never owned a mug made low offers, buyers' agents who were endowed with mugs made significantly higher offers. Apparently, as van Boven, Dunning and Loewenstein propose, buyers' agents who owned a mug were better able to imagine how they themselves would feel if they were in the other role; in other words, they were better able to simulate the valuations of the sellers. One might question whether this interpretation was clearly dictated by the experiment itself, but, at a minimum, the results of the experiment are suggestive.

Note that in this last experiment, attributors assumed that their targets shared their preferences or evaluations. This suggests that there is more to default attribution than belief attribution; mindreaders tend to assume that others share additional mental states as well. Default preference (or desire) attribution was evidenced in an experiment by Ross, Greene and House (1977). Subjects were asked to wear a large sandwich-board reading "Eat at Joe's" while walking around campus for half an hour, and later were asked how many of their peers would agree to the same request. Subjects who had agreed predicted that 62% of their peers would also agree, while subjects who had refused predicted that only 33% of their peers would agree (and the rest would refuse). Thus, subjects tended to predict that others would feel the same way they did.

Although simulation theory is primarily concerned with third-person mindreading, it is also relevant to self-attribution in a predictive or hypothetical stance. The simulation-favored technique of putting yourself in another person's shoes could be permuted into imagining yourself in some future, or hypothetical, shoes. Is this a technique people actually use?

Gilbert, Gill and Wilson (2002) suggest that it is. They talk about imagining oneself in a certain situation to see what affective responses one would undergo. If we wish to predict how we would feel if we found our spouse in bed with the letter carrier on New Year's Eve, we might imagine the event and notice how we react to the mental image. In other words, one pretends to be in that situation, feeds this pretend state into an affect-generating cognitive mechanism, and lets the mechanism produce affective outputs, e.g., waves of jealousy and anger. These outputs – which Gilbert et al. call "proxy reactions" – can be used to predict how one would react in the hypothetical situation.

Gilbert et al. do not claim that predictors make unmodified use of proxy reactions. Like other investigators discussed earlier, they suggest that people use a more complex, three-step procedure. First, they imagine events as described above. Second, they use their affective or hedonic reactions as the basis for a preliminary prediction. Third, they correct or adjust their preliminary forecast by considering the target event's temporal location and other known features about it. The entire procedure exemplifies the anchoring-and-adjustment heuristic proposed by Tversky and Kahneman (1974). The third step introduces an element of theorizing, which fits the hybrid form of simulationism, our overall preferred view (see Goldman, 2006).

9 CONCLUSION

The theme of mental simulation has roots in the speculations of historical philosophers and has been developed in recent decades by analytic philosophers of mind and cognitive scientists, especially neuroscientists. It is principally used as a theory of mindreading, but is also directly applicable to empathic phenomena that may or may not give rise to mindreading. At a low level of cognition, largely below the level of consciousness, there are striking examples of mental mirroring

or resonance in such domains as action representation, feeling and emotion. At a higher level of cognition, putting oneself into another person's shoes seems to play an important role in mindreading, though evidence indicates that people often project their own states onto others with minimal and inadequate amounts of perspective taking. Even in the latter cases, however, the self is used as a model of a target, a theme that is central to simulation theory.

BIBLIOGRAPHY

[Adolphs and Tranel, 2000] R. Adolphs and D. Tranel. Emotion recognition and the human amygdala. In J. P. Aggleton, ed., *The Amygdala: A Functional Analysis*, 2^{nd} ed. (pp. 587-630). Oxford: Oxford University Press, 2000.

[Adolphs *et al.*, 1994] R. Adolphs, D. Tranel, H. Damasio, and A. Damasio. Impaired recognition of emotion in facial expressions following bilateral damage to the amygdala. *Nature* 372: 669-672, 1994.

[Baron-Cohen, 1995] S. Baron-Cohen. *Mindblindness: An Essay on Autism and Theory of Mind*. Cambridge, MA: MIT Press, 1995.

[Bartsch and Wellman, 1989] K. Bartsch and H. Wellman. Young children's attribution of action to beliefs and desires. *Child Development* 60: 946-964, 1989.

[Bartsch and Wellman, 1995] K. Bartsch and H. Wellman. *Children Talk about the Mind*. Oxford: Oxford University Press, 1995.

[Birch and Bloom, 2003] S. A. J. Birch and P. Bloom. Children are cursed: an asymmetric bias in mental-state attribution. *Psychological Science* 14: 283-286, 2003.

[Birch and Bloom, 2004] S. A. J. Birch and P. Bloom. Understanding children's and adults' limitations in mental state reasoning. *Trends in Cognitive Sciences* 8: 255-260, 2004.

[Bisiach and Luzzatti, 1978] E. Bisiach and C. Luzzatti. Unilateral Neglect of Representational Space. *Cortex* 14:129-133, 1978.

[Block, 1978] N. Block. Troubles with functionalism. In C. W. Savage, ed., *Perception and Cognition, Issues in the Foundations of Psychology. Minnesota Studies in the Philosophy of Science, vol. ix.* Minneapolis: University of Minnesota Press, 1978.

[Botterill and Carruthers, 1999] G. Botterill and P. Carruthers. *The Philosophy of Psychology.* Cambridge: Cambridge University Press, 1999.

[Buccino *et al.*, 2001] G. Buccino, F. Binkofski, G. R. Fink, L. Fadiga, L. Fogassi, V. Gallese, R. J. Seitz, K. Zilles, G. Rizzolatti, and H.-J. Freund. Action observation activates premotor and parietal areas in a somatotopic manner: An fMRI study. *European Journal of Neuroscience* 13(2): 400-404, 2001.

[Calder *et al.*, 2000] A. J. Calder, J. Keane, J. Cole, R. Campbell, and A. W. Young. Facial expression recognition by people with Mobius syndrome. *Cognitive Neuropsychology* 17(1–3): 73-87, 2000.

[Calder *et al.*, 2000a] A. J. Calder, J. Keane, F. Manes, N. Antoun, and A. W. Young. Impaired recognition and experience of disgust following brain injury. *Nature Reviews Neuroscience* 3: 1077-1078, 2000a.

[Calder *et al.*, 2001] A. J. Calder, A. D. Lawrence, and A. W. Young. Neuropsychology of fear and loathing. *Nature Reviews Neuroscience* 2: 352-363, 2001.

[Camerer *et al.*, 1989] C. Camerer, G. Loewenstein, and M. Weber. The curse of knowledge in economic settings: An experimental analysis. *Journal of Political Economy* 97: 1232-1254, 1989.

[Carey and Spelke, 1991] S. Carey and E. Spelke. Knowledge acquisition: Enrichment or conceptual change? In S. Carey and R. Gelman, eds., *Epigenesis of Mind: Studies in Biology and Cognition.* Hillsdale, NJ: Erlbaum, 1991.

[Carlson and Moses, 2001] S. M. Carlson and L. J. Moses. Individual differences in inhibitory control and children's theory of mind. *Child Development* 72: 1032-1053, 2001.

[Cronbach, 1955] L. J. Cronbach. Processes affecting score on "understanding others" and "assumed similarity." *Psychological Bulletin* 9: 597-608, 1955.

[Csibra *et al.*, 1999] G. Csibra, G. Gergely, S. Biro, O. Koos, and M. Brockbank. Goal attribution without agency cues: the perception of 'pure reason' in infancy. *Cognition* 72: 237-267, 1999.

[Currie, 1995] G. Currie. Imagination and simulation: Aesthetics meets cognitive science. In M. Davies and T. Stone, eds., *Mental Simulation* (pp. 151-169). Oxford: Blackwell, 1995.

[Currie and Ravenscroft, 1997] G. Currie and I. Ravenscroft. Mental simulation and motor imagery. *Philosophy of Science* 64: 161-180, 1997.

[Currie and Ravenscroft, 2002] G. Currie and I. Ravenscroft. *Recreative Minds*. Oxford: Oxford University Press, 2002.

[Damasio, 1999] A. Damasio. *The Feeling of What Happens*. New York: Harcourt Brace and Company, 1999.

[Davidson, 1984] Davidson, D. (1984). *Inquiries into Truth and Interpretation*. Oxford: Oxford University Press.

[Davies and Stone, 1995a] Davies, M. and Stone, T., eds. (1995a). *Folk Psychology*. Oxford: Blackwell.

[Davies and Stone, 1995b] Davies, M. and Stone, T., eds. (1995b). *Mental Simulation*. Oxford: Blackwell.

[Dennett, 1978] Dennett, D. C. (1978). Beliefs about beliefs. *Behavioral and Brain Sciences* 1: 568-570.

[Dennett, 1987] Dennett, D. C. (1987). *The Intentional Stance*. Cambridge, MA: MIT Press.

[Diamond and Taylor, 1996] Diamond, A. and Taylor, C. (1996). Development of an aspect of executive control: Development of the abilities to remember what I said and to 'Do as I say, not as I do'. *Developmental Psychology* 29: 315-334.

[Dimberg and Thunberg, 1998] Dimberg, U. and Thunberg, M. (1998). Rapid facial reactions to emotional facial expressions. *Scandinavian Journal of Psychology* 39: 39-45.

[Ekman, 1992] Ekman, P. (1992). Are there basic emotions? *Psychological Review* 99(3): 550-553.

[Epley *et al.*, 2003] Epley, N., Keysar, B., Van Boven, L., and Gilovich, T. (2003). Perspective taking as anchoring and adjustment. Manuscript submitted for publication.

[Fadiga *et al.*, 1995] Fadiga, L., Fogassi, L, Pavesi, G., and Rizzolatti, G. (1995). Motor facilitation during action observation: a magnetic stimulation study. *Journal of Neurophysiology* 73(6): 2608-2611.

[Finke, 1989] Finke, R. (1989) *Principles of Mental Imagery*. Cambridge, MA: MIT Press.

[Fischoff, 1975] Fischhoff, B. (1975). Hindsight $\neq$ foresight: The effect of outcome knowledge on judgment under uncertainty. *Journal of Experimental Psychology: Human Perception and Performance 1: 288-299.*

[Flavell *et al.*, 1986] Flavell, J. H., Green, F. L., and Flavell, E. R. (1986). *Development of knowledge about the appearance-reality distinction*. Monographs of the Society for Research in Child Development. No. 212, 51(1).

[Fletcher *et al.*, 1995] Fletcher, P., Happé, F., Frith, U., Baker, S., Dolan, R., Frackowiak, R., and Frith, C. (1995). Other minds in the brain: A functional imaging study of 'theory of mind' in story comprehension. *Cognition* 57: 109-128.

[Fodor, 1983] Fodor, J. A. (1983). *The Modularity of Mind*. Cambridge, MA: MIT Press.

[Fodor, 1992] Fodor, J. A. (1992). A theory of the child's theory of mind. *Cognition* 44: 283-296.

[Frith and Frith, 2003] Frith, U. and Frith, C. D. (2003). Development and neurophysiology of mentalizing. *Philosophical Transactions of the Royal Society, Series B* 358: 459-473.

[Frye *et al.*, 1995] Frye, D., Zelazo, P., and Palfai, T. (1995). Theory of mind and rule-based reasoning. *Cognitive Development* 10: 483-527.

[Gallese, 2001] Gallese, V. (2001). The 'shared manifold' hypothesis: From mirror neurons to empathy. *Journal of Consciousness Studies* 8 (5-7): 33-50.

[Gallese, 2003] Gallese, V. (2003). The manifold nature of interpersonal relations: the quest for a common mechanism. *Philosophical Transactions of the Royal Society, Series B* 358: 517-528.

[Gallese *et al.*, 1996] Gallese,V., Fadiga, L., Fogassi, L., and Rizzolatti, G. (1996). Action recognition in the premotor cortex. *Brain* 119: 593-609.

[Gallese and Goldman, 1998] Gallese, V. and Goldman, A. (1998). Mirror neurons and the simulation theory of mind-reading. *Trends in Cognitive Sciences* 2: 493-501.

[Gergely *et al.*, 1995] Gergely, G., Nadasdy, Z., Csibra, G., and Biro, S. (1995). Taking the intentional stance at 12 months of age. *Cognition* 56: 165-193.

[Gilbert *et al.*, 2002] Gilbert, D. T., Gill, M. J., and Wilson, T. D. (2002). The future is now: Temporal correction in affective forecasting. *Organizational Behavior and Human Decision Processes* 88(1): 430-444.

[Gilovich, 1990] Gilovich, T. (1990). Differential construal and the false consensus effect. *Journal of Personality and Social Psychology* 59: 623-634.

[Goldman, 1989] Goldman, A. I. (1989). Interpretation psychologized. *Mind and Language* 4: 161-185. Reprinted in Davies and Stone (1995a).

[Goldman, 1992a] Goldman, A. I. (1992a). In defense of the simulation theory. *Mind and Language* 7: 104-119. Reprinted in Davies and Stone (1995a).

[Goldman, 1992b] Goldman, A. I. (1992b). Empathy, mind and morals. *Proceedings and Addresses of the American Philosophical Association*. Reprinted in Davies and Stone (1995b).

[Goldman, 2000] Goldman, A. I. (2000). The mentalizing folk. In D. Sperber, ed., *Metarepresentation*. Oxford: Oxford University Press.

[Goldman, 2006] Goldman, A. I. (2006). *Simulating Minds: The Philosophy, Psychology and Neuroscience of Mindreading*. New York: Oxford University Press.

[Goldman and Sripada, 2005] Goldman, A. I. and Sripada, C. S. (2005). Simulationist models of face-based emotion recognition. *Cognition* 94: 193-213.

[Gopnik, 1993] Gopnik, A. (1993). How we know our minds: The illusion of first-person knowledge of intentionality. *Behavioral and Brain Sciences* 16: 1-14.

[Gopnik and Astington, 1988] Gopnik, A. and Astington, J. (1988). Children's understanding of representational change and its relation to the understanding of false belief and the appearance-reality distinction. *Child Development* 59: 26-37.

[Gopnik and Meltzoff, 1997] Gopnik, A. and Meltzoff, A. N. (1997). *Words, Thoughts and Theories*. Cambridge, MA: MIT Press.

[Gopnik and Wellman, 1992] Gopnik, A. and Wellman, H. (1992). Why the child's theory of mind really *is* a theory. *Mind and Language* 7: 145-171.

[Gopnik and Wellman, 1994] Gopnik, A. and Wellman, H. (1994). The theory theory. In L. Hirschfeld and S. Gelman, eds., *Mapping the Mind: Domain Specificity in Cognition and Culture*. Cambridge: Cambridge University Press.

[Gordon, 1986] Gordon, R. (1986). Folk psychology as simulation. *Mind and Language* 1: 158-171. Reprinted in Davies and Stone (1995a).

[Gordon, 1992] Gordon, R. (1992). The simulation theory: Objections and misconceptions. *Mind and Language* 7: 11-34. Reprinted in Davies and Stone (1995a).

[Gordon, 1995] Gordon, R. (1995). Simulation without introspection or inference from me to you. In M. Davies and T. Stone, eds., *Mental Simulation: Evaluations and Applications*. Oxford: Blackwell.

[Gordon, 1996] Gordon, R. (1996). 'Radical' simulationism. In P. Carruthers and P. Smith, eds., *Theories of Theories of Mind*. Cambridge: Cambridge University Press.

[Griffin and Ross, 1991] Griffin, D. W. and Ross, L. (1991). Subjective construal, social inference, and human misunderstanding. In M. Zanna, ed., *Advances in Experimental Social Psychology. Vol. 24* (pp. 319-339). New York: Guilford.

[Harman, 1978] Harman, G. (1978). Studying the chimpanzee's theory of mind. *Behavioral and Brain Sciences* 1: 576-577.

[Harris, 1991] Harris, P. (1991). The work of the imagination. In A. Whiten, ed., *Natural Theories of Mind*. Oxford: Blackwell.

[Harris, 1992] Harris, P. (1992). From simulation to folk psychology: The case for development. *Mind and Language* 7: 120-144. Reprinted in Davies and Stone (1995a).

[Heal, 1986] Heal, J. (1986). Replication and functionalism. In J. Butterfield, ed., *Language, Mind, and Logic*. Cambridge: Cambridge University Press.

[Heal, 1996] Heal, J. (1996). Simulation and cognitive penetrability. *Mind and Language* 11: 44-67.

[Heal, 1998] Heal, J. (1998). Co-cognition and off-line simulation: Two ways of understanding the simulation approach. *Mind and Language* 13: 477-98.

[Heal, 2003] Heal, J. (2003). *Mind, Reason and Imagination*. Cambridge: Cambridge University Press.

[Heyes, 1998] Heyes, C. M. (1998). Theory of mind in nonhuman primates. *Behavioral and Brain Sciences* 21: 101-148.

[Hirschfeld and Gelman, 1994] Hirschfeld, L. and Gelman, S. eds. (1994). *Mapping the Mind: Domain-specificity in Cognition and Culture*. Cambridge: Cambridge University Press.

[Hughes, 1998] Hughes, C. (1998). Finding your marbles: Does preschoolers' strategic behavior predict later understanding? *Developmental Psychology* 34(6): 1326-1339.

[Hughes and Russell, 1993] Hughes, C., and Russell, J. (1993). Autistic children's difficulty with mental disengagement from an object: Its implication for theories of autism. *Developmental Psychology* 29: 498-510.

[Hume, 1739/1978] Hume, D. (1739/1978). *A Treatise of Human Nature*, 2^{nd} ed., L. A. Selby-Bigge, ed. Oxford: Oxford University Press.

[Hutchison *et al.*, 1999] Hutchison, W. D., Davis, K. D., Lozano, A. M., Tasker, R. R., and Dostrovsky, J. O. (1999). Pain-related neurons in the human cingulate cortex. *Nature Neuroscience* 2(5): 403-405.

[Kant, 1781/1953] Kant, I. (1781/1953). *Critique of Pure Reason*, 1^{st} ed., trans. Norman Kemp Smith. London: Macmillan.

[Kanwisher *et al.*, 1997] N. Kanwisher, J. McDermott, and M. M. Chun. The fusiform face area: a module in human extrastriate cortex specialized for face perception. *Journal of Neuroscience* 17: 4302-4311, 1997.

[Keillor *et al.*, 2002] Keillor, J. M., Barrett, A. M., Crucian, G. P., Kortenkamp, S. and Heilman, K. M. (2002). Emotional experience and perception in the absence of facial feedback. *Journal of the International Neuropsychological Society* 8(1): 130-135.

[Keysar *et al.*, 1995] Keysar, B., Ginzel, L. E., and Bazerman, M. H. (1995). States of affairs and states of mind: The effects of knowledge of beliefs. *Organizational Behavior and Decision Processes* 64: 283-293.

[Keysar *et al.*, 2003] Keysar, B., Lin, S. and Barr, D. (2003). Limits on theory of mind use in adults. *Cognition* 89: 25-41.

[Keysers *et al.*, 2004] Keysers, C., Wicker, B., Gazzola, V., Anton, J.-L., Fogassi, L. and Gallese, V. (2004). A touching sight: SII/PV activation during the observation of touch. *Neuron* 42: 335-346, 2004.

[Kohler *et al.*, 2002] Kohler, E., Keysers, C., Umilta, M. A., Fogassi, L., Gallese, V. and Rizzolatti, G. (2002). Hearing sounds, understanding actions: Action representation in mirror neurons. *Science* 297: 846-848.

[Kosslyn, 1978] Kosslyn, S. M. (1978). Measuring the visual angle of the mind's eye *Cognitive Psychology* 7: 341-370.

[Kosslyn *et al.*, 1999] Kosslyn, S. M., Pascual-Leone, A., Felician, O., and Camposano, S. (1999). The Role of Area 17 in Visual Imagery: Convergent Evidence and PET from rTMS *Science* 284: 167-170.

[Lawrence, 2002] Lawrence, A. D., Calder, A. J., McGowan, S. M., Grasby, P. M. (2002). Selective disruption of the recognition of facial expressions of anger. *NeuroReport* 13(6), 881-884.

[Leslie, 1987] Leslie, A. M. (1987). Pretence and representation: The origins of 'theory of mind'. *Psychological Review* 94: 412-426.

[Leslie, 1994] Leslie, A. M. (1994). Pretending and believing: issues in the theory of ToMM. *Cognition* 50: 211-238.

[Leslie and German, 1995] Leslie, A. M. and German, T. (1995). Knowledge and ability in 'theory of mind': One-eyed overview of a debate. In M. Davies and T. Stone, eds., *Mental Simulation: Philosophical and Psychological Essays*. Oxford: Blackwell.

[Leslie *et al.*, 2005] A. M. Leslie, T. German, and P. Polizzi. Belief-desire reasoning as a process of selection. *Cognitive Psychology* 50(1):45-85, 2005.

[Leslie and Polizzi, 1998] Leslie, A.M. and Polizzi, P. (1998). Inhibitory processing in the false belief task: Two conjectures. *Developmental Science* 1: 247-253.

[Lewis, 1972] Lewis, D. K. (1972). Psychophysical and theoretical identifications. *Australasian Journal of Philosophy* 50: 249-258.

[Lhermitte *et al.*, 1986] Lhermitte, F., Pillon, B., and Serdaru, M. (1986). Human autonomy and the frontal lobes. Part I: Imitation and utilization behavior: a neuropsychological study of 75 patients. *Annals of Neurology* 19(4): 326-34.

[Meltzoff and Brooks, 2001] Meltzoff, A. N. and Brooks, R. (2001). "Like me" as a building block for understanding other minds: Bodily acts, attention, and intention. In B. Malle, L. Moses, and D. Baldwin, eds., *Intentions and Intentionality: Foundations of Social Cognition* (pp. 171-192). Cambridge, MA: MIT Press.

[Mitchell and Lacohee, 1991] Mitchell, P. and Lacohee, H. (1991). Children's early understanding of false belief. *Cognition* 39: 107-127.

[Nichols and Stich, 2003] Nichols, S. and Stich, S. (2003). *Mindreading: An Integrated Account of Pretence, Self-Awareness, and Understanding of Other Minds*. Oxford: Oxford University Press.

[Nickerson, 1999] Nickerson, R. S. (1999). How we know—and sometimes misjudge—what others know: Imputing one's own knowledge to others. *Psychological Bulletin* 125: 737-759.

[Nietzsche, 1881/1982] Nietzsche, Friedrich (1881/1982). *Daybreak: Thoughts on the prejudices of morality*, trans. R. J. Hollingdale. Cambridge: Cambridge University Press.

[O'Craven and Kanwisher, 2000] O'Craven, K. M. and Kanwisher, N. (2000). Mental Imagery of Faces and Places Activates Corresponding Stimulus-Specific Brain Regions. *Journal of Cognitive Neuroscienc* 12:1013-1023.

[Ozonoff *et al.*, 1991] Ozonoff, S., Pennington, B. F. and Rogers, S. J. (1991). Executive function deficits in high-functioning autistic individuals: Relationship to theory of mind. *Journal of Child Psychology and Psychiatry* 32: 1081-1095.

[Parsons, 1987] Parsons, L. M. (1987). Imagined spatial transformation of one's hand and feet. *Cognitive Psychology* 19: 178-241.

[Parsons, 1994] Parsons, L. M. (1994). Temporal and kinematic properties of motor behavior reflected in mentally simulated action. *Journal of Experimental Psychology* 20: 709-730.

[Parsons *et al.*, 1998] Parsons, L. M., Gabrieli, J., Phelps, E., and Gazzaniga, M. (1998). Cerebrally lateralized mental representations of hand shape and movement. *Journal of Neuroscience* 18: 6539-6548.

[Perner, 1991] Perner, J. (1991). *Understanding the Representational Mind*. Cambridge, MA: MIT Press.

[Perner *et al.*, 1987] Perner, J., Leekam, S., and Wimmer, H. (1987). Three-year-olds' difficulty with false belief: The case for a conceptual deficit. *British Journal of Developmental Psychology* 5: 125-137.

[Piaget, 1962] Piaget, J. (1962). Comments: Addendum to Vygotsky, L. S. In J. Hanfmann and G. Valcar, eds. and trans., *Thought and Language*. Cambridge, MA: MIT Press.

[Povinelli *et al.*, 2000] Povinelli, D. J., Bering, J. M. and Giambrone, S. (2000). Toward a science of other minds: Escaping the argument by analogy. *Cognitive Science* 24: 509-541.

[Premack and Woodruff, 1978] Premack, D. and Woodruff, G. (1978). Does the chimpanzee have a theory of mind? *Behavioral and Brain Sciences* 1: 515-526.

[Quine, 1960] Quine, W. V. O. (1960). *Word and Object*. Cambridge, MA: MIT Press.

[Quine, 1990] Quine, W. V. O. (1990). *Pursuit of truth*. Cambridge, MA: MIT Press.

[Rizzolatti *et al.*, 1996] Rizzolatti, G., Fadiga, L., Gallese, V., and Foggasi, L. (1996). Premotor cortex and the recognition of motor actions. *Cognitive Brain Research* 3: 131-141.

[Ross *et al.*, 1977] Ross, L., Greene, D., and House, P. (1977). The "false consensus" effect: An egocentric bias in social perception and attribution processes. *Journal of Experimental Social Psychology* 13: 279-301.

[Rozin *et al.*, 2000] Rozin, P., Haidt, J., and McCauley, C. (2000). Disgust. In M. Lewis and J. Haviland, eds., *Handbook of Emotions*. New York: The Guilford Press.

[Scholl and Leslie, 1999] Scholl, B. and Leslie, A. M. (1999). Modularity, development and 'theory of mind'. *Mind and Language* 14: 131-153.

[Shoemaker, 1984] Shoemaker, S. (1984). Some varieties of functionalism. In *Identity, Cause, and Mind: Philosophical Essays*. Cambridge: Cambridge University Press.

[Singer *et al.*, 2004] Singer, T., Seymour, B., O'Doherty, J., Kaube, H., Dolan, R., and Frith, C. (2004). Empathy for pain involves the affective but not sensory components of pain. *Science* 303: 1157-1162.

[Smith, 1759/1976] Smith, A. (1759/1976). *A Theory of Moral Sentiments*. D. D. Raphael and A. L. Macfie, eds. Oxford: Clarendon Press.

[Sprengelmeyer *et al.*, 1999] Sprengelmeyer, R., Young, A. W., Schroeder, U., Grossenbacher, P. G., Federlein, J., Buttner, T., and Przuntek, H. (1999). Knowing no fear. *Proceedings of the Royal Society, Series B* 266: 2451-2456.

[Steedman and Johnson-Laird, 1980] Steedman, M. J. and Johnson-Laird, P. N. (1980). The production of sentences, utterances and speech acts: Have computers anything to say? In B. Butterworth, ed., *Language Production: Vol. 1. Speech and Talk*. London: Academic Press.

[Tomasello and Call, 1997] Tomasello, M. and Call, J. (1997). *Primate Cognition*. Oxford: Oxford University Press.

[Tversky and Kahneman, 1974] Tversky, A. and Kahneman, D. (1974). Judgment under uncertainty: Heuristics and biases. *Science* 185: 1124-1131.

[Umilta *et al.*, 2001] M. A. Umilta, E. Kohler, V. Gallese, L. Fogassi, L. Fadiga, C. Keysers, and G. Rizzolatti. I know what you are doing. a neurophysiological study. *Neuron* 31(1): 155-65, 2001.

[van Boven *et al.*, 2000] Van Boven, L., Dunning, D. and Loewenstein, G. (2000). Egocentric empathy gaps between owners and buyers: Misperceptions of the endowment effect. *Journal of Personality and Social Psychology* 79: 66-76.

[Wellman *et al.*, 2001] Wellman, H. M., Cross, D. and Watson, J. (2001). Meta-analysis of theory-of-mind development: The truth about false belief. *Child Development* 72: 655-684.

[Wimmer and Perner, 1983] Wimmer, H. and Perner, J. (1983). Beliefs about beliefs: Representation and constraining function of wrong beliefs in young children's understanding of deception. *Cognition* 13: 103-128.

[Zaitchik, 1991] Zaitchik, D. (1991). Is only seeing really believing? Sources of the true belief in the false belief task. *Cognitive Development* 6: 91-103.

THE THEORETICAL AND METHODOLOGICAL FOUNDATIONS OF COGNITIVE NEUROSCIENCE

Valerie Gray Hardcastle

What we call cognitive neuroscience is really an amalgam field of eight different research areas in neuroscience proper: development, perception, action, attention, memory, higher cognitive functions (including language), affect, and plasticity. We ask of each subdiscipline: what can you tell us about how mammals, especially the higher primates, think? All the answers taken together comprise the field of cognitive neuroscience.

Obviously, these answers change on almost a daily basis, given how rapidly neuroscience itself is changing. It is now a truism that our knowledge of the brain and brain function is growing exponentially. In more or less a hundred years' time, scientists have outlined the basic structure of the nervous system; they have articulated the biochemical, physiological, and pharmacological features of brains, along with their perceptual, mnemonic, attentional, affective, and cognitive functionings. New and powerful technologies have given us recent extraordinary discoveries: reach neurons, mirror neurons, the neural mechanisms behind our sleep cycles, the neural locations of pain processing.

At the same time, there is still much we do not know. Indeed, many central aspects of the brain remains a deep mystery: we do not know exactly what the hippocampus does; we do not know why we sleep; we have not identified all our neurotransmitters; we do not even know how all the neurotrophins are exchanged among neurons. Suffice it to say that neuroscience as a discipline is in a period of tremendous flux.

This makes writing a useful handbook entry on this area rather difficult, for I am certain that by the time of publication, much that is currently considered groundbreaking research will have been discredited and the science will have moved on. Consequently, this chapter focuses less on specific research results — though I do sketch several examples — and more on the theoretical and methodological issues confronting neuroscientists working on cognition. For these, I believe, will be more enduring.

The first section below on the theoretical issues is necessarily brief, for most of the research results are already covered in other sections of this book. Here I aim just to give the broad flavor of the contributing fields as well as my views about how research in these areas as they pertain to cognitive neuroscience will unfold over the next decade or so. The second section uses research in brain plasticity

Handbook of the Philosophy of Science. Philosophy of Psychology and Cognitive Science
Volume editor: Paul Thagard
General editors: Dov M. Gabbay, Paul Thagard and John Woods
© 2007 Elsevier B.V. All rights reserved

to explore methodological concerns peculiar to neuroscience. This section is more detailed, for it outlines problems and issues largely unexplored in cognitive science as a whole.

1 TRENDS IN THE CONTRIBUTING RESEARCH AREAS

Developmental neurobiology gives cognitive neuroscience the rules that govern brain development, which then set the limits on its capacities for cognition and help to bridge the explanatory gaps between the molecular, cellular and behavioral levels in neuroscience. Developmental neurobiology itself is becoming increasingly more molecular as scientists can now examine directly the role of individual genes in the developmental pathway. It is also concentrating more and more on basic cellular events: the regulation of cell cycle kinetics, programmed cell death, cell lineages, and the mechanisms behind neuron migration. These changes in focus necessitate emphasizing simpler organisms and in vitro preparations, which in turn translates into less emphasis on how the synaptic circuitry actually develops in the mammalian brain. As a result, cognitive neuroscientists interested in developmental issues often turn to computational modeling and more qualitative descriptions of large-scale brain developmental patterns instead of "hard" data (cf. [Crick and Jones, 1993]). Hence, the direct connections between developmental neurobiology and cognitive studies appear to be weakening.

However, at the same time, there is a renewed emphasis on the fundamental importance of genes, environment, and their interactions as catalysts for growth and change in cognition. For example, the only way to explain the uniformity of structure in the cerebrum is in terms of a specific genetic program, but the only way to account for its actual functioning is via sensory stimulation. Scientists can now link critical periods in learning and behavior to specific stages of neuronal differentiation and genetically-driven changes in synaptic connectivity, yet our mental skills depend upon our brains being able to access and utilize appropriate environmental stimuli. More and better bridges are needed between the molecular and the behavioral levels — bridges that are data-driven and not model-driven — but these bridges are slowly being forged within cognitive neuroscience.

Our *perceptual systems* represent the external environment to us. Through complex computations we do not yet understand, our sensory systems derive stable images from the ever-fluctuating raw signals of our transducers. These representations are much richer than the basic measurements of intensity, force, and chemical composition that comprise their building blocks, for our brains represent meaningful objects, events, and states of affairs. The representations are also simpler than the initial measurements, for they are winnowed from vast sums of transducer data, much of which brains never use at all.

Questions neuroscientists concerned with perception and cognition try to answer include: How and where are our sensory signals encoded and stored? How do we separate "figure" from "ground"? How are incoming signals "mixed" with our memories, attention, and our understanding of the world so that we get full-blown

representational experiences? How do we combine information from different sensory modalities? How do other brain systems transform and use this information? How do they modulate the representations to meet our behavioral goals and biological needs? How do we use representations to regulate action, planning, and other outputs?

We perceive in order to act effectively in our environments, for without any perceptions, *action* would be impossible. Therefore, perception research is closely tied to research into motor systems. Our motor systems bring together representations from our different sensory modalities and in different coordinate frames and transform them into useful signals for motor planning and execution. We use our sensory inputs to help guide our hands in grasping an object, for example. The areas that perform these sorts of computations are thought to be reciprocally connected (e.g., [Sakata *et al.*, 1997]). We can see how complicated this procedure must be when we consider that direction, amplitude, and velocity planning information must somehow be sent to a large number of individually controlled muscles. Furthermore, each motor solution designed to get the hand to grasp an object is non-unique; it could be accomplished be an indefinite number of variations of shoulder, elbow, and wrist movement combinations, following an indefinite number of different trajectories through space. Coordination among all these options is no simple task.

We also have motor areas that store various behavioral "programs" that we can tap into as needed. This store helps simplify and streamline the sensory input-behavioral output processing chain. Most recently, perhaps as part of this storage system, neuroscientists have isolated the so-called mirror neurons, which fire when a monkey either performs a particular action or watches someone else perform the same action [Gallese *et al.*, 1996]. These mirror neurons perhaps allow monkeys (and other creatures) to learn about actions without actually having to execute them.

How creatures orient to and selectively enhance perceptual events has been the central focus of *attention* research relative to cognitive neuroscience. Our brains are bombarded by numerous competing stimuli inputs each moment in time, but somehow they figure out which certain few are the most important and they devote more resources to processing those. We can find the neural effects of attention all the way from neurochemical signatures to enhanced activity in single cells to changes in the time course of electrical activity across the scalp to altered interactions among brain regions. We can find early effects of attention in V1 [Moran and Desimone, 1985], often as less than 80 msec after stimuli presentation [Clark and Hillyard, 1996]. At the same time, unattended information appears to be processed quite deeply such that it can still influence behavior. What areas of the brain are involved in coordinating attention across our neural circuitry (the "source" of attention), and how these areas accomplish these tasks (at the "site" of attention) are still matters of investigation. Other research concerns how and whether attention is tied into our other cognitive systems, such as perception and memory.

The connection between attention and *memory* is particularly vexing, for both systems are so very complex. Just as we probably will not be able to discuss attention as a whole, but rather as a set of separate but interlocking systems, we too will only be able to understand memory by its smaller and more manageable components. Memory researchers divide memory into (roughly) the following categories: procedural and semantic or cognitive, explicit and implicit, recall and recollection. Each of these divisions might mark a completely separate brain system (or they all might be completely wrong). But in each of these areas it is clear that current mental experiences are altered as a result of prior experience-induced physical changes in the brain. Exactly how these changes occur is the central focus of investigation. Obviously, a single neuron by itself cannot encode or retrieve from memory a complex representation. Yet, just as obviously, physical changes happen at the level of individual neurons and their synapses. Scientists have learned a lot about changes in synapses relative to memory; they have identified about sixty molecules as participants in synaptic long-term potentiation, long presumed to be the cellular mechanism behind mnemonic consolidation (cf. [Kandel *et al.*, 2000]). The question is how the brain coordinates itself across neurons to produce global effects. We know much less about how all these changes sum to a genuine memory of some object or event.

Probably unique to human cognition is *language*. We do not know much about the evolutionary history of language, but we do know some things about which areas of the brain are engaged in linguistic processing. While in general phonation depends on evolutionarily older call systems in the caudal midbrain, in humans, the neocortex also controls it (cf., [MacNeilage, 1998]). Our left temporal lobes next to primary auditory cortex store phonetic codes. Several left-hemisphere perisylvian areas analyze and synthesize syntactic, morphological, phonological, phonetic, and semantic information. The impact of imaging studies on our understanding of the functional architecture of language has been great and will only continue to grow over the next several years.

Similar comments can be made about our other *higher brain functions*: mental imagery, mathematical reasoning, categorization, and rational decision-making. Especially through recent advances in imagining studies, scientists are learning where such functions occur in our brains. For example, fMRI studies indicate that object information is distributed across discrete cortical areas, with the features defining the object stored close to the primary sensory and motor areas activated when the subject encountered the object in the world [Martin *et al.*, 2000]. That is, color information is stored close to color perception areas and so forth. Thus, our so-called lower level processes support and organize our higher-level functions. Understanding one level, then, should aid in understanding the other.

Studies in *affect* were neglected in the cognitive neurosciences until the mid-1990s. Now it is one of neuroscience's largest areas of growth. Central to just about all neurobiological studies of emotion is the amygdala and its role in all our affective processing. However, apart from this interest in one brain region there is little agreement across the rest of emotion research. There is no agreement on how

to define emotions; scientists have little understanding regarding how cognition and affect interact with one another; they know little about the basic mechanisms of our positive emotions; they are not sure how the neural mechanisms for our emotions connect to our subjective feelings. In addition, it is dangerous to generalize too much from what we know about the amygdala, for these data come largely from conditioned fear studies in rats. How that relates to a feeling of job satisfaction, for example, is unclear. The next two decades should see exponential growth in our understanding of affect and its relation to cognition. Right now, though, more work needs to be done before we can claim a deep theoretical understanding of human emotion.

Finally, we literally could not think without brain *plasticity*, the collection of mechanisms that support the brain's ability to change rapidly. We find plasticity at multiple levels of organization, from long-term potentiation in molecules on up to large-scale systemic functions. How all these mechanisms are integrated is probably the most important (and challenging) question facing cognitive neuroscientists today. For common plastic mechanisms across different brain areas unite our disparate cortical and subcortical functional areas. Hence, analyzing the mechanisms of plasticity will give us insight into the basic organizational principles of cognitive neural architectures. Because of plasticity's absolute and fundamental importance in cognitive neuroscience, I will be using research from this area to discuss the unique methodological challenges confronting neuroscience today.

2 COGNITIVE NEUROSCIENCE METHODOLOGY

Most of the methodological concerns found in the more traditional sciences are recapitulated in cognitive neuroscience. In particular, neuroscientists worry about what counts as appropriate empirical justification for a theoretical claim, how to determine which level of organization is the correct one for a scientific explanation, what explanations should look like, whether all explanations will or should reduce to some primitives, and how what we learn about the mind/brain should affect the larger social, economic, and political arenas ([Bechtel *et al.*, 2001], [Schaffner, 1993]). It is difficult to focus on only one of these concerns to the exclusion of the rest. Most likely, as some particular aspect of the practice of neuroscience becomes understood, others will be as well. What follows are areas of methodological concern as they *differ* from traditional issues. This discussion therefore should be laid on top of and be seen to complement the very rich literature in traditional philosophy and methodology of science.

2.1 Neuroscientific Theories

Brains are complicated and messy affairs; theories about brains share these same traits. The difficulty is that in order to make a simple generalization about how some aspect of the brain functions, scientists have to retreat to such a broad level of abstraction that their assertions become almost empirically meaningless. In

order to make their claims testable in a laboratory, neuroscientists have to confine their ideas to particular animals, to particular experimental tasks, or to both. As a result, they end up with neuroscientific "theories" that contain two distinct parts: a broad statement of theoretical principle and a set of detailed descriptions of how that principle plays out across different animal models and experimental tasks. Though the detailed descriptions fall under the general principle, they are not immediately derivable from it. Moreover, as described below, the detailed descriptions can be incompatible with one another, though each will maintain a family resemblance with the others.

At a gross level, mammalian brains are remarkably similar to one another. Indeed, the central nervous system (CNS) in invertebrates is not all that different from the mammalian CNS either. There are innumerable homologous areas, cell types, neurotransmitters, peptides, chemical interactions, and so forth. However, there are important differences beyond these surface resemblances.

For example, consider the semicircular canal. All mammals have roughly the same five end organs in their ears to support their auditory and vestibular systems, and they all work to keep their lateral semicircular canals in their ears parallel to the horizontal plane relative to the Earth, for keeping it in that position allows them to get the best possible information about head position in space. (The lateral canal is maximally excitatory to a yaw (left to right) head motion; keeping the canal in line with the horizontal plane allows the organ to detect this motion with the greatest accuracy.) But rodents ambulate with their necks extended, which keeps their heads in an extreme dorsal position, while humans incline their heads about twenty degrees when walking naturally. In general, the differences in the shape of the semicircular canals in the ear with skull shape correlate with the position that an animal's head is normally in.

For another example, consider the retina. There are striking differences between herbivores and predators in brain structure, for creatures who munch on grasses and trees require much less precise environmental information than those who hunt moving targets in order to survive. As a result, rodents have no foveae. To maintain visual fixation on a point, they move their necks, using what is known as the vestibular-colic response. The vestibular system in their ears tells them how their head is oriented and they use that information to reorient their heads in order to keep whatever object currently fascinates them in their line of sight.

In contrast, primates have foveae and they move their eyeballs to keep their target within the foveal area, using the vestibular-ocular response. This is a much more precise orienting mechanism which allows them to move their eyes to compensate for changes in head position such that they can keep objects foveated for as long as they wish. For some indication of how important computing horizontal eye motion is to primate brains, consider that the abducens (or VIth) nerve in humans, which controls horizontal eye abduction, feeds into one of the biggest motor nuclei in the brain stem. This ocular nucleus, which controls only one very tiny muscle, is only slightly smaller than the nucleus that controls all of the twenty or so facial muscles.

In more striking contrast still, bats do not maintain ocular position in the same fashion as the rest of the mammals. Because they fly and so have greater freedom to move in three-dimensional space, maintaining body position relative to the horizontal is not an easy option. As a result, they use other sense organs, primarily hearing (the other half of the VIIIth nerve), to determine how their eyes should be oriented. Consequently, they need not rely on vestibular-ocular responses as we do, even though their bodies are equipped with such reflex machinery.

All of these anatomical and physiological differences are important when neuroscientists want to investigate something like how the brain learns to compensate for damage to the vestibular pathways. What may seem as small and insignificant differences from a broad mammalian perspective becomes hugely important as scientists seek to understand the particular mechanisms of brain plasticity. Can they use animals with no foveae and a vestibular-colic response to learn about how foveated mammals recover their vestibular-ocular response? More generally, how well do particular animal models translate across the animal kingdom? Should scientists be allowed to generalize from experiments on a single species (or set of species) to how nature functions?

In all vertebrates, a unilateral labyrinthectomy (UL), or a lesion of the labyrinthine structure in one ear, gives rise to two types of ocular motor disorders. There are static deficits, such as a bias toward looking toward the lesioned side when the head is not moving, and dynamic deficits, such as abnormal vestibular-ocular reflexes (VOR), which occur in response to head movements. In only two or three days following the UL procedure, the brain starts to compensate for its loss and the static deficits disappear. Since labyrinthine structures do not regenerate, and peripheral neurons continue to fire abnormally, whatever the brain is doing to recover has to be a central effect. Single neuron recordings from a variety of animals indicate that the vestibular nuclei (VN) on the same side of the brain as the lesion start to show normal resting rate activity as the brain learns to compensate for its injury. Scientists do believe that whatever the mechanism is, it is also likely to be a general procedure the brain uses for recovery, for there are similar resting rate recoveries of the sort seen with the ipsilateral vestibular nuclei following denervation in the lateral cuneate nucleus, the trigeminal nucleus, and the dorsal horn, among other areas. Exactly how an argument to defend these convictions is supposed to run, though, is unclear, since it is fairly easy to find significant differences in how organism's recover and compensate for vestibular damage across the animal kingdom. Frogs, for example, appear to rely on input from the intact labyrinth to regulate the resting activity of the vestibular nuclei. Mammals, however, do not. The recovery of their vestibular nuclei occurs independent of transcommissural inputs. In addition, static symptoms follow different time courses in different animals. In rats, spontaneous nystagmus disappears within hours after UL, while in the rabbit and guinea pig, it persists for several weeks. In humans, it may continue in one form or another for several years.

There is a fundamental tension in neuroscience between the big picture story and what we find in particular instances. All sciences strip away features of the

real world when they devise their generalizations. Physicists neglect friction; economists neglect altruism; chemists neglect impurities, and so on. However, what neuroscientists are doing is not analogous what the physicists, economists, and chemists are doing. In each of the other cases, the scientists are simplifying the number of parameters they must consider in order to make useful and usable generalizations. In contrast, if neuroscientists were to ignore the differences they find across species, then they would have no data left to build a theory with. There is nothing left over, as it were, once neuroscientists neglect the anatomical and physiological differences found in the brain across the animal kingdom. There is much left over when physicists neglect friction; most of classical mechanics is left, in fact. In distinction to the other sciences, there is a tension in neuroscience between the general rules one hopes to find that describe all brains and the particular cases neuroscientists happen to study.

What should the scope and degree of generalization for neuroscientific theories be? It is an unpleasant choice. Either scientists settle for large-scale abstract generalizations, which gloss over what may be important differences, or they focus on the differences themselves, at the expense of what may be useful generalizations. However, despite appearances, it is not an either-or proposition that has to be resolved before scientists can move ahead. For a proper neuroscientific theory contains both general (and fairly vague) abstractions as well as detailed comments on specific anatomies and physiologies. The paradigm theories for physics are simple elegant equations with universal scope. Theories in neuroscience read more like a list of general principles plus detailed commentaries. One feels the tug of the dilemma posed above only if one is operating with a restricted notion of what a scientific theory is. Some theories are pithy and succinct; some are not. Neuroscientific theories are not.

In neuroscience, what scientists start with is a theoretical description at the most general level; it might be called the "theoretical framework" — the most general component in a neuroscientific theory. Once they adopt the framework, they can make more precise hypotheses as a way of filling out our theoretical proposal. These claims can be local to particular phyla or species; hence, they are not intended as a more detailed specification of the general framework. Instead, they can be thought of as instances or examples of how the framework might be cashed out in particular cases.

However, it is not the case that all "fillings out" fail to generalize. For example, the dynamic symptoms of unilateral labyrinthectomy recover using a different mechanism (probably). One hypothesis is that brains use a form of sensory substitution to compensate for the vestibular-ocular reflex. In this case, the brain uses internally generated signals from the visual or somatosensory systems to compensate for the vestibular loss. It may substitute computations from the saccadic or a visual pursuit system, both of which (probably) reconstruct head velocity internally, for vestibular throughputs. Data drawn from experiments on frogs, cats, and humans indicate that they all apparently use the same mechanism, though it

remains to be seen whether this proposal will be applicable to all creatures and whether it can be generalized much beyond vestibular reflexes.

There are different degrees of abstraction one might use once some theoretical framework is adopted. Some discussions are going to be restricted to a single species, or maybe even one developmental stage within a species; others will include several unrelated species or phyla. Both are legitimate ways of cashing out the framework in particular instances, and neither is to be preferred over the other. The data will dictate the scope of sub-hypotheses, and scope can vary dramatically.

And this is how theories in neuroscience are built and structured. Detailed conclusions regarding a single animal model give rise to general theoretical principles. These principles inspire new experiments done with other animal models, which in turn give us new (and probably incompatible) details but also new general principles. These new principles then connect to other detailed studies using different protocols on still other animals, and so it goes.

At the end of the day, there is a set of related theoretical principles that jointly comprise a general theoretical framework. And these principles are held together by the detailed data coming out of a wide variety of animal studies. Neuroscience continually moves between two different ways of understanding the nervous system, first in broad and sweeping strokes and second by being submerged in the minutiae. General theoretical principles arise out of and then feed back into particular animal experiments done on different animal models. Because physiology differs across species, specific experimental protocols are appropriate only for specific models. Sometimes the data arising out of the different animal models are different experimental procedures overlap, but largely they do not. Hence, sometimes the detailed conclusions are consistent, but sometimes — a lot of the time — they are not. Neuroscientists weave a story through their animal models and experimental protocols united by a common guiding theoretical thread. They both find commonalties and define differences. And this entire exercise, taken together, fashions the theoretical structure of neuroscience.

2.2 Localization and Reduction

When scientists do single unit recordings from a set of neurons they assume that they are busy examining a discrete system. They have been wildly successful using this strategy, identifying at least 36 different topographical visual processing areas in cortex [De Gelder, 2000], differentiating the "what" from the "where" object processing streams ([DeYoe and Van Essen, 1988]; [Mishkin *et al.*, 1983]), and distinguishing motion detection from contour calculations [Barinaga, 1995], to name but a few examples. Maps of brain function are getting more and more complicated as more and more is learned about the processing capacities of individual cells. And all these projects are founded on the belief that brains have discrete processing streams that feed into one another.

Yet the most neurons scientists have ever been able to record from simultaneously are around 150; the most cells they can ever see summed local field potential

activity over are a few thousand. But brain areas have hundreds of thousands of neurons, several orders of magnitude more than they can access at any given time. And these neurons are of different types, with different response properties and different interconnections with other cells, including other similar neurons, neurons with significantly different response properties, and cells of other types completely. Any conclusions they draw about the behavior of whatever cells they are recording from are going to be limited to very basic stimulus-response and correlation analyses of whatever neuronal subtype they are currently examining. Hence, the functionality they ascribe based on these relatively meager sorts of experiments might be much more restricted than what the cells are actually doing.

They insert an electrode in or near a cell and then record what it does as they stimulate the animal in some fashion. They record from a cell in a vestibular nucleus and then move the animal's head about to see if that changes the activity of the neuron. If it does, then they move it some more or they move it differently and see how that changes the neuronal output. If it does not, then they either try another nearby cell or they try some other stimulus. But what they cannot do is record from all the neurons in some isolated area, even if the area is very small. And what they cannot do is test any given cell for all the known functional contributions of brain cells in general. So, what they conclude about any cell will only reflect the cells they or others have actually recorded from using stimuli they or others have actually used. This research strategy systematically underestimates when neurons actually respond and under what conditions.

This sort of unit study attempts to combine scores, hundreds, or even thousands of single-unit recordings together to try to analyze the population. Theoretically, they could, perhaps, in principle, delineate a nervous system region stereotaxically if it had reproducible correlations between afferent and efferent connections such that they could ultimately articulate the neurophysiological function of the defined region. However, the likelihood of success for this type of study decreases as the complexity of the organism increases. They can draw functional conclusions regarding the activities of neurons in the abdominal ganglia of *Aplysia*, or the segmental ganglia of the leech. But the architecture of these organisms' central nervous system is so different from mammals' that the probability of successfully using similar techniques for understanding humans is very low.

In addition, the actual processing of information that goes on in those cells involves lots of different kinds of excitatory and inhibitory inputs from other areas in the brainstem, cerebellum, and cerebral cortex. The dorsal horn is supposed to integrate afferent nociceptive information from the periphery and pass it onto the motor system (among other things), but it does not do that segregated from the rest of the brain and what the brain is trying to do. It is integrating and passing as we are trying to pursue prey or flee from an enemy. Moreover, the brain regions that perform these tasks are often connected to the very area scientists are recording from. The motor system feeds back down into the dorsal horn, as does the thalamus and significant parts of cortex.

The impact on cognitive processing of such rampant feedback connections in the

brain is only just now starting to be explored in neuroscientific research, though exactly how to do this remains a difficult question. Neurophysiologists design their experiments keeping in mind the known anatomic connections between and among the relevant structures. At the same time, any actual experimental observations of all the remote influences on the dorsal horn, for example, are impossible, despite however many individual neurons scientists record from. They simply do not have any way of conducting such extensive, invasive tests on live animals. At best, the particular influences assumed in any particular recording series are a matter of previously accepted gospel, dogma, and faith.

Ideally, neuroscientists try to conjoin their single-cell studies with some sort of lesion experiment. Once scientists construct a general flowchart of the relevant structures based on anatomy experiments, and they have estimated normal unit behavior from a series of single cell studies, they then try to knockout the hypothesized functions by placing lesions in otherwise normal animals. They run their experiments based on the assumption that these lesions, placed in regions known to be important, will change the unit behavior of cells they are studying in a consistent fashion. If they witness such a change, they use that information to explain the relative functional contributions of the lesioned region to the cells under scrutiny. In other words, they are using lesion studies to try to derive a functional boxology for the brain, just as cognitive psychologists use reaction time distributions and error measurements to find one for the mind.

But there is a larger theoretical concern. What neuroscientists know, but generally ignore, is that any functional change in the central nervous system will lead to compensatory changes elsewhere. Because it is highly plastic, lesioning the brain in one place will provoke to it compensate elsewhere (e.g., [Merzenich *et al.*, 1983]). Usually these other places are not components in the system or region being studied. But even if they are, neuroscientists ignore plasticity of the brain in favor of assuming a consistent functional alteration as caused by the lesion and nothing more. How are investigators supposed to evaluate some observed functional change when the difference they see might have been evoked by the brain's attempt to compensate for its loss and not by any specific deficit induced by the lesion?

The short answer is that they cannot if they are restricted to single-cell recordings and lesion studies. To answer this question they need to be able to see the activity of the entire brain at once and over time. The excitement over fMRI and other imaging techniques concerns exactly this point: they do have a way of looking at the activity of the whole brain at one time as tied to some cognitive activity or other. But magnetic resonance imagining, the best non-invasive recording device currently available, only has a spatial resolution of about 0.1 millimeter and each scan samples about a half second of activity. This imprecision forecloses the possibility of directly connecting single cell activity — which operates three to four orders of magnitude smaller and faster — with larger brain activation patterns.

Here is how most functional imagining studies work. Experimenters pick two experimental conditions that they believe differ along with respect to the cognitive or perceptual process under investigation. They then compare brain activity

recorded under one condition with what happens in the second condition, looking for regions whose activity levels differ significantly across the two. These areas, they believe, comprise the neural substrates of the task under scrutiny.

Let us set aside the fact that this so-called subtraction method has no way of determining whether the differences found are actually tied to the cognitive process and not to something else occurring concurrently but coincidentally. Methodological difficulties with current imagining techniques are now well known and shall not be rehearsed here (though see [Bechtel, 2000]; [Cabeza and Nyberg, 1997, 2000]). Notice that how well the subtraction method will work depends upon the sensitivity of the measuring devices, and therefore that the worse the instrument is the better the method seems to be for localization studies. Low signal-to-noise ratios (SNR) means that scientists will find only a few statistically significant differences across conditions. And these are the sorts of results neuroscientists need in order to bolster any claims identifying particular cognitive processes with discrete brain regions.

But as the imagining technology improves and the SNR increases, more and more sites will differ across trials. The more sites that differ, the more it looks as though essentially the entire brain is involved in each cognitive computation. And the more it looks as though the entire brain is involved in each thought, the any assumption of functional specificity in the brain can be justified.

Neuroscience is a victim of imprecise instrumentation. If scientists extrapolate from what they might learn with more sensitive measures, it can easily be seen that there will come a time when this whole approach just will not work anymore. Put in the harshest terms, brain imaging seems to support reductionism because the imaging technology is not very good yet.

For example, Brodman area 6 appears significantly and differentially active after subtraction in studies of phonetic speech processing, voluntary hand and arm movements, sight-reading music, spatial working memory, recognizing facial emotions, binocular disparity, sequence learning, idiopathic dystonia, pain, itch, delayed response alternation, and category-specific knowledge, to list only a subset of activities. It could be the case that if neuroscientists keep on doing the sort of subtraction studies that they currently are doing, then eventually they will find a unifying and pithy way to describe what premotor cortex is doing. In this instance, neuroscience would be on the right track to determining brain function, but scientists still have a long way to go yet. But it could also be true that how a region functions depends heavily on the "neural context". Its functional role in a cognitive economy depends on how it is connected to other areas and how those other areas are responding. (The function of these areas would also be dependent on their particular connectivity and the current patterns of activation. And so it would go.) If this is correct, then searching for "the" function of particular areas is misguided, for different brain regions play different roles depending upon the cognitive tasks at hand.

2.3 Theory-Laden Observations and Single-Cell Recordings

It is almost a truism in philosophy of science that there is no real distinction between observation and theory. That is, all scientific observations are filtered through and by a prior theoretical framework. Raw data become observations as they are interpreted regarding how they either fit or belie our hypotheses [Woodward, 1989]. In short: what counts as an observation and how that observation functions in the business of science is heavily mediated by theory. In neuroscience in particular, it is easy to change the fundamental nature of observations using accepted methodological techniques for manipulating raw data. The line between good data tinkering and fudging the data is quite thin.

Good data allow scientists to discriminate among competing claims about phenomena [Suppe, 1989]. The particular practices of the scientific sub-field define how to judge whether data are good. Sometimes these practices involve explicit calculations and formal derivations; sometimes they involve matters of personal judgment and skill. The cases in neuroscience involve both. In particular, it is a matter of personal judgment in the world of single-cell recording when to employ certain computational procedures. Different sorting techniques give rise to different data, so which techniques to employ is an important question. But it is not a question for which any good algorithm exists to answer it.

In 1791, Luigi Galvani demonstrated the link between neural communication and electrical impulses when he stimulated frogs' legs with electricity and made them twitch. But not too much could be made of this discovery until the 1920's, when scientists developed the technology that allowed them to measure nerve impulses directly using amplified signals sent via electrodes. Not surprisingly, single cell recording devices have improved much since then. Most importantly, perhaps, in 1968 E.V. Evarts developed techniques for recording from single nervous cells in alert moving animals. But it has only been during the last decade or so that neuroscientists been able to record from the extracellular space of a large number of neurons from awake and behaving animals.

When scientists record with an electrode near a single cell, they do pick up the cells' action potentials, which is what most people think of when they think of neuronal communication. But they also record things that look like action potentials, but are actually voltages generated by axonal bundles or the field potentials from parallel sets of dendrites. Moreover — and especially if the microelectrode has a relatively low impedance — extracellular electrodes pick up signals from several neurons at the same time, recording from all the cells in a nearby area.

The problem is how to differentiate the contributions of the different cells and cell parts when all we have is a single lump recording. In many cases scientists only care about one particular action potential; the rest, from their perspective, is background noise. The challenge is how to separate what they want from all the electrical signals they do not want. The challenge is how to move from the recordings of the electrode's output to genuine, reliable, and informative data.

This challenge is compounded by the noisy nature of the recordings themselves.

Some of the noise is mechanical and comes from the amplifiers themselves, but some is biological and comes from the neurons. Brain cells jitter around constantly (cf., [Connors and Gutnick, 1990]). Neurons are not quiet until they fire off a spike, as some might think. Instead, they are always producing some activity or other. All in all, scientists have to cull their data from quite a din.

Finally, because scientists cannot assume that anything in a recording remains constant, it is difficult to get a theoretical hook into the waveform. Spike shapes can change over time, electrodes can drift during recording session, changing position relative to the cells, which would also alter the spike amplitudes, and the electrical properties of electrodes vary with changes in tip condition or background impedance. Gathering data from single unit activity presents neuroscientists with a serious technical challenge.

In order to get useable data — to get genuine observations — out of what the electrode transmits, scientists must isolate each neuron's contributions to the recorded waveform. They first need to ascertain exactly how many neurons the recorded waveform reflects. How can they do this if they have a mess of overlapping action potentials and field potentials from a variety of cells at different and unknown distances from the electrode? This question becomes particularly vexing if other neurons in same area have spikes of the same or similar shape and amplitude.

There are several decomposition algorithms, albeit imperfect ones (see [Lewicki, 1998]). Each represents a different way to move from raw output to interpreted and interpretable data, giving scientists different ways of refining the waveforms they have recorded so that they can later interpret them. Each is what philosophers are thinking about when they talk about the theory-ladenness of data. Scientists have to choose what to do with their measurements in order to get something that can be scientifically useful. And how they choose is determined by previously accepted theories.

But even with all these advanced sorting techniques, it is still hard to predict the number of neurons eliciting the data. Ideally, scientists would like to claim that one neuron generates each cluster of spikes they have identified, but if the cells are firing in complex bursts, or if there is non-stationary noise, or if the spike trains overlap one another, they cannot get accurate classifications at all. It is simply an unsolved problem how to decompose coincident action potentials with variable spike shapes. The best scientists can do at this point is guess. Their guesses are informed by their years of experience, but they are guesses nonetheless.

Guessing is not quite what philosophers of science have in mind when they talk about the theory-ladenness of observation. Their vision of creating data is one of more "scientific method". That is, to pull data out of the dial movements or changes in color or squiggles on the page, philosophers generally hold that there is some explicit background theory, devised in some other scientific inquiry, that scientists learn and then use to interpret what they are seeing or measuring as something useful for their studies. But there is a theoretical gap in the move from raw recordings to genuine data, a gap that cannot be filled with any sort of

decision-making algorithm. The best scientists can do at this point is simply leap across the gap, on blind faith, with an eye to where they want to go.

Neurophysiology travels in a cognitive circle; scientists use what they know to cull data that support what they believe to be the case. Nevertheless, progress is not stymied. Knowledge accrues in small increments, with each set of single cell recordings altering the face of what is known a wee bit at a time. Because neurophysiological sorting techniques rely so heavily on previously accepted neurophysiological hypotheses, there will likely never be an abrupt or dramatic conceptual revolution. But what is known can evolve slowly but surely until the final resting position is quite far removed from where it started.

2.4 Representation in the Brain

Important to keep in mind is that brains are evolved products — they were formed to help animals feed, flee, fight, and reproduce. This contrasts with digital computers, say, which are designed to compute. Early cognitive science generally took minds to be importantly analogous to computers and so tried to build theories of representation that spanned both what computers and humans do. Contemporary neuroscientists think of representation in terms of brains only.

Most operate implicitly under the assumption that individual neurons are the representational engine that drives our brains. This assumption is not universally accepted in neuroscience, but it does provide a good starting point nonetheless. As I have already explained, the behavior of individual neurons is extremely complicated. Hence, many working in the area of brain representation carry out their research looking at artificial neural nets. They carry the advantage of being much simpler and easier to control than what nature has provides, so it is much easier to design and carry out experiments on them. They have the disadvantage of being much simpler and easier to control so that scientists cannot be sure that the behaviors we get out of them are relevantly similar to what is seen in the brain. But what representations look like in them are multi-dimensional phase spaces, something quite different from what traditional philosophers envision when they speak about representations. Whether a phase-space approach will replace traditional approaches to representation in philosophy remains to be seen.

Part and parcel of the problem of understanding representations is understanding learning, understanding how we get representations in the first place. Brain organization is remarkably constant across structures. Cortex is cortex is cortex. Therefore most specialized areas have to be carved out via experience. The postulated mechanism for producing such changes is nothing more than Hebbian learning, a mechanism first articulated in the 1940s: repeated activation will cause future activation to be easy; decreased activation will make future responses more difficult [Hebb, 1949].

But to get directed learning one must combine a learning mechanism with some sort of reward system. Animals need a reason to repeat an event in order to learn about it. Things that feel good they repeat. Things that do not they do

not. Much is now known about reward networks, especially those involving fear conditioning in rats. Here in particular is where studies in the neurobiology of affect are relevant, for they give scientists some clues about how all sorts of reward-based learning might be going on in brains, though whatever the final story is, it will be unquestionably more complicated than what is known now.

Cognitive neuroscience is an intricate combination of several different research areas, each with their own regions of concern, bent on understanding the processes of an intricate organic machine. But they all work together — or at least in parallel — toward that common goal. Bound by common methodologies and common methodological difficulties, the subdisciplines that comprise cognitive neuroscience are making tremendous headway. It is my sincere hope that the complete story of how a brain thinks will be known in my lifetime. Right now, that hope does not seem unreasonable.

BIBLIOGRAPHY

[Barinaga, 1995] M. Barinaga. Remapping the motor cortex. *Science*, 268: 1696–1698, 1995.

[Bechtel, 2000] W. Bechtel. From imaging to believing: Epistemic issues in generating biological data. In R. Creath and J. Maienschein (eds.), *Epistemology and Biology*. Cambridge, MA: Cambridge University Press, 138–163, 2000.

[Bechtel *et al.*, 2001] W. Bechtel, P. Mandik, J. Mundale and R. Stufflebeam (eds.). *Philosophy and the Neurosciences: A Reader*. New York: Blackwell Publishing, Inc., 2001.

[Cabeza and Nyberg, 1997] R. Cabeza and L. Nyberg. Imaging cognition: An empirical review of PET studies with normal subjects. *Journal of Cognitive Neuroscience*, 9: 1–26, 1997.

[Cabeza and Nyberg, 2000] R. Cabeza and L. Nyberg. Imagining and cognition II; An empirical review of 275 PET and fMRI Studies. *Journal of Cognitive Neuroscience*, 12: 1–47, 2000.

[Churchland and Sejnowski, 1988] P. S. Churchland and T. J. Sejnowski. Perspectives on cognitive neuroscience. *Science*, 242: 741–745, 1988.

[Clark and Hillyard, 1996] V. P. Clark and S. A. Hillyard. Spatial selective attention affects early extrastriate but not striate components of the visual evoked potential. *Journal of Cognitive Neuroscience*, 8: 387–402, 1996.

[Connors and Gutnick, 1990] B. W. Connors and M. J. Gutnick. Intrinsic firing patterns of diverse neocortical neurons. *Trends in Neuroscience*, 13: 99–104, 1990.

[Crick and Jones, 1993] F. Crick and E. G. Jones. Backwardness of human neuroanatomy. *Nature*, 361: 106–110, 1993.

[De Gelder, 2000] B. De Gelder. More to seeing than meets the eye. *Science*, 289: 1148–1149, 2000.

[DeYoe and Van Essen, 1988] E. A. DeYoe and D. C. Van Essen. Concurrent processing streams in monkey visual cortex. *Trends in Neuroscience*, 11: 219–226, 1988.

[Emiliano *et al.*, 2000] M. Emiliano, C. D. Frith and J. Driver. Modulation of human visual cortex by crossmodal spatial attention. *Science*, 289: 1206–1208, 2000.

[Evarts, 1968] E. V. Evarts. A technique for recording activity of subcortical neurons in moving animals. *Electroencephalography and Clinical Neurology*, 24: 83–86, 1968.

[Gallese *et al.*, 1996] V. Gallese, L. Fadiga, L. Fogassi and G. Rizzolatti. Action recognition in the premotor cortex. *Brain*, 119: 593–609, 1996.

[Goldman-Rakic, 1996] P. S. Goldman-Rakic. The prefrontal landscape: Implications of functional architecture for understanding human mentation and the central executive. *Philosophical Transactions of the Royal Society of London — Series B: Biological Sciences*, 351: 1445–1453, 1996.

[Hardcastle, 1995] V. G. Hardcastle. *How to Build a Theory in Cognitive Science*. Albany, N.Y.: SUNY Press, 1995.

[Hebb, 1949] D.O. Hebb. *The Organization of Behavior*. New York: Wiley, 1949.

[Kandel *et al.*, 2000] E. R. Kandel, J. S. Schwartz and T. J. Jessell. *Principles of Neural Science, 4^{th} Edition*. New York: McGraw-Hill, 2000.

[Lewicki, 1998] M. S. Lewicki. A review of methods for spike sorting: The detection and classification of neural action potentials. *Network: Computational Neural Systems*, 9: R53–R78, 1998.

[Livingstone and Hubel, 1988] M. Livingstone and D. Hubel. Segregation of form, color, movement, and depth: Anatomy, physiology, and perception. *Science*, 240: 740–749, 1988.

[MacNeilage, 1998] P. F. MacNeilage. The frame/content theory of the evolution of speech production. *Behavioral and Brain Science*, 21: 499–511, 1998.

[Marcus, 2004] S. J. Marcus (ed.). *Neuroethics: Conference Proceedings*. New York: Dana Press, 2004.

[Martin *et al.*, 2000] A. Martin, L. G. Ungerleider and J. V. Haxby. Category specificity and the brain: The sensory/motor model of semantic representations of objects. In M. Gazzaniga (ed.) *The New Cognitive Neurosciences: Second Edition*. Cambridge, MA: The MIT Press, 1023–1036, 2000.

[Merzenich *et al.*, 1983] M. M. Merzenich, J. H. Kaas, M. Sur, R. J. Nelson and D. J. Felleman. Progression of change following median nerve section in the cortical representation of the hand in areas 3b and 1 in adult owl and squirrel monkeys. *Neuroscience*, 10(3): 639–665, 1983.

[Mishkin *et al.*, 1983] M. Mishkin, L. G. Ungerleider and K. A. Macko. Object vision and spatial vision: Two cortical pathways. *Trends in Neuroscience*, 6: 414–417, 1983.

[Moran and Desimone, 1985] J. Moran and R. Desimone. Selective attention gates visual processing in the extrastriate cortex. *Science*, 229: 782–784, 1985.

[Quirk and Wilson, 1999] M. C. Quirk and M. A. Wilson. Interaction between spike waveform classification and temporal sequence detection. *Journal of Neuroscience Methods*, 94: 41–52, 1999.

[Sakata *et al.*, 1997] H. Sakata, M. Taira, M. Kusunoki, A. Marata and Y. Tanake. The TINS Lecture: The parietal association cortex in depth perception and visual control of hand action. *Trends in Neuroscience*, 20: 350–357, 1997.

[Schaffner, 1993] K. F. Schaffner. *Discovery and Explanation in Biology and Medicine*. Chicago: University of Chicago Press, 1993.

[Stufflebeam and Bechtel, 1997] R. S. Stufflebeam and W. Bechtel. PET: Exploring the myth and the method. *Philosophy of Science: Supplement*, 1997.

[Suppe, 1989] F. Suppe. *The Semantic Conception of Theories and Scientific Realism*. Chicago: University of Illinois Press, 1989.

[Woodward, 1989] J. Woodward. Data and phenomena. *Synthese*, 79: 393–472, 1989.

COMPUTATIONAL NEUROSCIENCE

Chris Eliasmith

1 OVERVIEW

'Computational neuroscience' has come to denote two significantly different disciplines. On the one hand computational neuroscience is sometimes considered a part of 'bioinformatics' or 'computational biology,' a field that is concerned with the application of computers to the collection, organization, and analysis of biological data. Perhaps the best-known example of this kind of work is the recent human genome project, where the sequencing of genes was largely done automatically using computers. On the other hand, 'computational neuroscience' is taken to refer to what has also been called 'theoretical neuroscience'. In this sense, computational neuroscience is the application of quantitative theories relating to computation and information processing to neurobiological systems. Notably, "relating to computation" in this definition is intended to be weak enough so that even those who wish to claim that the brain is not computational, but who use mathematical formalism (e.g., sets of differential equations), are taken to be engaged in computational neuroscience. It is this second sense that I am primarily interested in here.

I take the distinction between computational and experimental neuroscience to be analogous to that between theoretical and experimental physics. But, as the difference in terminology suggests, computational neuroscience relies heavily on computer simulation and modeling and so is not, perhaps, as 'purely' theoretical. To date, this most often means that particular neural systems have been subjected to detailed mathematical analyses which are often tested through numerical simulations. For example, the tools of nonlinear systems theory have been applied to systems of coupled oscillators in an attempt to understand the neural mechanisms underlying lamprey and leech locomotion [Wilson,1999, ch. 13] Such systems are then often simulated to determine or demonstrate the effects of certain parameter regimes. These same sorts of mathematical tools are more often, and similarly, applied to a much lower level of description of neural systems: the characterization of the dynamics of action potential generation[1] in single cells [Wilson,1999, ch.

[1]Neural action potentials are the stereotypical, rapid voltage changes induced in neurons as a result of their being excited past some threshold. The production of these potentials is also refered to as the neuron 'firing'. These changes are so rapid, that they are often considered instantaeous 'spikes' of voltage. A series of such spikes is called a 'spike train' (e.g., the black dots in Figure 1).

Handbook of the Philosophy of Science. Philosophy of Psychology and Cognitive Science
Volume editor: Paul Thagard
General editors: Dov M. Gabbay, Paul Thagard and John Woods
© 2007 Elsevier B.V. All rights reserved

5-10] In large part this is because there has been a great deal of success (ever since Hodgkin and Huxley's [1952] famous mathematical description of the giant squid axon) at applying such tools to understanding the dynamics of ion conductance in single cells. More recently, researchers have begun to construct detailed, biologically plausible models of large systems of interconnected neurons. Building such complex networks often makes mathematical analysis intractable, demonstrating the indispensability of computers for simulating and analyzing these models. Much of this work can be seen as a direct descendant of connectionism (see [Elman, 1998] for a brief history and discussion of connectionism).

Given the analogous role that computational neuroscience plays as compared to theoretical physics, it should come as no surprise that computational neuroscience has significant consquences for conceptual debates about the mind/brain, just as theoretical physics has significant conceptual consequences for theories of what there is. From a philosophical point of view, computational neuroscience stands to enter debates regarding mental representation, mental categories, the relation between higher-level (cognitive) and single neuron descriptions of the brain, and most other debates relating to a naturalistic understanding of cognition. In subsequent sections, I describe some results from computational neuroscience that touch on these and other metaphysical questions. First, however, let me present some brief historical and conceptual background.

As recently as 15 years ago, computational neuroscience was referred to as a "new and fledgling" discipline [Koch and Segev, xi]. However, computational neuroscience has historical roots that reach as far back as those of connectionism, at least to the beginning of the last century. It was, nevertheless, about 15 years ago when computational neuroscience split from its connectionist roots. This was about the same time at which the field of 'neural computing' split from its connectionist roots as well. In fact, a similar split is evident in cognitive science, with artificial intelligence (AI) and cognitive psychology (which has largely kept the 'cognitive science' label) focusing on different research problems. In both cases, one branch (AI/neural computing) has become mainly concerned with technical applications of the theory, interested in solving difficult engineering problems, while the other branch (cognitive science/computational neuroscience) has become concerned with the application of the theory to understanding natural systems.

So, in contrast with neural computing, in computational neuroscience a premium is placed on biological realism. That is, models are not constructed merely to realize some function (however useful that function might be from an engineering perspective), but rather to realize a particular function *as it is realized* in some biological system. Computational neuroscientists are not interested in how one might implement working memory in a collection of identical, interconnected, computational nodes, but rather in how working memory is actually implemented by real, heterogeneous neurons in a living brain. It is already well-known that there are multiple ways in which collections of neuron-like units can be made to store memories. The pressing question for computational neuroscientists is: Which of those ways is the one relevant for neurobiology? So, the most obvious connection

between computational neuroscience and cognitive science is their shared interest in trying to come up with naturalistic explanations of the behaviour of complex natural systems.

Nevertheless, the central methodologies of these approaches are vastly different. Computational neuroscientists take as paradigmatic, data from experiments in single cell physiology (e.g., intercellular or extracellular microelectrode recordings, dendritic patch clamps, multielectrodes) and neuroanatomy (e.g., cellular and tract staining, immunocytochemistry, tract tracing, deoxyglucose uptake). Cognitive scientists, in contrast, tend to consider data from a wide variety of multi-subject psychology studies (e.g., studies that measure reaction time, gross motor behaviour, linguistic behaviour, etc.). More recently, neuroscience-related techniques like functional magnetic resonance imaging (fMRI), positron emission tomography (PET), and event related potentials (ERPs) have begun to inform cognitive science. The biggest difference between the standard methodologies in cognitive science and those standard in computational neuroscience is their scale: computational neuroscientists have focused on the microscopic and cognitive scientists have focused on the macroscopic. This has largely been because, until recently, the kind of mathematical descriptions provided by computational neuroscientists have not been "scaled up" to be informed by the data at the macroscopic level.

Because of these methodological differences, cognitive science and computational neuroscience also tend to focus on different kinds of behaviours as targets of explanation. Cognitive scientists, not surprisingly, focus on largely cognitive phenomena including: reasoning, object recognition, imagery, and language processing. Computational neuroscientists focus on more universally biological phenomena including: motor control, low-level perception (e.g., proprioception, retinal and early visual processing), and self-organization (neuron-level learning). As a result, implementational issues (i.e., issues like how many neurons do you have in working memory systems; what signal to noise ratio is supported by neural systems; how interconnected are the neurons; what are the time constants of dendritic potentials observed in working memory; etc.) have proven much more significant to computational neuroscientists.

From the perspective of a cognitive scientist, it might seem that a preoccupation with low-level implementational issues is a good way to waste your time just trying to get the irrelevant details right — details that have nothing to do with cognition [Fodor and Pylyshyn, 1988]. But perhaps computational neuroscientists agree with Mies van der Rohe who suggested that "God is in the details". That is, perhaps the designs discovered by mother nature are much more effective, robust, efficient, and so on, than those invented by a clever engineer. Perhaps a functional decomposition that is sensitive to the same design constraints as mother nature will look very different from one that is not. If so, time might be well spent looking at the details, no matter how devilishly obscure they may be. After all, we have proof that mother nature is able to construct sophisticated, robust, and extremely complex systems — systems exquisitely designed to adapt to rapidly

changing, dynamic environments. Unfortunately, the same cannot be said of human designed systems. Similarly, we know that mother nature's design can scale up appropriately, we do not know that for our own designs. So, since we seem to be at the beginning of an era where the tools for reverse-engineering the brain are becoming available, it would be wise to exploit those tools.

Some exploitation of these tools has already taken place (indeed, information theory has been used since its inception to understand neural systems). However, most work in computational neuroscience has focused on improving our understanding of single cells. And, while the success there has been impressive — there are general methods for modeling both the electrical and morphological properties of individual neurons — the most remarkable behaviours demonstrated by natural neurobiological systems stem from the organization of billions of such cells. It is only when we begin considering large groups of neurons that we approach the point of contact between neuroscience and cognitive science. So, major challenges lie in trying to apply mathematical tools to neural systems of a sufficient size to tackle issues of interest to those concerned with cognitive function [Eliasmith, 2003].

In the remainder of this article I address some methodological, metaphysical, and conceptual issues that arise when we begin to consider the integration of higher-level (cognitive) and lower-level (neural) descriptions of mental phenomena. More specifically, I discuss accounts of representation, computation, dynamics, and cognition from a computational neuroscience perspective. This discussion, I suggest, shows that computational neuroscience is indispensable for generating a complete naturalistic characterization of mental/cognitive function.

2 REPRESENTATION

The concept of a 'mental representation' is a central one in cognitive science [Dietrich, this volume]. As a result, much time has been spent trying to better understand, analyze, and operationalize the notion of representation. Debates in cognitive science on what kinds of representations are used for what cognitive tasks [Kosslyn, 1994], how to determine representational content [Churchland, 1989; Fodor, 1990; Fodor and Lepore, 1992], and even on the relevance of representation to cognition [van Gelder, 1998] are currently ongoing. As a result, it is essential to address what, if anything, computational neuroscience has to say about representation.

One useful distinction out of the philosophical discussion of representation is that between the contents and the vehicles of representations [Fodor, 1981; Cummins, 1989]. The contents, or meaning, of a representation is the semantic value of a representation for an animal. In general, contents are thought to be determined by the object in the world that the representation picks out, the relation of that representation to other representations, or a combination of both. The vehicles of representations are the syntactic structures, or physical realization of objects that play the role of representations (i.e., carrying a content) in a cognitive system. Currently, computational neuroscience can contribute most to improving

our understanding of representational vehicles. Nevertheless, I discuss some of the consequences of recent work in computational neuroscience for our understanding of semantics as well.

Unfortunately, in neuroscience the concept of representation itself has received much less consideration than in the cognitive sciences, despite an equally wide usage.[2] As a result, the use of the term in neuroscience is often problematic. In general, if a neuron 'fires' relatively rapidly when an animal is presented with a certain set of stimuli, the neuron is said to "represent" the property that the set of stimuli share (see, e.g., [Felleman and Van Essen, 1991]. This kind of experiment has been performed since Hubel and Wiesel's [1962] classic experiments in which they identified cortical cells selective to the orientation and size of a bar in a cat's visual field. The 'bug detector' experiments of Lettvin *et al.*. [1988/1959], perhaps better known to philosophers, take a similar approach. In the 'bug detector' experiments, retinal ganglion cells were found that respond to small, black, fly-sized dots in a frog's visual field. More recently, this method has been used to find 'face-selective cells' (i.e., cells that respond strongly to faces in particular orientations) in monkey visual cortex [Desimone, 1991]. In all of these cases, what is deemed important for representation is how actively a neuron responds to some known stimuli.

The difficulty is that such usage assumes that single neurons are the carriers of content, and that content can be determined by what has been called the 'naïve causal theory' — a view well knwon to be highly problematic [Dretske, 1988]. So, there are no principled means of determining what the representational vehicles are, or how they might be related.

Nevertheless, much of the work in *computational* neuroscience can help refine the notion of representation such that it is clearly relevant for understanding neurobiological systems and avoids the difficulties of these assumptions. This is partly because, unlike the typical treatment of representation in cognitive science, considerations of representation in computational neuroscience cannot avoid implementational constraints on representation. This proves to be a strength because such constraints highlight significant aspects of representation that have otherwise been missed (e.g. time, precision, etc.).

One of the most significant conceptual contributions of computational neuroscience to a neuroscientific understanding of representation is its emphasis on decoding. As mentioned, characterizing the responses of neurons to stimuli in the environment has been the mainstay of neuroscience. This, however, describes only an encoding process. That is, the process of encoding some physical environmental variable into neural action potentials, or 'spike trains'. By adopting an information theoretic view of representation, computational neuroscience holds that if we truly understand the encoding process, we must be able to demonstrate that we

[2] A search of PubMed for the term 'representation OR representations OR represent' returns well over 130000 hits. So far this year, in the Journal of Neuroscience alone, 8% of publications (42/543) mention representation in the title or abstract.

can decode a spike train to give us the originally encoded signal. As a result, to fully define representations, we must understand both encoding and decoding.[3]

As well, computational neuroscience has focused on two distinct aspects of representation: temporal representation; and population representation. The former deals with how neurons represent time-varying signals. The latter deals with issues of distributed representation (that is, how we should understand the contribution of a single cell's response to a complex representation over a large group of neurons). In the next two sections, I describe computational neuroscientific characterizations of encoding and decoding over time and neural populations.

2.1 Temporal representation

What is often considered one of the major flaws of classical cognitive science can simply not be ignored when considering neural systems: the importance of time. The extent and impact of this flaw for classical cognitive science has been vociferously argued by proponents of dynamicism [Port and van Gelder, 1995; van Gelder, 1998]. They have suggested that both symbolic and connectionist models often completely ignore time constraints, or include them only after the fact. This, they suggest, is completely unrealistic. Real systems are significantly constrained by the highly dynamic environments in which they are embedded. To ignore time is to ignore one of the most salient driving forces behind the impressiveness of cognitive systems, the limitations of cognitive systems, and the evolution of cognitive systems.

Indeed, neuroscientists, who are interested in exploring real world neural systems, have always been confronted with the essentially temporal nature of neurobiological systems. This is clearly reflected in much of the standard vocabulary of neuroscience: neuroscientists speak of spike times, firing rates, stimulus onset, persistent activity, adaptation, membrane time constants, and so on. So, unlike the often static considerations of representation of symbolicists and connectionists, neuroscientists have always considered time-varying representations. In this respect, any neuroscientific endeavour, be it experimental or theoretical, would be remiss if it ignored time. So it is no surprise that the representation of time-varying signals is central to computational neuroscience.

Perhaps the best understood aspect of how neural systems represent time-varying signals is the encoding process. In some ways, this should not be too surprising since the focus of neuroscience in general has been on encoding. This is because the encoding process can be characterized with respect to a single cell. So, the highly successful work on quantifying the dynamics of action potential generation in single cells — including mathematical descriptions of voltage sensitive ion channels of various kinds [Hodgkin and Huxley, 1952], the use of cable equations to describe dendritic and axonal morphology [Rall, 1957; 1962], and the introduction of canonical models of a large class of neurons [Hoppensteadt and Izhikevich, 2003] — supports a highly mechanistic understanding of encoding. However, this

[3]Note that this does not demand that the decoding take place explicitly within an animal.

mechanistic understanding cannot be focussed soley on the properties of individual neurons. Fully describing the encoding process also necessitates the identification of the particular, perhaps external, parameters a neuron may be sensitive to (partially in virtue of its relation to other neurons in the brain). These more holistic considerations are implicitly captured by the ubiquitously reported 'neuron tuning curves' in the neuroscientific literature. Improving our understanding of the encoding process is largely an empirical undertaking, one which has a long, successful history in both experimental and theoretical neuroscience. However, this is not true of temporal decoding.

There are two main kinds of theory of temporal decoding in neuroscience. These are referred to as the "rate code" view and the "timing code" view. Generally speaking, rate code theories are those that assume that information about temporal changes in the stimulus is carried by the average rate of firing of the neuron responding to that stimulus [Shadlen and Newsome, 1994; 1995; Buracas *et al.*, 1998]. In contrast, timing code theories assume that information about the stimulus is carried by the distance between neighbouring spikes in the spike train generated by the neuron responding to the stimulus [Softky and Koch, 1993; de ruyter van Steveninck *et al.*, 1997; Rieke *et al.*, 1997].

Many results in neuroscience are reported under the assumption that rate coding is an appropriate way to characterize temporal representation in neurons. For this reason, it is common to see firing rates, or spike histograms, of neurons over time. To generate these figures, neuroscientists usually determine the mean firing rate of a spike train over a relatively long, sliding time window (usually about 100 milliseconds). However, there are a number of problems with adopting such a view of neural representation. Perhaps the most obvious is that response times of many animals to a given stimulus are much quicker than 100 milliseconds. If the motor system was averaging over 100 milliseconds time window in order to determine what information was in the spike train arriving from sensory systems, such rapid responses would be impossible. In fact, there is a large volume of evidence that behavioural decisions are often made on the basis of one or two neural spikes arriving a few milliseconds apart (see [Rieke *et al.*, 1997, pp. 55-63] for a review). As well, there is evidence that spike trains with exactly the same average firing rate but different placement of spikes within the hundred millisecond window produce very different output from a neuron they are presented to [Segundo *et al.*, 1963]. And finally, it has been demonstrated that rate codes cannot transmit information at the rates observed in neurons [Rieke *et al.*, 1997], although the timing code can [MacKay and McCulloch, 1952].

Given these difficulties, many computational neuroscientists prefer to understand temporal coding in neurons as dependent on the timing of individual spikes. A standard way to characterize this code is to take the inverse of the inter-spike intervals in a neuron spike train. So, when the inter-spike interval is short, the encoded signal was high, and when the inter-spike interval is long, the encoded signal was low (although this is an over-simplification). However, it would be misleading to say that it is the *precise* time of spikes that carries information about the

stimulus in most such characterizations. There is good evidence that the precise timing of spikes is not mandatory for the successful transmission of signals with neurons [Bialek *et al.*, 1991]. In some ways, this is not a surprising result, since a code which was extremely sensitive to the timing of individual spikes would not be robust to noise.[4]

When we look carefully at the difference between typical rate codes and timing codes, we notice that they are variations on the same theme. Both codes assume that we choose some time window and count how many spikes fall in that window. In the case of rate codes the time window is about 100 milliseconds, and in the case of timing codes the size of the window varies depending on the distance between spikes. It should not be too surprising, then, that methods have been developed for understanding temporal decoding that vary smoothly between rate codes and timing codes [Rieke *et al.*, 1997]. So, in the end, the distinction between rate codes and timing codes is not a significant one for understanding temporal representation.

These methods are surprisingly simple, as they are linear (i.e. rely only on weighted sums). Suppose we are trying to understand the representational role of a particular neuron. To do so, we present it with a signal and then record the spikes that it produces in response to the signal. These spikes are the result of some (well-characterized) highly nonlinear encoding process. As mentioned, if we truly understand representational abilities of this neuron, we should be able to use those spikes to reconstruct the original signal. However, to do this we need to identify a decoder. We can begin by assuming that the particular position of a given spike in the spike train does not change the meaning of that spike; i.e., the decoder should be the same for all spikes. Essentially, every time a spike occurs, we place a copy of the decoder at the occurrence time of the spike. We can then sum all of the decoders to get our estimate of the original input signal (see figure 1). There are well-tested techniques for finding optimal decoders of this sort. As a testament to the effectiveness of these assumptions, surprisingly this kind of decoding captures nearly all of the information that *could* be available in the spike trains of real neurons [Rieke *et al.*, 1997, pp. 170-176].

A limitation of this understanding of temporal representation is that it is not clear how *our* ability to decode the information in a spike train relates to how that spike train is actually used by the organism. Recently, it has been suggested that the postsynaptic currents (PSCs) observable in the dendrites of receiving neurons can act as temporal decoders [Eliasmith and Anderson, 2003, ch. 4]. While the amount of information lost increases under this assumption, it is biologically plausible (unlike the acausal optimal decoders described earlier), and increasing the number of neurons in the representation can make up for any information lost. An example of this kind of decoding for two neurons is shown in figure 2a. That example shows how a rapidly fluctuating signal can be decoded from a neural spike train by a receiving neuron using a timing code (it is a timing code because significant jitter in the position of the spikes would greatly change the estimate).

[4]There seems to be a large variety of sources of noise in nervous systems [Eliasmith, 2001].

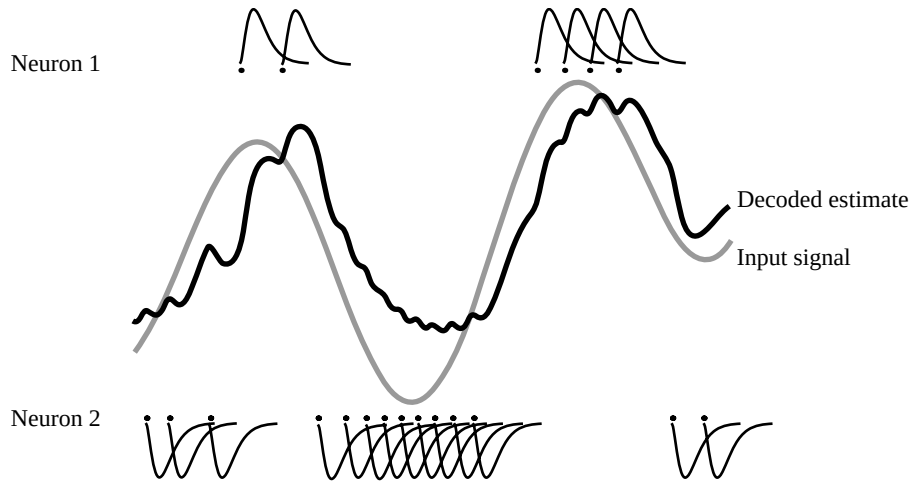

Figure 1. Temporal decoding. This diagram depicts linear decoding of a neural spike train (dots) using stereotypical decoders (skewed bell-shaped curves) on an input signal (grey line). The result of the decoding from two neurons (black line) is a reasonable estimate of the input signal. This estimate can be indefinitely improved with more neurons. (Adapted from [Eliasmith and Anderson, 2003].)

An example of decoding a much slower signal, which is encoded using something more like a rate code is shown in figure 2b.

Together, these diagrams show that linear temporal decoding with PSCs is biologically plausible and continuous between rate and timing codes, depending on the signals. This way of understanding temporal coding in neurons is both biologically plausible and applicable to understanding the representation of signals at a wide variety of time scales. However, the signals being represented are extremely simple. All of these examples are time-varying scalar values. To support representations of sufficient complexity to handle cognitive functions, we need to understand how large groups of neurons can cooperate to effectively encode complex, real-world objects.

2.2 Population representation

In a well-known series of experiments, Ausonio Georgopoulos explored the idea that the representation of physical variables in the cortex could be understood as a weighted sum of the individual neuron responses [Georgopoulos *et al.*, 1986; 1989]. By recording from a population of neurons in motor cortex, he demonstrated that a good prediction of a monkey's arm movement could be made by multiplying neuron firing rates by their preferred direction of movement and summing the result over the population. Essentially, Georgopoulos discovered a decoding method for extracting information carried by the neural firing rates that captured reasonably

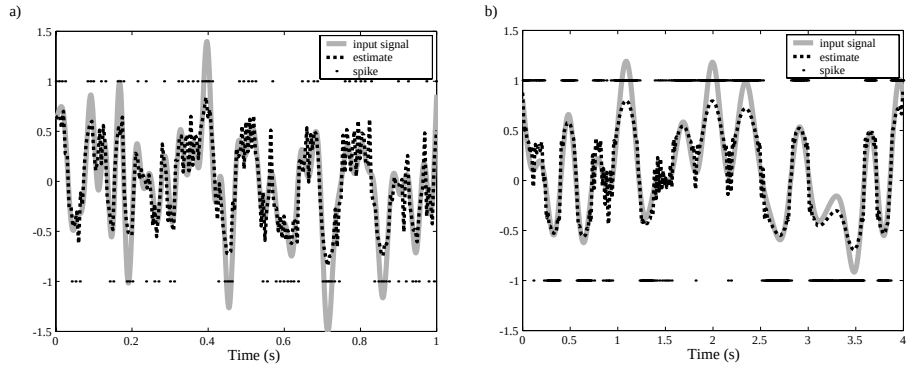

Figure 2. Biologically plausible timing and rate coding. a) A high frequency signal effectively decoded using postsynaptic currents (PSCs) as the decoders. This demonstrates a timing code. b) A low frequency signal (note the difference in time scale) similarly decoded. This demonstrates a typical rate code. (Adapted from [Eliasmith and Anderson, 2003].)

well how this information was used by the actual motor system.

As a result, it is generally agreed that Georgopoulos provided a demonstration of how to decode a scalar variable (arm angle) encoded by a population of neurons. This kind of decoding, we should notice, is identical to that described in the temporal case. It is a simple linear decoding where the temporal decoder is replaced by a population one (i.e., preferred direction). However, the particular decoding chosen by Georgopoulos is far from optimal. Nevertheless, it is a simple matter to determine the optimal linear decoder [Eliasmith and Anderson, 2003]. Furthermore, it is easy to generalize this kind of understanding of neural representation to more complex mathematical objects. For instance, instead of understanding neurons in motor cortex as encoding a one dimensional scalar (i.e., direction), we can take them to be encoding a two-dimensional vector (i.e., direction and distance of arm movements). Indeed, there is evidence that the neurons Georgopoulos originally recorded from carry information about both of these dimensions [Schwartz, 1994; Moran and Schwartz, 1999; Todorov, 2000].

However, scalar and vector representation alone is not sufficient for capturing a wide variety of neural responses. One of the most common kinds of tuning curves observed in cortex is a Gaussian-shaped 'bump' (sometimes called 'cosine tuning') around some preferred stimulus. For example, in lateral intraparietal cortex (LIP), neurons have these bump-like responses centered around positions of objects in the visual field [Andersen *et al.*, 1985; Platt and Glimcher, 1998]. On first glance, it may be natural to see them as encoding a scalar value which indicates the current estimate of the position of an object in the visual field. However, there is evidence to suggest that these representations are more sophisticated. For instance, the representation in this area can encode multiple object positions simultaneously,

and can have differing heights of bumps at those positions [Platt and Glimcher, 1997; Sereno and Maunsell, 1998]. And, although this has not been directly tested in LIP, there is evidence that this kind of representation is used to encode information about the uncertainty of the estimate of the parameter being encoded [Britten and Newsome, 1998]. So, this kind of representation seems best understood as the representation of a function, not a scalar. The width of this function can be used to encode the uncertainty of the representation (wider functions indicate greater uncertainty). And, more complex functions like bimodal Gaussians, could be used to encode position and certainty information about multiple objects simultaneously. So, it is natural to suggest that LIP is representing functions, not scalars or vectors.

It seems essential, given the noisy, complex, and uncertain environment in which neurobiological systems reside, for such systems to be able to make decisions with partial data. An increasing number of neuroscientists and computational neuroscientists have been taking seriously the idea that neural systems must perform some kind of statistical inference in order to operate effectively in such an environment [Eliasmith and Anderson, 2003; Kording and Wolpert, 2004; Rao, 2004; Deneve and Pouget, 2004]. So, it is only to be expected that the representations in neural systems be able to support statistical inference. Clearly, the kind of function representation described for LIP can fulfill this role.

Conveniently, function representation can be understood analogously to scalar and vector representation. Rather than a preferred direction vector in some parameter space, we can take neurons to have preferred functions. This would (approximately) be the function that best matched the neuron's tuning curve over the parameter space (e.g. object position). It is then possible to find the optimal linear functional decoder for estimating some set of functions that the neural population can represent [Eliasmith and Anderson, 2003].

These examples demonstrate the wide variety of kinds of mathematical objects that can be represented in a neurobiologically plausible way. Nevertheless, a question that might remain for cognitive scientists is: is this kind of representation appropriate for understanding *cognitive* behaviour? That is, do we have reason to think that things like functions, vectors, and scalars are sufficient for explaining imagery, object recognition, and language use? After all, the standard kinds of representations used to explain these cognitive phenomena are things like pictures, graphs, and symbolic representations. I would like to suggest that the differences between these kinds of representations are terminological. That is, representations like images *just are* functions, or vectors. Take, for instance, the representation of an image on a computer. A computer represents an image as simply a long binary string. A long binary string is a vector. Similarly, if we consider a non-digitized version of a projected image, we can express the image as changes in light intensities that are a function of spatial position. That is, we can write down a mathematical expression which captures all of the information contained in the projected image. So, despite the particular vocabulary we might use for identifying classes of representations, they contain the same information about objects,

spatial relations, color patterning, and so on. Similar kinds of examples can be generated for the representation of objects and language-like symbols [Eliasmith, 2003].

So, there is nothing inherently limiting in understanding neurons as able to represent mathematical objects rather than the typical classes of representation identified by cognitive scientists. And, the characterization of a wide variety of representations as variations on a theme (i.e. nonlinear encoding and linear decoding), rather than distinct classes, greatly unifies our understanding of representation in neurobiological systems.

In some ways, this discussion of representation has a great affinity with the discussion of distributed representation in connectionism. Clearly, each of these kinds of population representation are distributed representations. However, there are some important differences. For instance, connectionists only consider the case of vector representation. For another, they have no principled relation between the representations of individual nodes and 'higher-level' representations over a set of nodes (i.e., each node is an element of the represented vector for connectionists, but not for the distributed representation described above). For example, if we take a neural system to be representing images, we can define a representation of that image as a function (a 'higher-level' representation) embedded in a population of neurons (at the 'lower-level'). Connectionists cannot do this. And, related to these two limitations, connectionists almost always try to learn representations rather than having a hypothesis about what representations are relevant for a particular task and then determining if that hypothesis is consistent with other (computational, dynamic, or intrinsic) properties of the population.

The constraints introduced by trying to learn representations can be very severe. It is well known that the order in which items are presented, the number of occurrences of items in the presentations, and the particular choice of which items are presented all greatly influence the results of learning. Furthermore, while these difficulties affect the typical three layer network, they become unmanageable for more complex kinds of networks. In contrast, understanding neural representation in terms of encoding and decoding avoids such difficulties. Very complex networks, with any number of layers, can have representations and transformations embedded in them using these methods. As well, since the implementation of the representations is a neural one, biologically plausible learning can be introduced at any point in the process, as necessary. So, while both connectionism and computational neuroscience approaches understand representations as distributed, only the latter provides principled methods for moving between levels of description of the representational capacities in a neurobiologically realistic network.

To this point, I have described both population representation and temporal representation independently. However because both descriptions are cases of nonlinear encoding and linear decoding, it is a simple matter to combine these two kinds of representation. That is, rather than having a separate temporal decoder and a separate population decoder, we can define a single population-temporal decoder which can be used to decode a spiking, population-wide encoding

of some mathematical object that captures the properties to be represented. This, then, completes a computational neuroscientific description of the representational vehicles employed by neurobiological systems.

2.3 Semantics

As I mentioned earlier, the tools of computational neuroscience are best suited to characterizing representational vehicles rather than representational content. Nevertheless, the characterization of vehicles and contents is not independent (contra [Cummins, 1989]): if we think that all vehicles can only support scalars, then only the kinds of contents that can be carried by scalars can be contents of a system with those vehicles. In general, the ways in which we can characterize vehicles ought to tell us something about the kinds of contents those vehicles can carry.

There are three broad classes of semantic theories: causal, conceptual role, and two-factor theories. Causal theories of meaning have as their main thesis that mental representations are about, and thereby mean, what causes them [Dretske, 1981; 1995; Fodor, 1990; 1998]. In the context of the previous discussion this means that the encoding process alone determines meaning. Conceptual role theories hold that the meaning of a term is determined by its overall role in a conceptual scheme [Loar, 1981; Harman, 1982]. Under such theories, the meaning of a term is determined by the inferences it causes, the inferences it is the result of, or both. Here, the focus is on the decoding of whatever information happens to be in some neural state. One theoretical move, to avoid the difficult problems that arise when adopting either a causal theory or a conceptual role theory, is to combine them into a 'two-factor' theory [Field, 1977; Block, 1986]. On two-factor theories, causal relations and conceptual role are equally important, independent elements of the meaning of a term: "the two-factor approach can be regarded as making a conjunctive claim for each sentence" [Block, 1986, p. 627]. So, only two-factor theories explicitly acknowledge both encoding and decoding.

In the preceding characterization of representational vehicles, a representation is only defined once both the encoding and decoding processes are identified. This means that, contrary to both causal and conceptual role theories of content, both how the information in neural spikes is used in (decoding), as well as how it is related to previous goings-on (encoding) are relevant for determining content. So given the characterization of vehicles I have presented, two-factor theories of content seem most plausible.

However, it is assumed by past two-factor theories that the factors are independent. This property raises a grave difficulty for such theories. In criticizing Block's theory, Fodor and Lepore [1992] remark "We now have to face the nasty question: *What keeps the two factors stuck together?* For example, what prevents there being an expression that has the inferential role appropriate to the content *4 is a prime number* but the truth conditions appropriate to the content *water is wet*?" [p. 170]. If, in other words, there is no relation between the two factors (i.e., they are simply a conjunction), it is quite possible that massive misalignments between

causal relations and conceptual roles can occur.

However, in the computational neuroscientific characterization of representation presented above, there is a tight relation between the encoding and decoding processes. Broadly speaking, the population-temporal decoders are found in order to estimate some function of the encoded parameter. While for simple representation this function is identity, it need not be (as I discuss below). That is to say, all of the inferences derivable from some particular neural encoding depend on the information carried by that encoding. As a result, if there is no relation between the 'wetness of water' and '4 being a prime number', it would it be impossible for the latter to be part of the conceptual role of the encoding of the former. While there remains much to be said regarding the precise relation between encoding and decoding, such considerations suggest that this characterization can avoid the main weakness of past two-factor theories.

3 COMPUTATION

Conveniently, the above characterization of representation leads naturally to an understanding of computation in neural systems. Of course, without an account of computation, any characterization of representation is useless. The vast variety of complex and interesting behaviours observable in animals arise not from simply representing the environment, but performing complicated, and currently ill-understood, computations with these representations.

Before considering neural computation in more detail, it will be beneficial to address a terminological problem that often arises during such discussions. In the past, the notion of 'computation' has been closely allied to processing by serial digital computers. For this reason, many theorists take computational characterizations to be applicable only to machines that have discrete symbolic representations and process them via rules [van Gelder, 1995]. Unfortunately, this use of the term belies the existence of an entire field of research on analog computation [Uhr, 1994; Douglas and Mahowald, 1995; Hammerstrom, 1995]. The reason the strict notion of computation is often adopted, is because it is unclear whether there can be a principled distinction between computational and non-computational mechanisms if analog computation is considered to be a kind of computation. Nevertheless, I will adopt the widely accepted view in computational neuroscience at the notion 'computation' is applicable to neural systems in virtue of the kinds of descriptions we apply to them [Churchland *et al.*, 1990]: indeed, this area of research would not be called *computational* neuroscience otherwise.

This terminological issue aside, then, we can define neural computation — just as we did neural representation — in terms of a nonlinear encoding and a (different) linear decoding. Essentially, representation consists of trying to compute the identity function. That is, whatever is encoded into the neural spikes trains is what we are trying to decode when we take neurons to be representing their input. I refer to the decoders used for computing this function as 'representational decoders'. More generally, we can identify decoders for computing any function of

the encoded input: I refer to these other decoders as 'transformational decoders'. So, for example, if we define the representation in LIP to be a representation of the position of an object, we can find representational decoders that estimate the actual position given the neural firing rates. However, we can also use exactly the same encoded information to estimate where the object would be if it was translated 5 ° to the right. For this we could identify a transformational decoder. This particular example is merely linear transformation of the encoded information, and so not especially interesting. However, exactly the same methods can be used to find the transformational decoders for estimating nonlinear computations as well (e.g., perhaps the system needs to compute the square of the position of the object; [Eliasmith and Anderson, 2003]).

This account of computation is successful largely because of the nonlinearities in the neural encoding of the available information. When decoding, we can either attempt to eliminate these nonlinearities by appropriately weighting the responses of the population (as with representational decoding), or we can emphasize the nonlinearities necessary to compute the function we need (as with transformational decoding). In either case we can get a good estimate of the appropriate function, and we can improve that estimate by including more neurons in the population encoding the information.

4 DYNAMICS

Given the previous characterizations of representation and computation, it is possible to build neurally realistic circuits that take time time-varying signals as input and compute 'interesting' functions of those signals. However, these techniques, as they stand, apply only to feedforward computations. As is well-known, recurrence, or backward projections, are ubiquitous in neural systems. This kind of complex interconnectivity suggests that feedforward computation is not sufficient to understand neurobiological function. As a result, computational neuroscientists need a means of characterizing the sophisticated, possibly recurrent, internal dynamics for the representations they take to be present in neural populations.

Eliasmith and Anderson [2003] suggest that neural dynamics can be best understood by taking neural representations to be control theoretic state variables. Control theory is a set of mathematical techniques developed in the 1960s to analyze and synthesize complex, analog, physical systems [Kalman, 1960]. For linear, time-invariant (LTI) systems, control theory provides a canonical way of expressing, optimizing, and analyzing the set of possible behaviours of the system. More complex dynamics, such as nonlinear and time-varying dynamics, can also be expressed using control theory, although analysis of the systems is no longer guaranteed to be tractable.

The standard state-space form for control theoretic descriptions of physical systems is a set of differential equations defined over variables called the "state variables" (figure 3a). For any system so described, the current value of the state variables and the set of differential equations governing their dynamics completely

determines the future behaviour of the system. In neural systems, the set of differential equations can be taken to describe how the representation in a neural population changes over time. The value of the variables at any particular time is determined by the (spiking) neural representation at that time, and the governing equations are determined by the connection weights between that population and any others providing input to it (possibly including that population itself).

Notably, the standard control system depicted in figure 3a assumes that the dynamics of the physical system being described can be characterized as integration (hence the transfer function being an integral). However, neurons have their dynamics determined by intrinsic properties (e.g., ion channel speed, membrane capacitance, etc.), and do not naturally support integration. As a result, it is necessary to be able to translate the standard control theoretic equations into a form appropriate for neural systems (figure 3b). Fortunately, this translation can be done in the general case [Eliasmith and Anderson, 2003]. Such a translation allows any standard control theoretic description of a system to be written in an equivalent 'neural' control theoretic form. The ability to affect such a translation can prove a great benefit to theorists, as they can draw from the vast resources of control theory when hypothesizing about which neural architectures may be able to realize some function: a function that may already be well-understood by control theorists.

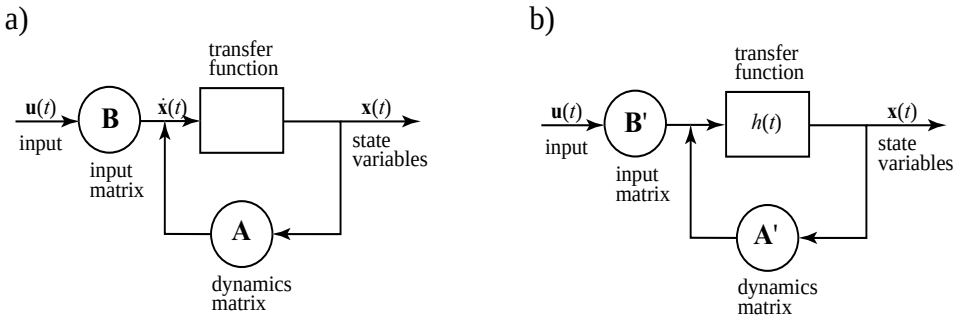

Figure 3. Diagram of the dynamics equation for LTI control theoretic descriptions of a) a standard physical system, $\dot{\mathbf{x}}(t) = \mathbf{A}\mathbf{x}(t) + \mathbf{B}\mathbf{u}(t)$ and b) a neural system, $\mathbf{x}(t) = h(t) * (\mathbf{A}'\mathbf{x}(t) + \mathbf{B}'\mathbf{u}(t))$. The input signal, $\mathbf{u}(t)$, can be modified by the parameters in the input matrix, $\mathbf{B}$, before being added to any recurrent signal which is modified by the parameters in the dynamics matrix, $\mathbf{A}$. The result is then passed through the transfer function which defines the dynamics of the state variable, $\mathbf{x}(t)$. In a), the canonical form, the transfer function is integration. In b), the neural form, the transfer function is determined by intrinsic neural dynamics. Fortunately, given the canonical form and the transfer function, $h(t)$, $\mathbf{A}$' and $\mathbf{B}$' can be determined for any $\mathbf{A}$ and $\mathbf{B}$.

Notably, adopting this description of the dynamics of neural systems is reminiscent of the 'dynamicist' view in cognitive science [Port and van Gelder, 1995]. Indeed, both characterizations share an emphasis on the importance of time, and both suppose that sets of differential equations are the best mathematical tools available for quantifying cognitive systems under this assumption. However, there is an extremely important difference between these two views. For dynamicists, the variables over which the differential equations are defined are not explicitly related to the physical system itself [Eliasmith, 1997; 2003]. So, for example, the "motivation" variable in motivational oscillatory theory (MOT), which is intended to characterize some high-level property of the animal, is never related to any specific physical property of the system [Busemeyer and Townsend, 1993] — and it is entirely unclear how it could be. In contrast, the approach described above is explicit on the relation between higher-level neural representations and the activations of single cells. And, it is precisely these representations that serve as the variables over which the dynamics are defined. This makes the computational neuroscience approach more directly responsive to (and informative for) the experimental results generated by neuroscientists and cognitive neuropsychologists (see [Hardcastle, this volume]).

5 SYNTHESIS

The previous sections have defined three important principles for characterizing neurobiological systems (principles of representation, transformation (or computation), and dynamics). However, it may not yet be clear how these principles interact, and, more importantly, how they are intended to map onto the observable properties of real neural systems.

Figure 4 depicts how these principles can be integrated in order to describe the functioning of neurobiological systems at various levels of description. Specifically, figure 4 shows the components of a generic neural subsystem, including temporal decoders, population decoders, control matrices, encoders, and the spiking neural nonlinearity. A series of such subsystems can be connected in order to describe larger neural systems, since both the inputs and outputs of the subsystems are neural spikes.

Additionally, figure 4 depicts what it means to suggest that neural representations are control theoretic state variables. The state variables are defined by the temporal and population decoders and encoders, the dynamics of the control system are defined by the control matrices, and any functions that must be computed in order to implement the control system can be estimated by replacing the appropriate representational decoders with transformational decoders.

This figure also captures how the theoretical elements of this description map onto real neural systems. In particular, the control matrices, decoders, and encoders can be used to analytically compute the connection weights necessary to implement the desired control system in the neural population. The temporal decoders, as noted earlier, are mapped onto the postsynaptic currents (PSCs) pro-

duced in dendrites as a result of incoming neural spikes. Finally, the weighted dendritic currents arriving at the soma (cell body) of the neuron determine the output of the neural nonlinearity, i.e., the timing of neural spikes produced by neurons in this population.[5]

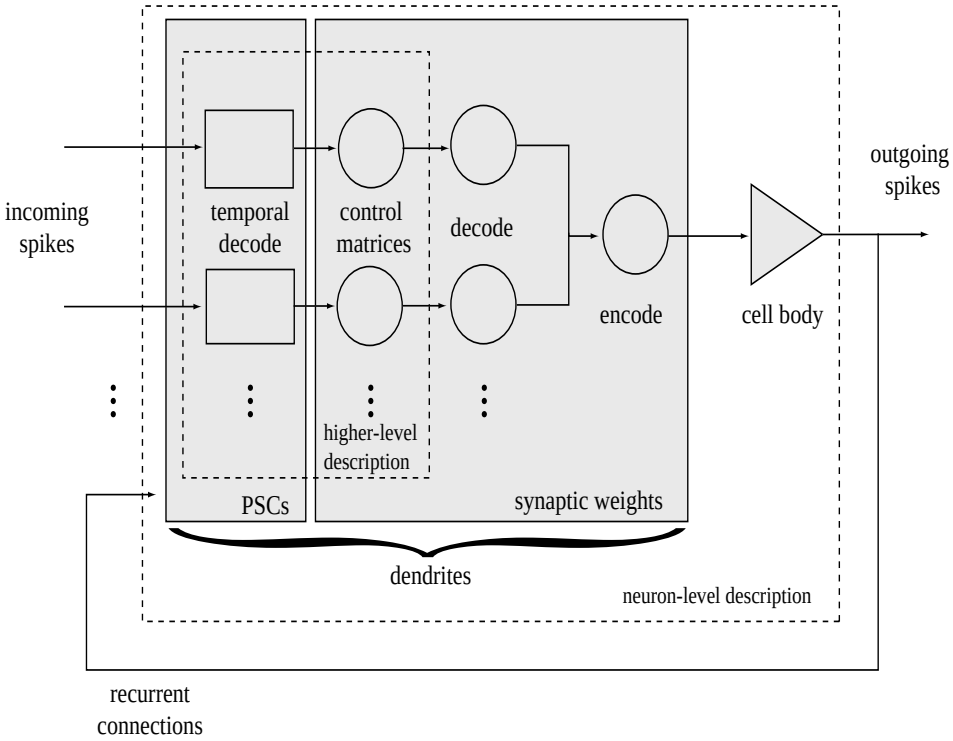

Figure 4. A generic neural subsystem (adapted from [Eliasmith, 2003]). The outer dotted line encompasses elements of the neuron-level description, including PSCs, synaptic weights, and the neural nonlinearity in the soma. The inner dotted line encompasses elements of the control theoretic descriptions at the higher-level. The grey boxes identify experimentally measurable elements of neural systems. The elements inside those boxes denote the theoretically relevant components of the description. See text for details.

Notably, the theoretical elements in this description are not identical to physically measurable properties of neural systems. As a result, there is a sense in which neural systems themselves never decode the representations they employ.

[5] This nonlinearity can be captured by a set of differential equations that describes the dynamics of the channel conductances that control the flow of ions through the cell membrane resulting in action potentials, or it could be a simpler reduced model of neural spiking (like the common leaky integrate-and-fire model).

This is because decoding, encoding, and the dynamics determined by the control matrices are all included in the synaptic weights. Nevertheless, if our assumptions regarding representation or dynamics of the system are incorrect, the model which embodies these assumptions will make incorrect predictions regarding the responses of individual neurons. So, while we cannot directly measure decoders, we can justify their inclusion in a description of neural systems insofar as the description is a successful one at predicting the properties we can measure. This, of course, is a typical means of justifying the introduction of theoretical entities in science.

This approach has been successfully applied to modeling the number of neurobiological systems, including sensory, cognitive, and motor systems [Eliasmith and Anderson, 1999; 2000; 2001; Nenadic *et al.*, 2000]. Of greatest interest to cognitive scientists, however, is whether or not this kind of description can contribute to improving our understanding of high-level cognitive function.

6 COGNITION

On occasion, cognitive scientists have expressed hostility towards the idea that our understanding of cognitive processes can be improved by a better understanding of how the brain functions [Lycan, 1984; Fodor, 1999]. They have suggested that knowing where things happen in the brain, or how they are implemented in the brain, is not relevant for answering the more important question of what the cognitive architecture is. This, however, is no longer the widely-held view that it once was. The immense recent increase in cognitive research employing techniques that measure physiological signals in the brain (such as functional magnetic resonance imaging (fMRI), positron emission tomography (PET), event related potentials (ERPs), etc.), suggest that a more common view is that improving our understanding of neural organization will greatly benefit our understanding of the nature of cognitive function.

Despite the seeming prominence of this view, there are no well-established modeling techniques for bridging the gap between the functioning of single neurons and the functioning of the cognitive systems of which they are a part. Models that incorporate biologically realistic single neurons tend to focus on low-level perception (e.g., receptive fields, motion, contour sensitivity, etc.), motor control (saccade generation, vestibular ocular reflex, invertebrate locomotion, digestion, etc.), and single-cell learning (e.g., retinal wave effects, receptive field learning, cortical column organization, etc.). In contrast, cognitive scientists are largely interested in models of cognitive functions like reasoning, memory, language, and so on. So, the question naturally arises as to whether the methods of computational neuroscientists can actually prove useful to cognitive scientists.

For some time now, computational neuroscientists have built sophisticated models of brain areas whose functions map naturally onto certain psychological categories. Perhaps the most successful example of such models are those related to short-term/ working memory [Camperi and Wang, 1998; Laing and Chow, 2001;

Brody *et al.*, 2003]. In general, these models have helped make the notion of working memory much more precise. Specific implementational constraints have been derived from some of these models, including the kinds of time constants and receptors that are likely necessary to support working memory function. Preliminary results regarding the source of working memory limitations [Trappenberg, 2003], and the cost of representational complexity [Eliasmith and Anderson, 2003] in such networks have also been generated. Such results have vastly improved our understanding of how working memory might be organized, its strengths, expected limitations, and so on. Such insights should inform our functional decomposition of other high-level cognitive categories. So, in this instance, computational neuroscience serves as a useful means of filling in important details for improving our understanding of a specific cognitive function, and may serve as a springboard for better understanding other cognitive functions.

However, working memory itself does not seem to be particularly representative of what counts as truly cognitive behaviour, since it is found in a wide variety of nonhuman animals. As well, the available models of working memory tend not to be systemic. That is, they are focused on small, highly-constrained neural areas, and do not have complex dynamics or support sophisticated representations. So, it could well be argued that such examples do not demonstrate that the methods described earlier will be able to tackle interesting problems in cognitive science.

Recently, however, significantly more sophisticated models of clearly high-level cognitive function have been built [Eliasmith, 2005]. As an example, I consider a model of the Wason card selection task [Wason, 1966]. This task requires the manipulation of language-like structures and is performed only by humans. Thus it is, if anything is, a high-level cognitive phenomenon. A considerable amount of research in psychology has focussed on exploring this task [Griggs and Cox, 1982; Reich and Ruth, 1982; Cheng and Holyoak, 1985; Cosmides, 1989; Gigerenzer and Hug, 1992; Chater and Oaksford, 1996; Liberman and Klar, 1996; Fiddick *et al.*, 2000]. The task itself consists of subjects being presented with a conditional statement such as "if P then Q," which they must test for validity given some data. So, for example, suppose that a subject is presented with a rule such as "if a card has a vowel on one side, then it has an even number on the other". They are then give a data set that consists of four cards laying on a table with one of 'A', 'B', '2', or '3' written on the visible side of each card. The subject is then asked to choose the card or cards that must be turned over in order to test this conditional. The correct choice in this example are the 'A' and '3' cards. However, this combination is usually chosen only by a small minority of subjects (5-40%). More often, subjects will choose the 'A' card or the 'A' and '2' cards (60-90%).

In an interesting variation on this task, the rules chosen were made significantly more concrete, or familiar. For instance, the researcher might use the rule, "if someone drinks alcohol then that person is over 21," with the data being analogous to that presented earlier. This shift in the rule's content massively increased success on the card selection task. Much of the discussion regarding the Wason card selection task focuses on attempting to explain the differences between this

'concrete' task and the original 'abstract' task.

The neural-level model generated using the techniques described earlier demonstrates how a specific hypothesis regarding these differences can be shown to be neurally plausible. Specifically, the model shows that domain general, context-sensitive inference [Cheng and Holyoak, 1985; Cosmides, 1989; Gigerenzer and Hug, 1992] can account for the differences in performance in human subjects on these tasks in a manner consistent with known neuroanatomy and neurophysiology. Some highlights of the model are that: 1) it is a large-scale model, consisting of nine separate, interconnected populations; 2) it performs structure-sensitive processing (it uses one of the available vector binding operations [Plate 1994; 1997], for encoding the rules into language-like structures and transforming the rules appropriately); 3) it performs context-sensitive processing (i.e., different transformations are applied depending on whether the rule is concrete or abstract); 4) it learns which transformations to perform (i.e., based on the provision of the correct answer, it updates the appropriate transformation in the current context to determine which transformation is appropriate for which context); 5) the model is mapped to known neuroanatomy; 6) these functions are carried out using populations of spiking neurons that encode complex representations (100 dimensional vectors); and 7) it successfully performs syntactic generalization (often considered a hallmark of systematicity, and hence cognition).

In sum, this model has both cognitive and neural plausibility: that is, it explains cognitive behaviour and does so in a neurally realistic way. Unlike past attempts at understanding human performance on this task, it can be used to make specific predictions regarding physiological measures of people performing these tasks, as well as behavioural predictions. Also unlike other suggestions, it has been shown that the hypothesized dynamics, representations, and computations can actually be carried out in neural hardware. More generally, this model serves to demonstrate that we can, indeed, begin to bridge the gap between neural and psychological explanations of cognitive systems.

Given such examples, I suspect that the methods of computational neuroscience will not only prove useful, but will be essential for giving a comprehensive explanation of how cognitive systems operate. It is not only the success of such models, but also our changing understanding of the nature of cognitive systems, that will render such accounts indispensable. Let us consider two recent shifts in our understanding of cognition that support this contention.

First, it has been suggested by a number of psychologists and philosophers, perhaps Gibson [1955] being the most famous, that the distinction between perception/action and cognition may be a misleading one. They suggest that in order to understand the complicated behaviours we see arising from neurobiological systems — no matter how complicated — we must first understand the tight connection between perception and action. Indeed, there has been a significant amount of recent work in dealing with what is often called "embedded" or "embodied" cognition [Varela *et al.*, 1991; Haugeland, 1993; Hutchins, 1995; Clark, 1997; Kelly, 2001]. This is the view that in order to understand cognition, we must un-

derstand an organism as placed within its environment. These theorists thus take the organism, its actions, and its environment as tightly linked: linked through the perceptual/motor loop of the organism itself. For these theorists, thinking of cognitive behaviours as somehow specially internal, or separable from action and perception, is misguided. Related experimental work suggests that much of our cognitive life is thus tightly tied to more 'basic' perceptual and action-oriented processes [Barsalou 1999; Richardson *et al.*, 2001; Matlock 2004]. Thus, if we cannot easily distinguish cognitive processes from non-cognitive ones, then work in computational neuroscience on simple motor and perceptual systems is highly relevant for understanding cognitive function.

Second, not only has the distinction between cognition and perception/action come under scrutiny in recent years, but so has the distinction between cognition and emotion [Damasio, 1994]. There is ample neurological evidence that people with damage to emotionally related brain structures (e.g. ventromedial prefrontal cortex) become partly cognitively incompetent. While many cognitive functions are preserved in such patients, tasks such as decision-making, risk assessment and planning become severely impaired [Bechara *et al.*, 2000; Spinella 2003]. As a result, the traditional cognitivist view that emotions inhibit our ability to be rational seems implausible. From a computational perspective, this suggests that distinguishing cognitive systems from other systems may not allow us to properly explain what we typically take to be cognitive functions.

Both of these challenges to the idea that cognition is somehow separable from other neural functions suggest the need for a method that is continuous between typically cognitive systems and supposedly non-cognitive ones. On the face of it, this need should not be surprising since all such systems fundamentally use spiking neural cells to perform their functions. As a result, if we have methods for embedding high-level functions into complex networks of neural cells, we can use those methods to consistently characterize what are intuitively different high-level functions (e.g., cognitive versus non-cognitive). This consistency thus makes it possible to interface models of traditionally non-cognitive systems with those of more cognitive brain systems. As a result, the methods of computational neuroscience, unlike other methods, make it possible to unify a wide variety of neural functions — a unification that seems essential if we are to understand the whole system and its continuously impressive variety of abilities.

BIBLIOGRAPHY

[Andersen *et al*, 1985] R. A. Andersen, G. K. Essick, and R. M. Siegel. The encoding of spatial location by posterior parietal neurons. *Science* **230**: 456-458, 1985.

[Barsalou, 1999] L. Barsalou. Perceptual symbol systems. *Behavioral and Brain Sciences* **22**: 577-609, 1999.

[Bechara *et al.*, 2000] A. Bechara, H. Damasio, and A. Damasio. Emotion, decision making and the orbitofrontal cortex. *Cerebral Cortex* **10**(3): 295-307, 2000.

[Bialek *et al.*, 1991] W. Bialek, F. Rieke, R. de Ruyeter van Steveninck and D. Warland. Reading a neural code. *Science* **252**: 1854-57, 1991.

[Block, 1986] N. Block. Advertisement for a semantics for psychology. In *Midwest Studies in Philosophy*. P. French, T. Uehling and H. Wettstein (eds). Minneapolis, University of Minnesota Press. **X:** 615-678, 1986.

[Britten and Newsome, 1998] K. Britten and W. Newsome. Tuning bandwidths for near-threshold stimuli in area MT. *Journal of Neurophysiology* **80**(2): 762-770, 1998.

[Brody *et al.*, 2003] C. D. Brody, R. Romo, and A. Kepecs. Basic mechanisms for graded persistent activity: discrete attractors, continuous attractors, and dynamic representations. *Current Opinion in Neurobiology* **13**: 204-211, 2003.

[Buracas *et al.*, 1998] G. T. Buracas, A. M. Zador, M. DeWeese and T. Albright. Efficient discrimination of temporal patterns by motion-sensitive neurons in primate visual cortex. *Neuron* **20**: 959-969, 1998.

[Busemeyer and Townsend, 1993] J. R. Busemeyer and J. T. Townsend. Decision field theory: A dynamiccognitive approach to decision making in an uncertain environment. *Psychological Review* **100**(3): 432-459, 1993.

[Camperi and Wang, 1998] M. Camperi and X. J. Wang. A model of visuospatial working memory in prefrontal cortex: recurrent network and cellular bistability. *Journal of Computational Neuroscience* **5**: 383-405, 1998.

[Chater and Oaksford, 1996] N. Chater and M. Oaksford. Rational explanation of the selection task. *Psychological Review* **103**: 381-391, 1996.

[Cheng and Holyoak, 1985] P. W. Cheng and K. J. Holyoak. Pragmatic reasoning schemas. *Cognitive Psychology*. **17**, 391-416, 1985.

[Churchland, 1989] P. Churchland. *A neurocomputational perspective*. Cambridge, MA, MIT Press, 1989.

[Churchland *et al.*, 1990] P. Churchland, C. Koch, and T. Sejnowski. What is computational neuroscience? In *Computational Neuroscience*. E. Schwartz (ed). Cambridge, M.A., MIT Press, 1990.

[Clark, 1997] A. Clark. *Being there: Putting brain, body and world together again*. Cambridge, MA, MIT Press, 1997.

[Cosmides, 1989] L. Cosmides. The logic of social exchange: has natural selection shaped how humans reason? Studies with the Wason selection task. *Cognition* **31**: 187-276, 1989.

[Cummins, 1989] R. Cummins. *Meaning and mental representation*. Cambridge, MA, MIT Press, 1989.

[Damasio, 1994] A. R. Damasio. *Descartes' error: Emotion, reason, and the human brain*. New York, NY, Grosset/Putnam, 1994.

[Stevenick *et al.*, 1997] R. R. de ruyter van Steveninck, G. D. Lewen, S. Strong, R. Koberle, and W. Bialek. Reproducibility and variability in neural spike trains. *Science* **275**: 1805-1808, 1997.

[Deneve and Pouget, 2004] S. Deneve and A. Pouget. Bayesian multisensory and cross-modal spatial links. *Journal of Neurophysiology*. **98**:249-58, 2004.

[Desimone, 1991] R. Desimone. Face-selective cells in the temporal cortex of monkeys. Journal of Cognitive Neuroscience **3**: 1-8, 1991.

[Douglas and Mahowald, 1995] R. Douglas and M. Mahowald. Silicon neurons. In *The handbook for brain theory and neural networks*. M. Arbib (ed). Cambridge, MA, MIT Press, 1995.

[Dretske, 1981] F. Dretske. *Knowledge and the flow of information*. Cambridge, MA, MIT Press, 1981.

[Dretske, 1988] F. Dretske. *Explaining behavior*. Cambridge, MA, MIT Press, 1988.

[Dretske, 1995] F. Dretske. *Naturalizing the Mind*. Cambridge, MA, MIT Press, 1995.

[Eliasmith, 1997] C. Eliasmith. Computation and dynamical models of mind. *Minds and Machines* **7**: 531-541, 1997.

[Eliasmith, 2003a] C. Eliasmith. Moving beyond metaphors: Understanding the mind for what it is. *Journal of Philosophy* **100**(10): 493-520, 2003.

[Eliasmith, 2003b] C. Eliasmith. Neural engineering: Unraveling the complexities of neural systems. *IEEE Canadian Review* **43**: 13-15, 2003.

[Eliasmith, 2005] C. Eliasmith. Cognition with neurons: A large-scale, biologically realistic model of the Wason task In *Proceedings of the 27^{th} Annual Conference of the Cognitive Science Society*, B. Bara, L. Barsalou, & M. Bucciarelli (eds). Lawrence Erlbaum, 624-629, 2005.

[Eliasmith and Anderson, 1999] C. Eliasmith and C. H. Anderson. Developing and applying a toolkit from a general neurocomputational framework. *Neurocomputing* **26**: 1013-1018, 1999.

[Eliasmith and Anderson, 2000] C. Eliasmith and C. H. Anderson. Rethinking central pattern generators: A general framework. *Neurocomputing* **32**: 735-740, 2000.

[Eliasmith and Anderson, 2001] C. Eliasmith and C. H. Anderson. Beyond bumps: Spiking networks that store smooth n-dimensional functions. *Neurocomputting* **38**: 581-586, 2001.

[Eliasmith and Anderson, 2003] C. Eliasmith and C. H. Anderson. *Neural engineering: Computation, representation and dynamics in neurobiological systems*. Cambridge, MA, MIT Press, 2003.

[Elman, 1998] J. Elman. Connectionism, artificial life, and dynamical systems. In *A Companion to Cognitive Science*. W. Bechtel and G. Graham (eds). Oxford, UK, Blackwell. 488-505, 1998.

[Felleman and van Essen, 1991] D. J. Felleman and D. C. Van Essen. Distributed hierarchical processing in primate visual cortex. *Cerebral Cortex* **1**: 1-47, 1991.

[Fiddick *et al.*, 2000] L. Fiddick, L. Cosmides, *et al.*. No interpretation without representation: the role of domain-specific representations and inferences in the Wason selection task. *Cognition* **77**: 1-79, 2000.

[Field, 1977] H. Field. Logic, meaning, and conceptual role. *Journal of Philosophy* **74**: 379- 409, 1977.

[Fodor, 1981] J. Fodor. *Representations*. Cambridge, MA, MIT Press, 1981.

[Fodor, 1990] J. Fodor. *A theory of content and other essays*. Cambridge, MA, MIT Press, 1990.

[Fodor, 1998] J. Fodor. *Concepts: Where cognitive science went wrong*. New York, Oxford University Press, 1998.

[Fodor, 1999] J. Fodor. Diary. *London Review of Books* **21**(19), 1999.

[Fodor and Lepore, 1992] J. Fodor and E. Lepore. *Holism: A shopper's guide*. Oxford, UK, Basil Blackwell, 1992.

[Fodor and Pylyshyn, 1988] J. Fodor and Z. Pylyshyn. Connectionism and cognitive architecture: A critical analysis. *Cognition* **28**: 3-71, 1988.

[Georgopoulos *et al.*, 1989] A. P. Georgopoulos, J. T. Lurito, M. Petrides, A. Schwartz, and J. Massey. Mental rotation of the neuronal population vector. *Science* **243**: 234-236, 1989.

[Georgopoulos *et al.*, 1986] A. P. Georgopoulos, A. B. Schwartz, *et al.*. Neuronal population coding of movement direction. *Science* **243**(1416-19), 1986.

[Gibson, 1955] J. J. Gibson and E. J. Gibson. Perceptual learning: Differentiation or enrichment? *Psychological Review* **62**: 324-341, 1955.

[Gigerenzer and Hug, 1992] G. Gigerenzer and K. Hug. Domain specific reasoning: Social contracts, cheating and prospective change. *Cognition* **43**: 127-171, 1992.

[Griggs and Cox, 1982] R. Griggs and R. Cox. The elusive thematic material effect in Wason's selection task. *British Journal of Psychology* **73**: 407-420, 1982.

[Hammerstrom, 1995] D. Hammerstrom. Digital VLSI for neural networks. In *The handbook of brain theory and neural networks*. M. Arbib (ed). Cambridge, MA, MIT Press, 1995.

[Harman, 1982] G. Harman. Conceptual role semantics. *Notre Dame Journal of Formal Logic* **23**: 242-56, 1982.

[Haugeland, 1993] J. Haugeland. Mind embedded and embodied. In *Proceedings of Mind and Cognition: An International Symposium*, Taipei, Taiwan, Academia Sinica, 1993.

[Hodgkin and Huxley, 1952] A. L. Hodgkin and A. Huxley. A quantitative description of membrane current and its application to conduction and excitation in nerves. *Journal of Physiology* **117**: 500-544, 1952.

[Hoppensteadt and Izhikevich, 2003] F. Hoppensteadt and E. Izhikevich. Canonical neural models. *The handbook of brain theory and neural networks*. M. Arbib (ed). Cambridge, MA, MIT Press, 2003.

[Hubel and Wiesel, 1995] D. Hubel and T. Wiesel. Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. *Journal of Physiology* **160**: 106-154, 1995.

[Hutchins, 1995] E. Hutchins. *Cognition in the wild*. Cambridge, M.A., MIT Press, 1995.

[Kalman, 1960] R. E. Kalman. Contributions to the theory of optimal control. *Boletin de la Sociedada Matematica Mexicana* **5**: 102-119, 1960.

[Kalman, 2004] R. E. Kalman. A new approach to linear filtering and prediction problems. *ASME Journal of Basic Engineering* **82**: 35-45, 1960.

[Kelly, 2001] S. D. Kelly. Demonstrative concepts and experience. *Philosophical Review* **110**(3): 397-420, 2001.

[Kording and Wolpert, 2004] K. P. Kording and D. M. Wolpert. Bayesian integration in sensorimotor learning. *Nature* **427**: 244-247, 2004.

[Kosslyn, 1994] S. Kosslyn. *Image and brain: The resolution of the imagery debate*. Cambridge, MA, The MIT Press, 1994.

[Laing and Chow, 2001] C. R. Laing and C. C. Chow. Stationary bumps in networks of spiking neurons. *Neural Computation* **13**: 1473-1494, 2001.

[Lettvin *et al.*, 1988/1959] J. Lettvin, H. Maturana, W. McCulloch and W. Pitts. What the frog's eye tells the frog's brain. *Embodiments of mind*. W. McCulloch (ed). Cambridge, MA, MIT Press, 1988/1959.

[Liberman and Klar, 1996] N. Liberman and Y. Klar. Hypothesis testing in Wason's selection task: social exchange cheating detection or task understanding. *Cognition* **58**: 127-156, 1996.

[Loar, 1981] B. Loar. *Mind and meaning*. London, UK, Cambridge University Press, 1981.

[Lycan, 1984] W. Lycan. *Logical form in natural language*. Cambridge, MA, MIT Press, 1984.

[MacKay and McCulloch, 1952] D. MacKay and W. S. McCulloch. The limiting information capacity of a neuronal link. *Bulletin of Mathematical Biophysics* **14**: 127-135, 1952.

[Matlock, 2004] T. Matlock. Fictive motion as cognitive simulation. *Memory and Cognition*. **32**(8):1389-400, 2004.

[Moran and Schwartz, 1999] D. W. Moran and A. B. Schwartz. Motor cortical representation of speed and direction during reaching. *Journal of Neurophysiology* **82**: 2676-2692, 1999.

[Nenadic *et al.*, 2000] Z. Nenadic, C. H. Anderson, and B. Ghosh. Control of arm movement using a population of neurons. *Computational neuroscience: Trends in research 2000*. J. Bower (ed), Elsevier Press, 2000.

[Plate, 1997] T. Plate. A common framework for distributed representation schemes for computational structure. *Connectionist systems for knowledge representation and deduction*. F. Maire, R. Hayward and J. Diederich (eds). Brisbane, AU, Queensland University of Technology**:** 15-34, 1997.

[Plate, 1998] T. A. Plate. *Distributed representations and nested compositional structure*. Computer Science. Ph.D. Dissertation. Toronto, University of Toronto, 1994.

[Platt and Glimcher, 1998] M. L. Platt and G. W. Glimcher. Responses of intraparietal neurons to saccadic targets and visual distractors. *Journal of Neurophysiology* **78**: 1574-1589, 1998.

[Platt and Glimcher, 1998] M. L. Platt and G. W. Glimcher. Response fields of intraparietal neurons quantified with multiple saccadic targets. *Experimental Brain Research* **121**: 65- 75, 1998.

[Port and van Gelder, 1995] R. Port and T. van Gelder, eds. *Mind as motion: Explorations in the dynamics of cognition*. Cambridge, MA, MIT Press, 1995.

[Rall, 1957] W. Rall. Membrane time constant of motoneurons. *Science* **126**: 454, 1957.

[Rall, 1962] W. Rall. Theory of physiological properties of dendrites. *Annual New York Academy of Science* **96**: 1071-1092, 1962.

[Rao, 2004] R. Rao. Bayesian computation in recurrent neural circuits. *Neural Computation* **16**, 2004.

[Reich and Ruth, 1982] S. S. Reich and P. Ruth. Wason's selection task: Verification, falsification and matching. *British Journal of Psychology* **73**: 395-405, 1982.

[Richardson *et al.*, 2001] D. Richardson, M. Spivey, S. Edelman, and A. Naples. Language is spatial: Experimental evidence for image schemas of concrete and abstract verbs. *Twenty-third Meeting of the Cognitive Science Society*. Erlbaum, Mawhah, NJ. 873-878, 2001.

[Rieke *et al.*, 1997] F. Rieke, D. Warland, R. de ruyter van Steveninck, and W. Bialek. *Spikes: Exploring the neural code*. Cambridge, MA, MIT Press, 1997.

[Schwartz, 1994] A. B. Schwartz. Direct cortical representation of drawing. *Science* **265**: 540-542, 1994.

[Segundo *et al.*, 1963] J. P. Segundo, G. P. Moore, L. Stensaas, and T. Bullock. Sensitivity of neurones in Aplysia to temporal pattern of arriving impulses. *Journal of Experimental Biology* **40**: 643-667, 1963.

[Sereno and Maunsell, 1998] A. B. Sereno and J. H. R. Maunsell. Shape selectivity in primate lateral intraparietal cortex. *Nature* **395**: 500-503, 1998.

[Shadlen and Newsome, 1994] M. Shadlen and W. Newsome. Noise, neural codes and cortical organization. *Current Opinion in Neurobiology* **4**: 569-579, 1994.

[Shadlen and Newsome, 1995] M. Shadlen and W. Newsome. Is there a signal in the noise? *Current Opinion in Neurobiology* **5**: 248-250, 1995.

[Softky and Koch, 1993] W. Softky and C. Koch. The highly irregular firing of cortical cells is inconsistent with the temporal integration of random EPSPs. *Journal of Neuroscience* **13**: 334-350, 1993.

[Spinella, 2003] M. Spinella. Evolutionary mismatch, neural reward circuits, and pathological gambling. *International Journal of Neuroscience* **113**(4): 503-512, 2003.

[Todorov, 2000] E. Todorov. Direct cortical control of muscle activation in voluntary arm movements: A model. *Nature Neuroscience* **3**: 391-398, 2000.

[Trappenberg, 2003] T. P. Trappenberg. Why is our capacity of working memory so large? *Neural Information Processing-Letters and Reviews* **1**(3): 97-101, 2003.

[Uhr, 1994] L. Uhr. Digital and analog microcircuit and sub-net structures for connectionist networks. In *Artificial intelligence and neural networks: Steps toward principled integration*. V. Honavar and L. Uhr (eds). Boston, MA, Academic Press**:** 341-370, 1994.

[van Gelder, 1995] T. van Gelder. What might cognition be, if not computation? *The Journal of Philosophy* **XCI**(7): 345-381, 1995.

[van Gelder, 1998] T. van Gelder. The dynamical hypothesis in cognitive science. *Behavioral and Brain Sciences* **21**(5): 615-665, 1998.

[Varela *et al.*, 1991] F. Varela, E. Thompson, and E. Rosch. *The embodied mind: Cognitive science and human experience*. Cambridge, MA, MIT Press, 1991.

[Wason, 1966] P. C. Wason. *Reasoning*. New horizons in psychology series. B. M. Foss. Harmondsworth, Penguin, 1966.

[Wilson, 1999] H. Wilson. *Spikes decisions and actions: Dynamical foundations of neuroscience*. Oxford, UK, Oxford University Press, 1999.

PSYCHOPATHOLOGY: MINDING MENTAL ILLNESS

George Graham and G. Lynn Stephens

Psychopathology is the study of the origin, symptoms, consequences and treatment of mental illness and disorder. Individual psychopathologies are the types of mental illnesses themselves, e.g. panic disorder, major depression, mixed personality disorder with histrionic traits, and so forth.

Psychopathology and psychopathologies have long been recognized to be philosophically perplexing and theoretically troubling (cf. [Fulford, Thornton, and Graham, 2006]). This chapter reviews recent philosophical work on psychopathology and psychopathologies. However it goes well beyond a literature review in its defense of a thesis about the most central and widely discussed philosophical question regarding psychopathology. This is the question: What is mental illness?

1 WHAT IS MENTAL ILLNESS?

Or as that question may be put with a linguistic turn: To what do we refer when we speak of mental illness?

It has proven hard to say what mental illness is. It is remarkable that at this late date in the ongoing attempt to describe the nature of mental illness, there is little consensus about whether mental illness is a distinguishable general kind of illness. What makes for this hardness of understanding and absence of consensus? Three factors are critical, among others.

The first is semantic. Although it is one sentence, the question 'What is mental illness?' actually is two questions. The first is:

Q1. When is something mental an *illness*?

The second is:

Q2. When is an illness something *mental*?

Philosophers and philosophically minded mental health professionals often direct their attention, in asking 'what is mental illness?', to the first question and thus to the following issue: If something is mental, how can it be ill? Although Q1 may be a frequently discussed philosophical issue about mental illness, it is not the only source of philosophical perplexity, at least not for physically minded anti-dualists about mind and body, who claim that the mental is not categorically

Handbook of the Philosophy of Science. Philosophy of Psychology and Cognitive Science
Volume editor: Paul Thagard
General editors: Dov M. Gabbay, Paul Thagard and John Woods
© 2007 Elsevier B.V. All rights reserved

non-physical. On a dualist view, a human being or person is a duality, a mind and body paired. Non-mental illness is bodily illness. Mental illness is non-bodily illness. On an anti-dualist view of the physical type, although we may usefully distinguish between mental and non-mental physical characteristics of a person, the mental/physical distinction is just a distinction and not a dichotomy. Each and every characteristic classified as mental is, in some ultimate sense, physical. So, mental illness should not be understood as non-bodily or non-physical. It should be described in terms that are compatible with appreciating a constitutive role for physical states and processes in mental illness. Alas, though, it is not clear how to discharge this descriptive responsibility. It is not obvious how an illness that is spoken of as mental could be something physical, such as a neurobiological disorder or brain disease. There is no question that mental illness is closely associated with physical processes in the brain, but there seem to be truths about mental illness that are not readily reducible to facts about brain structure, biology, and function.

So 'the' question about mental illness actually is two questions. Difficulties associated with each help to explain why consensus about the nature of mental illness has not been achieved. (It may be noted, parenthetically, that in this chapter we plan to focus on the second [Q2] and not the first question [Q1], although we have things to say about the first question. This explains why the chapter bears the sub-title of minding mental illness. The emphasis is on the place of the mental in mental illness and not on the very idea of illness.)

The second factor contributing to difficulty in describing mental illness is the checkered history of medical models of mental illness and particularly the recent failure of the Freudian or psychoanalytic theory of mental illness. In a once popular and unifying picture, famously championed by Freud, mental illness is depicted as unconscious and wayward mental forces, whose lively, fertile, and often repressive dynamics helps to account for everything from episodes of major depression to occasional verbal slips. We assume that Freudianism has proven to be a failure as a general theory of mental illness.[1] There is good reason to believe, for example, in light of anti-Freudian attacks, that too much human mental disturbance or distress qualifies as illness from a Freudian perspective. In the words of Rutgers' sociologist Allan Horwitz [2002]: 'The [Freudian] assumption of a near universality of psychopathology [makes] the abnormal less strange and at the same time [heightens] the strangeness of the normal' [p. 42]. That is, on Freud's conception of mental illness, the contrast between mental illness and non-ill forms of mental distress threatens to evaporate or dissolve. The ill appears un-ill and the non-ill appears ill. Presumably, that's not a welcome feature for a theory of mental illness. Presumably, a good theory of mental illness should preserve a theoret-

[1] The claim that Freudianism is a failed theory of mental illness has been defended in a variety of different ways. See Grunbaum [1984], Kitcher [1992], and Horwitz [2002]. See also Beam [2002] and Hobson and Leonard [2001]. The claim that Freudianism is a failed theory is consistent, of course, with the claim that one or more insights of Freud into mental illness remain valuable or even necessary for the description, explanation or treatment of mental illness.

ical or conceptual contrast, albeit perhaps inescapably imprecise or contestable around its edges, between mental states or conditions that are illnesses and those that, although distressful and disturbing, are not illnesses. (Compare: It would not be wise for a theory of bodily illness if it preserved no essential difference, however fuzzy around the edges, between physical illness and unfortunate bodily conditions that are not illnesses.) However, despite the failure or demise of the Freudian picture of mental illness, Freudianism did provide a general conception of mental illness applicable across a broad range of conditions, whereas as numerous anti-Freudian criticisms have been shed on the causes, consequences, symptoms, and treatment of mental illness, the Freudian picture has not been replaced with an equally broad and unifying alternative conception.

The third factor making it hard to identify the general character of mental illness concerns social institutions and clinical practice. The category of mental illness is firmly established in social and institutional reality. Each year many thousands of people are diagnosed as suffering from some form of mental illness. The legitimacy of such diagnoses is accepted by the vast majority of the public and, frequently, by diagnosed persons themselves. A recognized medical specialty claims the treatment of mental illness as its professional competence and its practitioners receive compensation for their services from private and public payers. Fully accredited medical schools, Ph.D. programs, and other institutions provide training for mental health professionals and support research aimed at understanding and ameliorating mental illness. However despite robust social investment in and clinical commitment to the reality of mental illness, all is not rosily united or smoothly integrated around a coherent general concept of mental illness. The social institutions and clinical practices that supposedly buttress or confirm the reality of mental illness do not speak with one voice or consistently reinforce or confirm the reality of the same sort of condition and this contributes to conceptual instability in the very idea of mental illness. Mental illness's supporters do not agree on how to identify and treat mental illness or isolate its causes. This proliferation of descriptive, therapeutic and etiological conceptions fuels suspicion among some of mental illness's critics that there is little more to the concept of mental illness than a group of poorly organized ideas and theoretical constructs that ought to be abandoned.

Various well-intended critics, many from within the mental health profession, regard the concept of mental illness as a misbegotten notion that obscures our understanding of human health and well-being. Critics offer two sorts of objections. From one direction, critics like Thomas Szasz [1960; 1961] deny that mental illness represents a genuine illness. They charge that the very idea of a mental condition qualifying as illness is part of an attempt to medicalize and clinically or medicinally over-treat the human condition: to classify familiar human distresses, failings, excesses, eccentricities, and other deviations from social expectations as diseases or maladies requiring medical intervention. They also charge that this re-description of 'problems of living' as mental illness lacks scientific support and intellectual coherence. Other critics like Nancy Andreasen [1984; 2001] and J. Allen Hobson

[Hobson and Leonard, 2001] among others (cf. [Taylor, 1999]) acknowledge that mental health professionals often confront real illnesses, but deny that it serves any useful purpose to classify these disorders as distinctively mental. They maintain that 'mental' illness when truly an illness is, more accurately, a disease of the brain, not of the mind. Some such critics perceive no good reason to distinguish mental illness from other behavioral manifestations of brain malfunction such as movement disorders and aphasias.

Although advocates of these two positions often work at cross purposes, it is quite possible, and perhaps reasonable, to combine them. That is, one might argue that some so-called mental illness, perhaps including certain cases of personality disturbance or substance abuse, represent problems of living that are harmful for their subjects but should not be medicalized, whereas others, such as Alzheimer's-disease, are genuine brain diseases and must be treated medically. One might also argue that mental illness as a category is a conceptual mistake or confusion, without offering or even believing that it is desirable or necessary to offer a unified account of conditions currently included within that classification as illnesses.

Very well, then: grant that mental illness is hard to categorize or understand and that the concept of mental illness is difficult to regiment into a coherent concept about which there might be uniform consensus among informed observers, nonetheless how, if regimented, should it be regimented? Suppose we discount wholesale criticisms of the notion of mental illness aimed to motivate discarding the notion, we must still confront, in a wide range of cases, the problem of distinguishing mental illness from problems of living as well as from brain diseases that are not mental illnesses. The concept of mental illness is empirically empty unless it provides principled guidance in discriminating mental illness from instances of imprudent or disruptive or distressful behavior, on the one hand, and non-mental brain disease, on the other.

2 MENTAL QUA MENTAL ILLNESS

So, what is mental illness? If the concept of mental illness is not simply to be abandoned, it must be recast or regimented in ways that permit plausible responses to critics. This must include showing how the concept makes relevant distinctions between mental illness and problems of living. It must also show that, although mental illness, ultimately, in some sense, may be physical it is nevertheless worthwhile to distinguish mental illness proper from non-mental diseases of the brain, such as temporal lobe epilepsy and Parkinson's disease that have psychological consequences but are not mental in nature. Finally, although there is little point in formulating a notion of mental illness that does not have substantial overlap with illnesses of the sort commonly cited as mental and with clinical practice, a suitably regimented notion of mental illness must contain the resources for revisionism at least around the edges of received conceptions and clinical practice.

So, again, to what do or should (for remember, this is regimentation and regimentation carries a prescriptive or legislative component) we refer when we say

'such-and-such a condition is a mental illness'?

To begin with, we assume that we know at least in serviceable terms what makes a condition qualify as an illness (of which more below). So, we read our question of what is mental illness as Q2. When is such-and-such an illness a mental illness?

Here is our answer. It will be explained and defended in the rest of the chapter.

> An illness is *mental*, prototypically speaking, when the conscious representational content of the illness, roughly, how both self and world consciously appear like to its subject or victim, matters critically to four aspects or dimensions of the illness. These include its: (1) origin or causes, (2) symptoms, (3) consequences (including development and behavioral variation), and (4) proper and effective treatment.[2]

3 CLARIFICATION OF PROPOSAL

There are three sets of detailed comments to be made in clarification of the above proposal (in addition to those made in the second footnote).

First: Above we refer to the concept of mental illness 'prototypically speaking'. Here is what we mean by this locution. We assume (without argument here) that the concept of mental illness does not conform to what is called a *classical definition*. That is, it does not encode a tractably specifiable set of necessary and sufficient conditions for its application. Instead, we presuppose that the concept is best understood as possessing what is called a *prototype* semantic structure. According to this picture of the semantics of the concept, the concept encodes features or conditions that exemplary or clear-cut cases of mental illness possess. Conditions (1) through (4) above constitute, we propose, the prototype of mental illness. So, when each and every such feature is present in an illness, this makes for a paradigmatic case of mental illness. So, too, given our assumption that the concept of mental illness has a prototype structure, we speak of illness as mental 'prototypically speaking'.[3]

Second: To get a quick gist of the meaning or purport of above proposed *mental illness prototype* ('MIP' for short), let us pick a candidate mental illness entry

[2]Mattering 'to origin or causes' means that conscious representational content helps casually to explain the occurrence of the illness. 'To symptoms' means that symptoms include conscious content. 'Matters to consequences' means that conscious contents help to produce or control the variations in and behavior associated with the illness. 'To treatment' means that the illness can be treated or therapeutically ameliorated by direct address of its conscious content. For the immediate present we use the expression 'conscious representational content' in an unexamined or intuitive way. The expression is clarified later in the chapter.

[3]In assuming that mental illness fails to admit of a classical definition, we are agreeing with Andreasen [1984] that the concept possesses 'debateable boundaries' [p. 35]. However, in assuming that it possesses a prototype structure, we are disagreeing with Gorenstein [1992, p. 14] that it merely is a 'catchall' concept. On our view, efforts at conceptual regimentation can be useful and successful. Conceptual boundaries, at prototypical peaks, can be defended against unwelcome intruders and suspect conceptual innovations (see [Graham and Horgan, 1998; Ramsay, 1998; Laurence and Margolis, 1999]).

from the Diagnostic and Statistical Manual of Mental Disorders, 4th Edition [APA, 1994]. One such entry is panic disorder (being prone to panic attacks) with agoraphobia [pp. 396-99].

Instances of panic disorder with agoraphobia (literally 'fear of the market place') contain robust and vivid conscious representational content. The content includes, during panic attacks or active episodes of the disorder, fear and anxiety directed at or responsive to situations ('panicky situations') where escape appears to be difficult or impossible to a victim. Typically, too, it consciously seems from the victim's point of view as if she is ill, going crazy, or perhaps even dying. According to the prototype concept of mental illness that we propose, viz. MIP, if such contents of panic disorder help to explain its origin, help to constitute its symptoms, play a causal role in the behavior that results from a panic attack, and can be foci of effective therapeutic treatment and intervention (a topic to be examined much later in the chapter), then panic disorder counts as a case of mental illness. Indeed, it qualifies as a prototypical case of mental illness.

Several disorder entries in DSM-IV conform to MIP. In our judgment these include among others and in addition to panic disorder: Stress-related adjustment disorders, conversion disorders, situation-sensitive mood disorders, some cases of impulse control disorder, and personality disorders that reflect personality changes after traumas or 'stressors that are ... so serious, so overwhelming in their impact, that we must give them (special) status' [Wheaton, 1999, p. 188]. Some of these claims or classifications will be explained later in the chapter.

Although disorder entries in conformity with MIP are not hard to find, some entries in DSM-IV depart in various aspects from MIP. In some cases, such as various movement disorders (see below), the entries conform to what would be a contrary prototype of non-mental physical illness. So, if MIP is adopted, they should be de-classified as mental illnesses or disorders and categorized as non-mental physical illnesses. Still other DSM-IV entries share some but not all features of MIP. So, it may be worth classifying them as mental and non-mental but in different respects or along different dimensions. Alzheimer's type dementia [APA, 1994, pp. 139-143], for example, is both mental and non-mental. Conscious representational content plays a role in Alzheimer's symptoms. Victims may suffer, for example, from thoughts of theft and paranoia, which, in the context of the amnesia and confusion associated with the condition, may be interpretations of mislaying or misplacing objects. In advanced Alzheimer's patients, however, therapies that directly address conscious representational content associated with the disorder are ineffective given the cognitive disabilities or impairments of victims [Lovestone, 2000, pp. 393-394]. Additionally, on the question of mattering critically to consequences of the disorder, reference to content may help to account for the depressed mood sometimes associated with Alzheimer's, perhaps reflecting changes in a victim's perception of disturbed social relationships and diminished personal autonomy. However, in origin or etiology Alzheimer's appears due to plagues and neurofibrillary tangles that form in the hippocampus and spread throughout the cerebral cortex. These are areas of the brain associated with the onset of both de-

mentia and amnesia (cf. [Andreasen, 2001, pp. 263-265; Lovestone, 2000, p. 392]). Conscious representational content plays no causal role in Alzheimer's origin. So, Alzheimer's is removed in certain significant ways from MIP, although not so far removed as to count as non-mental along each and every dimension.

Third: MIP presupposes, as noted above, that we may be warranted in believing of certain conditions that they are illnesses. This presupposition (viz. that we have a sound concept of illness) is not something that we can fully defend here. However, fortunately for advocacy of MIP, there is an active conceptual cottage industry whose purpose is to regiment the notion of illness, some of whose output is directed at the notion as applied to mental illness (cf. Gert and Culver, 2004; Woolfolk 1999). Often, this is done in terms of the twin concepts of functional impairment and harm (cf. [Wakefield, 1992; 1999; Mechanic, 1999; Murphy and Stich 2000]; but compare with [Murphy and Woolfolk, 2001]). An illness is said to consist in the impaired and harmful functioning (or hurtful malfunctioning) of an organ system, part or faculty (including psychological faculty) of a person. Illustration: A heart's function is to pump blood because hearts' pumping blood in the past has given them a selection advantage and so led to the survival of more animals with hearts. Heart tumors that cause life-threatening heart malfunctions are illnesses (or diseases).

We use the word 'illness' as a general term for a negative state of unhealthiness that causes (or severely increases the likelihood of[4]) persons in them being harmed and that reflects impairment or disability in the operation of one or more parts or faculties of persons. We picture the relevant harm as, in some sense, stemming directly from the impairment in question. The impairment serves as triggering and sometimes also sustaining cause of harm. We use 'parts' and 'faculties' broadly. Bones are parts of persons. So, bone fractures are negative states of unhealthiness just when they impair the natural operation of bone structure (keeping the body erect, for example) and therein harm a person (causing severe pain or immobility, for example). To take another example: Visual perception is a faculty of persons. So, blindness is a state of unhealthiness since it means that the operation of visual perception is dramatically impaired. Blindness, of course, also undermines a person's freedom of movement and personal liberty and therein is harmful to its subject.

This is not a chapter on illness or the concepts of dysfunction and harm.[5] We plan to concentrate on MIP. So, here, as noted, we cannot fully defend or detail

[4]Reference to 'increased likelihood' is meant to cover a case like the following: a broken bone that has not begun to hurt but will under conditions of further decay or continued use. The unhealth of the break is a function, in part, of the relevant increased likelihood.

[5]If this was a chapter on illness per se, regimentation of the notion of illness would have to include discussion of the relationships between notions like the following: that of illness, disorder, disease, malady, and affliction. Each sometimes is or may be used to identify a negative or unwanted health state but not used synonymously. For example, one may wish to speak of some negative states of unhealth as diseases, in a semantically strict or narrowly defined sense of 'disease', which applies only to physical organ systems of a person (e.g. a diseased liver), and to reserve the word 'ill' for the whole person or for the subjective experience of disease (e.g. 'he feels ill with liver disease'). See Fulford [1989] and Mechanic [1999] for discussion.

our construal of the term 'illness' or explicate associated terms (e.g. 'harm'). But it may be noted in favor of our construal that by embracing our picture of illness a vexing problem associated with mental illness and mentioned earlier in the chapter might be solved. We will return to clarification of how we are using 'conscious representational content' in MIP in several moments.

4 MENTAL DISTRESS AND PROBLEMS OF LIVING

The problem that we are about to mention appears in one guise or another in different places in the mental health literature. These include discussions of philosophical counseling (cf. [Marinoff, 2000; Raabe 2000]), histories of concepts of mental health and illness (cf. [Radden, 2000]), cross-cultural analyses of mental illness (cf. Kleinman and Good 1985) and critiques of the over-medicalization of mental disturbance and distress (cf. [Horwitz, 2002; Wakefield, 1999]).

Mental illness or disorder, on our view, is a type of mental distress or disturbance. It is different from other types of distress. Some types are best classified as problems in living. The human condition is riddled with problems in living. Some are harmful or injurious. Abnormally prolonged or intractable grief over the death of a loved one is one such problem. In excess it can ruin a person's life. Some among other types of problems in living include certain sub-cultural patterns of social deviance. Various types of religious hysteria and vagabondage fall into this category of sub-cultural deviance. How should mental illness be distinguished from forms of mental distress that are problems in living and not illnesses?

The answer to this question that our use of the word 'illness' and related terms helps to warrant goes something like this: Provided that relevant psychological faculties of persons, such as memory, perception, and self-control (including control of self-destructive emotions; cf. [Flanagan, 2000]), among others, are in proper operation, mental distress (in which conscious representational content critically is positioned) fails to qualify as mental illness. If one or more faculties are both impaired and responsible for harm or the seriously increased risk of harm to a person, the distressful condition in question (depending upon the purport of its influence stream) might be better classified as a form of mental illness.

Consider, as an example of a complex form of mental distress that is not an illness, not grief or vagabondage, but the belief system of someone in *philosophical despair*. By this we mean, following Richard Garrett [1994], that a person in philosophical despair embraces pessimism about whether anyone's life is good or meaningful or has the potential of being good or meaningful. Human existence to a philosophically despairing person is believed to be utterly absurd or worthless or tragic. Bertrand Russell [2000] advocated philosophical despair. So, too, did Arthur Schopenhauer [2000], Richard Taylor [1970], and Joel Feinberg [1980].

Philosophical despair may be more than mildly distressing or disturbing. Indeed, some elements of the mood and belief system of a philosophically despairing person may be indistinguishable from those of someone who clinically is depressed. Nonetheless, the presence of similarities or analogies between depression and philo-

sophical despair does not mean that someone with this philosophic view is mentally ill. It is eminently possible to formulate a competent intellectual justification for a philosophically despairing attitude i.e. for the attitude that life is not worthwhile or is tragic, and to do this in competent exercise of intact psychological faculties (perhaps even escaping melancholic mood (cf. [Feinberg, 1980]). For this reason one should not presume that proponents of the attitude are impaired or ill.

Nancy Andreasen [1984, p. 144] appears to make just such a mistaken presumption in comments she makes about St. Augustine. She writes of Augustine as follows:

> No medical texts are extant from the early medieval period, but literary and historical evidence indicates that the absence of texts does not bespeak the absence of illness. St. Augustine has confessed to his struggles with the hopelessness and despair of depression.

Andreasen is referring to Augustine's soul-searching quest described in *The Confessions* to discover whether life truly is worthwhile. However because Andreasen seems to be confusing philosophical despair with depression, she describes the venerable saint's struggle in clinical or medical rather than spiritual or philosophical terms. Augustine did not confess to struggling with the hopelessness of depression. He confessed to his struggles with the spiritual emptiness of philosophical despair. Augustine was not on a prescient search for a compassionate psychotherapist or effective dose of fluoxetine (Prozac). He sought a cosmic meaning or purpose to human existence.

Philosophical despair per se does not reflect impairment. It may be a form of distress, certainly, but it is not an illness. Whether it should be classified as a problem of living may be wondered. Much depends upon just how the expression 'problem of living' is regimented and whether it is best reserved for cases of severe or hurtful distress. What, however, of mental distress that unlike philosophical despair is harmful but like philosophical despair does not reflect an impairment or disability in a faculty or part of a person? This might be a promising way in which to describe a condition as a problem in living. Problems in living, from a severity perspective, are harmful distresses that do not reflect impairment. Let's consider as illustration, briefly in this context, the question of whether alcoholism is a problem of living or mental illness.

Assume that human beings possess a psychological faculty for exercising prudent control over consumption of food and drink. Just how we describe this faculty is not something that we can discuss here, but, roughly, let's refer to it as a combination of conscious foresight and self-discipline. Clearly, failure to exercise prudent control and the disposition to be intemperate about alcohol consumption harms or seriously risks harm to its subject. The question of whether alcoholism is an illness, then, comes down, we assume, to the issue of whether the relevant faculties of foresight and self-discipline are impaired in persons with alcoholism. What has to be decided to settle this question?

Showing that alcoholics fail to exercise prudent control over their drinking does

not demonstrate that foresight and self-discipline are impaired. All sorts of plausible explanations may be given other than impairment to account for failure to exercise control. An alcoholic might have false beliefs about the costs and benefits of drinking. Or various emotional distractions might prevent them from giving sufficient attention to negative health consequences of drinking. The assumption we are making here is that faculty impairment is a matter of damaged competence or disability rather than of misinformed or distracted performance. Failure to perform is not proof of incompetence or impairment. If, therefore, it is shown, as some critics of the alcoholism-is-a-mental-illness hypothesis have argued that it can be shown (cf. [Fingarette, 1985]), that alcoholics do sometimes exercise prudent control over their drinking, then this would be a step towards showing that impairment or disability is not present and that, at least in some cases, alcoholism is a serious and harmful problem of living and not an illness. It is an inability rather than disability.

A quick disclaimer is needed. The purpose of the above remarks is not to claim that alcoholism fails to be a genuine illness. Indeed, alcoholism is an illness, on our view, if it reflects harmful impairment. It is only to suggest that the definition of 'illness' that we have adopted, by way of a background assumption for MIP, is on the right conceptual track insofar it helps to distinguish mental illness from other forms of mental distress including problems in living. In any case, let us turn to the central clause in MIP.

5 CONSCIOUS REPRESENTATIONAL CONTENT

What about the expression 'conscious representational content'? This expression is the linguistic centerpiece of MIP. The conscious representational content of an illness is said by MIP to be critical to classifying it as mental.

Presumably, of course, there is a difference between being conscious and failing to be conscious. You and I are conscious. A stone is not. Presumably, also, there are different types of conscious states or processes. Precisely what counts as a type of conscious state or process is not entirely clear, since one can distinguish types in a variety of different ways and for different purposes. We assume that the most interesting and pervasive types of conscious states are those that are identified or individuated in terms of their conscious representational content (sometimes called 'Intentional content'). Conscious content state types represent the world or self as being one way or another. They help to constitute a person's conscious conception of things or (to use a technical expression) phenomenal intentionality (cf. [Horgan and Tienson, 2002; Horgan *et al.*, 2005]). If the conscious representational content is veridical, the world or self is the way that the content represents it as being. If the content is in error or misleading (e.g. hallucinatory), the world or self is other than the content represents it as being. We mean to refer to such contents when we speak of 'conscious representational content'. Conscious representational contents consist of consciously representing the world or self as being one way or another.

Illnesses that are prototypically mental are constituted by conscious representational content. Normally, such content is more or less specific in type or kind to the disorder in question. So, the conscious representational content of depression is distinguishable (despite some overlap) from that of panic disorder. When a person suffers from a major depressive episode, for instance, the experience of being depressed typically includes thoughts of death, suicide or personal worthlessness. In panic disorder a person consciously believes themselves to be losing control or going crazy. In personality disorders with grandiose delusions people are convinced that they have great talents or supernatural powers. In symptoms of depersonalization parts of one's body or actions are consciously experienced as if not belonging to oneself. And so on and so forth.

Consider, to further illustrate the notion, an example of conscious representational content that is not associated with illness. Suppose I am dining in a restaurant.[6] Suppose I visually perceive something in my soup. Suppose I visually perceive this something as a bug. Or perhaps, alternatively, I perceive it as a leaf or as a small, unidentifiable foreign body. Either way or any way insofar as I consciously represent something in my soup, I experience it as something (as a bug, etc.). It is experienced or represented by me as being one way or another.

Now consider the concept of conscious representational content applied to a case that is a full-blown case of mental illness. Consider the case of Virginia Woolfe (1842-1941), English novelist and essayist. Woolf committed suicide by drowning herself, fearing a depressive breakdown (having had mental crises in 1895 and 1915) from which she could not recover. Apparently, Woolfe suffered from chronic depressive illness of a bipolar or manic variety. In a suicide note to her husband she wrote:

> I feel certain that I am going mad again: I feel we cant go through another of those terrible times. And I shant recover this time. I begin to hear voices, and cant concentrate. So I am doing what seems to be the best thing to do .... I don't think two people could have been any happier till this terrible disease came. I cant fight it any longer .... Everything has gone from me but the certainty of your goodness. I can't go on spoiling your life any longer.[7]

Quentin Bell wrote of her depressive episodes: 'Her sleepless nights were spent in wondering whether her art, the whole meaning and purpose of her life, was fatuous, whether it might be torn to shreds by a discharge of cruel laughter.'[8] Of her manic episodes, Leonard Woolfe wrote: 'she was extremely excited ... talked volubly and, at the height of the attack, incoherently; she ... heard voices.'[9]

It should be noted that the overall character of Virginia Woolfe's manic depression is a structurally rich (albeit profoundly disturbing) conscious representational

[6]The first person singular is occasionally be used as a stylistic device.
[7]Quote as cited together with primary source in Slavney and McHugh [1987, 31].
[8]Cited in Slavney and McHugh [1987, 31].
[9]Cited in Slavney and McHugh [1987, 31].

experience of both self and world. The representational content is of an apparent whole array of events, activities, properties, and relations — including the experience of voices, sleepless nights, fear of laughter, doubts about her art, and so on. It is the complex experience for her of appearing to be a disturbed and disturbing wife and writer. One may engage in the behavior of writing a suicide note, for example, because one is bored, play-acting, or exploring new literary styles. Or one may do this because one is Virginia Woolfe, depressed and intending to discontinue living. The note was designed by Woolfe as a suicide note and not a literary exercise.

Woolfe's illness 'shaped her development as a person and as a writer, affected her closest relationships, and eventually claimed her life' [Slavney and McHugh, 1987, 116]. The conscious representational content specific to her illness critically and, indeed, tragically, mattered to the illness. The content contributed to its onset or origin. The content also figured among its central symptoms and helped to lead to her suicide. (It might also have been the focus of effective treatment.)

When conscious representational content enters into the onset or origin of an illness, and then ramifies through symptoms and consequences, this constitutes something we call a *stream of influence* through different phases of an illness, its symptoms, and consequences. The manner in which both self and world appeared to Woolfe was pivotal or critical not just in producing her depressive episodes, but in accounting for their symptom expression and consequences. The content exerted a stream of influence from one phase of her illness to the next.[10]

It should be noted that in proposing that conscious representational content is the form or mode of mentality to which reference should be make in a description of the prototype of mental illness, we do not mean to exclude non-conscious mentality from playing important roles in mental illness. Perhaps Woolfe had unconscious or implicit representational content as part of her depression. Perhaps these included unconscious emotional memories triggered by real or imagined failure in her literary career. But it should be noted that MIP's reference to conscious representational content as prototypically central to mental illness does assume that conscious representational content is the form of mentality that operates in clear-cut or exemplary cases of mental illness. Non-conscious activity, as such, on our view, is not prototypically characteristic of mental illness.

6 RESTRICTING THE MENTAL

Remember that in an effort to regiment the notion of mental illness, we aim to identify what makes mental illness mental. We are trying to distinguish mental illness, paradigmatically speaking. So, we need to identify a form of mentality that

[10] Remember, on our view, if Woolfe truly was mentally ill she also had to have one or more of her faculties impaired or disabled — in a manner revealed by her symptoms. Fuller exploration of her case than can be offered here would need to explore which faculties these might have been. Additional conceptual issues would be raised about what makes a faculty a faculty as opposed, say, to a mere disposition or pattern of performance.

clearly is mental and such that if this form of mentality matters critically to an illness, then the illness is sufficiently mental so as to count as mental illness. Since conscious representational content counts as mental if anything does, then, if we are right, no one should deny that an illness in which conscious representational content matters critically (along the four dimensions mentioned) is a case of mental illness.

Outside the scope of the prototype just how generous is the very idea of mentality in mental illness? The extension of 'mental' in discussions of mental illness sometimes possesses much wider range and referential generosity than that of conscious representational content. DSM-IV, for example, appears to permit classifying various motor deficits or movement disorders as mental disorders (e.g. [APA, 1994, 53-55]). In such cases 'mental' seems to be conceptually functioning as applying also to 'sub-personal, sub-conscious informational' and so mental illness is presumed constituted by deficits in sub-personal information processing. We doubt, however, that sub-personal, sub-conscious informational deficits may suffice to characterize an illness as mental. Should the very idea of the mental in mental illness be *that* broad?

The philosopher Fred Adams [2003] has argued that some theories of representational content attribute content to things that they shouldn't. Some theories attribute too much content, as Adams notes. Adams describes this as the problem of *semantic promiscuity* for a theory of content. In the case of a theory of mental illness, we wish to call the analogous problem the problem of *mental promiscuity* for a theory of mental illness. Classifying sub-personal, sub-conscious informational content as 'mental' is too promiscuous and not consistent with a notion of mental qua mental illness that is even modestly regimented, exemplary or prototypical.

This is not to deny, of course, that there may be both utility and wisdom in referring to sub-personal, sub-conscious information processing to help to account for certain features of mental illness as well as, of course, human and animal behavior generally (cf. [Bechtel *et al.*, 1998]). However let's be careful here. Does the general utility of reference to sub-personal information processing mean that we should identify mental illness *per se* by reference to non-conscious information processing deficits full stop? That kind of conceptual generosity towards the very idea of mental if taken as constitutive of mental illness makes it a mystery, we believe, how mental illness can be distinguished from non-mental illness and opens up the unwelcome taxonomic possibility that all sorts of illnesses (even those outside of the already porous boarders of DSM; cf. [Poland, *et al.*, 1994]) will qualify as mental illnesses. Let us explain.

Some types of informational content enter the world exclusively as the result of objective probability relations between events [Dretske, 1981]. The movements of a barometer needle (when operating properly) carry information about changes in atmospheric pressure. The posture of a joint carries information about perturbations in the shifting equilibrium points of muscles [Brooks, 1986, 204-205]. However barometers cannot be mentally ill. Meanwhile ideomotor apraxia is a disorder in the coordination of movement. So, describing such a disorder as mental

would be an unwarranted verbal gloss on a preferably more strictly neuromuscular description of the disorder as a non-mental physical illness. We don't need to apply a notion of mental illness to that sort of disorder anymore than we need a notion of mental disorder to illuminate the dysfunction in broken barometer needles.

Sub-conscious, sub-personal information processing deficits of different sorts surely play causal-explanatory roles in mental illness. However, we contend, such processes do not play a role, prototypically, in the concept of mental illness. Such deficits do not help to distinguish illness that is mental in a clear-cut manner. If they did, then an illness such as osteoarthritis should count as mental, since, presumably, disorder in sub-personal information processing occurs in the joint and skeletal system in victims of the illness. However osteoarthritis is not, by any stretch of the theoretical imagination, a mental illness.

7 BROKEN BRAIN

Reflecting on the unneeded and unhelpful conceptual promiscuity of including deficits in sub-personal, sub-conscious information processing as characteristic of mental illness may produce the contrary worry, in some readers, that clinging to a concept of mental illness like that represented MIP, which is wedded to the proposition that conscious representational content is critical to mental illness, is an unwelcome theoretical conservatism. The worry is that MIP is not at the hard-nosed medico-scientific cutting edge in describing mental illness. It might be charged in two sentences: 'No illness is ever mental.' 'So-called mental illness when a genuine illness just is a disease of the brain.' Or as Michael Alan Taylor [1999] puts it, 'psychiatry and neurology [constitute] one field.' 'Mental illness is . . . not 'mental' at all, but the behavioral disturbances associated with brain dysfunction and disease' [p. viii].

Mental illnesses are brain diseases or disorders. That looks like a proposition that an anti-dualist of physicalist disposition would surely embrace. But one must be careful.

In the first place, advocates like Taylor for replacing the concept of mental illness in favor of that of behavioral disturbance caused by a (to appropriate Nancy Andreasen's [1994] expression) broken or diseased brain presuppose, of course, that it can be determined just when the brain is broken or diseased. Advocates, however, sometimes overlook the fact that whether the brain is broken or diseased is unintelligible without normative judgments about the brain's contribution to our leading healthy conscious and satisfying social lives. And it must be remembered: we do lead conscious social lives.

An important conceptual constraint on conceiving of mental illness as brain disorder is that the very idea of brain disorder is governed by our sense of just what it means for the brain to be healthy or properly function. Meanwhile, our sense of just what it means for the brain to properly function is bound up with our picture of healthy conscious and socially situated lives and of how the brain contributes to those lives. Consider the following analogy. The very idea of a 'healthy' or properly

functioning automobile engine is not derived from our picture of the basic physics or quantum mechanics of locomotion. We explain how the engine works by taking it apart and examining its functional units (cylinders, pistons, rods, etc.) and the roles that each of these units play in generating motion in environments for which the engine is suited. Each unit makes a contribution that we articulate in terms of producing locomotion. If locomotion occurs we assume that the engine — its cylinders, pistons, etc. — is functional and not broken. If locomotion fails to occur we look for either a dysfunction or impairment (e.g. busted rod or cylinder) or the engine's foiled attempt at locomotion in an environment for which it is ill suited (e.g. it is trying to move a car through 3 feet of northern Ontario snow).

Something analogous is true of the brain. Whether we understand the brain as healthy or diseased depends, in part, upon judgment of whether our conscious representational and social lives — of motivation, affect, decision, belief, social relationships, interpersonal behavior and so forth — appear to be healthy and well. If our conscious and socially situated lives seem healthy and well, we assume that to this extent (independent of whatever brain activity is associated with non-mental forms of wellness) our brain is properly functioning and healthy. If our conscious socially situated lives appear unhealthy or unwell, we may appeal to brain dysfunction or neurological disease to help to explain poor mental or social health. Or we may search for a disconnection or mismatch between an otherwise healthy brain, on the one hand, and environmental circumstances, on the other, for which the brain is ill-suited (e.g. the brain is skidding over social landscapes packed ice tight with learned helplessness reinforcement schedules). Either way, rather than conceiving of the brain as an organ that operates in a conscious representational vacuum, we should recognize that the healthiness of neural activity is judged in goodly measure in terms of assumptions about lives of conscious representational sorts. So, if mental illness is a brain disease this actually may require conceiving of mental illness as a disorder in representational consciousness (in the manner of MIP). Absent disorder in our conscious representational lives, if other forms of wellness are in place, there is no reason to believe that the brain is diseased.

Of course mental symptom expression (of unhealthy conscious life) is insufficient, by itself, to decide whether an illness is mental. Untreated phenylketonuria leaves individuals with an IQ less than 50. But the mental deficiency is due to a defect in the gene for phenylalanine hydroxylase (cf. [Heninger, 1999]). Causes of the illness are not states with conscious representational content. So, the illness fails to qualify as mental — at least by the lights of MIP.

We don't know how the physical states of the brain give rise to states with conscious representational content. So if conscious representational content matters to illness, though we might postulate intimate connections or even identities between conscious representational states and brain states, there are so many questions about the precise formulation and explanatory significance of such connections, that there is plenty of epistemic elbow room to maintain that conscious representational content plays a causal role in the onset of some illnesses viz. mental illness as understood by MIP. But does it?

Overall we believe that the proposition that some illnesses include states with conscious representational content among their causes is a coherent, plausible view without fatal counter-evidence and with much supporting positive evidence. We will concentrate on one class of such illnesses: those in which there is a strong social environmental interactive contribution to the onset of the illness. Truth be told, there is no way to understand how the social environment plays a role in the onset and development of some illness that does not make reference to a victim's conscious representation of his environment.

To illustrate: Suppose I have a panic attack. My brain activity somehow carries or produces conscious representational content about the environment as possessed of such-and-such panicky properties. Whether the environment does offer such properties depends upon the representational content of my consciousness. The facts by reference to which the onset of a panic attack is explained therein include reference to how the environment is represented by me. And, it will be recalled, this is one dimension along which, according to MIP, conscious representational content figures in prototypical mental illness. How the world (or self) appears to the subject of illness is constitutive of the onset of mental illness.

Let us consider two ways in which reference to conscious representational interaction with the environment plays a role in the explanation of the onset or origin of mental illness. One we shall call *situation sensitive* conscious interaction. The other we shall call *culturally scaffolded* conscious interaction.

8 CONSCIOUS REPRESENTATIONAL CONTENT AND ENVIRONMENTAL INTERACTION

Barbara Grizutti Harrison, recipient of the O. Henry Prize for short fiction, was a victim of panic disorder. She describes the conscious representational content of her panic attacks as follows:

> I am living inside my fear . . . I cannot bring myself to walk the familiar path . . . I shrink the world around me and cocoon myself. . . A panic attack as the force of an oncoming train. Panic is anxiety severely, grotesquely heightened. . . Life is hard; perhaps you've noticed. [Harrison, 1998, 3-7].

DSM-IV includes symptoms like those mentioned by Harrison in its list of symptoms of panic disorder, adding that victims often feel that 'they are 'going crazy' or losing control or are emotionally weak' [APA, 1994, 398]. A sense of 'pain . . . grotesquely heightened', to use Harrison's expression, is part of the specific conscious representational content of the disorder as well as of what DSM-IV means when it reports that victims feel that they are losing control or going crazy.

What of the environmentally situated onset of attacks? James Ballenger [2000] notes that the highest reported rates of panic disorder 'occur in widowed, divorced, or separated individuals living in cities' [p. 807]. Assuming that reported rates are comparable to actual or real rates, why is that? A promising hypothesis

offered by Ballenger (which we will endorse here without argument, referring the reader to Ballenger and elsewhere[11] for its defense) is that such individuals live in environments with robust 'stressors' or situations that cause or induce panic. These stressors help to produce panic attacks in people of certain personality types or with certain learning histories.[12] Relevant city-wise stressors for some people (especially those who have panic attacks with specific sources of anxiety such as agoraphobia) may include public transportation, shopping malls, tunnels, elevators, and crowded city streets.

Mention of the explanatory need to refer to stressors prompts an observation and requires a distinction. What stressors are found in cities or elsewhere depends upon the definition of 'stressor', an issue more conceptually elusive than it may first appear.[13] There is a big difference between physically defined and psychologically salient stressors. The stressors just mentioned — tunnels, elevators, crowded city streets, etc. — concern, in some sense, physically defined stressors. Each has distinctive physical features or configurations of parts or people in those configurations. Tunnels are designed one way. Crowded streets are configured another.

However there obviously is another side to stressors. Why should certain objects or situations induce 'stress' in some people and be responsible for onset of panic attacks? The answer, we presume, doesn't lie in physics, optics or neurophysiology. It does not consist in reference to architectural features or to the numerical population density of tunnels or city streets. It requires reference to the consciously represented interactions of susceptible persons who interact with stressors. Stressors, in the salient sense, consciously mean or signify something to persons whom they distress. Being in a tunnel for someone may appear like an uncontrollable threat to her physical safety. Walking on a crowded street may imply to a person that he is living in isolation and loneliness among strangers and from which there is no escape. Either of these contents may contribute to panic. Consequently, if an illuminating appreciation of the role of stressors or situational

[11]See also Cockerham [2003, 76-89] and Wheaton [1999].

[12]'Stressors', so called, should not be confused with what are sometimes called 'environmental insults'. Insults are identifiable environmentally induced neuromolecular traumas (such as malnutrition, viruses and toxins) which cause cell death or interfere with critical transduction pathways and leave distinguishable neurochemical footprints. Although some insults are stressors (insofar as persons react to the insults as stressors) and some stressors are insulting (in the previous mentioned technical sense), there is no necessity (other perhaps than the expectation on the part of investigators) that a stressor will contribute to cell death or interfere with brain circuits or that insults will register in behavioral disturbance.

[13]We are about to introduce a distinction between 'stressor' defined in non-stress nominally physical terms versus in stress related terms. We must also distinguish between stressors that are essential or internal parts of an illness (such as a frightening delusion to a schizophrenic) and external stressors that help to produce instances or episodes of an illness (such as uncertain threats to physical safety as antecedents of a panic attack). Another useful distinction is between proximate stressors, which occur just prior to an episode or instance of illness, and distal stressors, which have occurred previously but with long-term effects (see [Zuckerman, 1999, 424-425].) We aim to be talking, in what follows, about proximate external stressors that help to produce instances or episodes of an illness.

factors in understanding a disorder such as panic attacks is sought, it must be noted that there is nothing intrinsically stressful about the physicality of objects like tunnels or crowded streets. Scores of people are exposed to these situations but never experience panic. Clearly most people cope, master, or buffer against anxiety. A relative few, however, do not. For them panic is a response to those situations: to the environment as represented in the conscious content of a panic attack. Certain aspects of those situations 'say' or 'mean' something panicky to such people.

What role or function is there for the brain (broken, impaired or otherwise) in stressor induced or situation sensitive panic attacks? Regardless of the precise and ultimate neuroscientific details, whatever the brain does in such situations, it encodes information and carries memories about stressors as well as relevant learning histories. Some of this information is somehow translated or readied for presentation in consciousness in the form of conscious representational content as when, for example, a victim explicitly believes that her physical safety is threatened in a tunnel and this leads to panic.

If conscious representational content is not at least in some measure responsible for the situational sensitivity or environmental embedding of a panic attack, then two other hypothetically explanatory possibilities may come into play. One is that a derivative of conscious representational content (e.g. unconscious memory of prior attacks themselves conscious) is potent in an attack. The other is that reference to a person's sub-personal neural activity in tunnels and streets together with reference to physical properties of those situations might somehow account for the attack. But note that the first explanatory option continues to give conscious content etiological traction in the account, since it backtracks to previous instances of panic and to memory traces of those experiences. Whereas the second option must cite the *situational identity* (to a victim of panic attacks) of circumstances like tunnels and crowded streets in order to hold the stressor class together as a distinguishable or identifiable stimulus class across a wide array of realizations or instances. We seriously doubt whether causally potent identities of stressor classes are explicable by reference, for example, just to the optical properties of the ambient illumination of tunnels or crowded streets given the multiple realizations and diverse lighting conditions associated with such circumstances. So, reference to what stressors 'say' or 'mean' to susceptible people seems unavoidable as a means, either immediately or derivatively, to account for the onset of panicky disturbance. Such reference grounds our ability to group or classify city-wise configurations into instances of relevant stressors.

We are not saying that mental illness (or even panic attacks) is always situation sensitive. We are not saying that the sole purpose of reference to conscious representational content in the origin or onset of mental illness consists just in explaining situational onset or sensitivity. Rather, we are saying that some mental illnesses owe their origin or onset to situational circumstances and these are characterized by how environmental situations are consciously represented.

Obsessive-compulsive disorder provides a second example of a mental illness with situation sensitive and representationally informed onset. The door that may be unlocked or the hands that may be unclean are consciously experienced in the case of some victims of OCD as an occasion for 'that awful anxiety (that) keeps me checking', as one victim reports [Rappaport, 1998, 16]. Other illnesses with situational sensitivity and thus contribution from conscious representational content may include some cases of major depression, post-traumatic stress disorder, and borderline personality disorder.

Meanwhile situational sensitivity is not the only manner in which the conscious representational content of environmental interactions is linked causally or influence stream-wise to illness onset and development. Secondly, there are cases of (what we call) *cultural scaffolding* of mental illness. Here, again, reference to how the environment appears to a victim is necessary in order to explain the origin or occurrence of an illness and, in particular, its symptom expression. We must be brief.

Consider impulse control disorder. Types of this disorder appear to partition along gendered lines. A compulsive male might become a victim of obsessive compulsive disorder (OCD) in one social environment or a gambling addict in another. A compulsive female might develop anorexia nervosa. Social context including culturally conditioned gender expectations seem to exert a powerful influence on types of disorder and symptom character. Linking particular types of disorder to cultural contexts without reference to conscious representational content appears impossible. One female victim of anorexia nervosa reports: 'I knew that I had a certain strength ... that would really show up somewhere.' 'I skipped breakfast.' 'I just couldn't fit the calories into my regimen.' 'I always 'watched it" [Costin, 1998, 243-244]. A male gambler notes: 'Winning big made me feel important ... and potent' [Heineman, 1998, 161]. How a thin female body or a male gambler's winning schedule appears to the person — what it consciously means to its subject — plays a critical role in the occurrence and consequence of these compulsions. Different people have different reactions to the appearance of thinning bodies or to winning at horses. Most people of course do not become compulsive about such activities. If anorexia or heavy betting is to have any explanation, it seems, reference to aspects of the environment (to the fly in the soup as it were) as they 'consciously speak' to a subject must be part of the story.

The cultural scaffolding of mental illness has received considerable scrutiny in the sociology of mental illness and cross-cultural studies of symptom expression (cf. [Horwitz, 2000; Cockerham, 2003; Schumaker, 2001]). Horwitz refers to it as the structuring of mental illness [Horwitz, 2000]. 'In the broadest sense,' he writes, 'symptom profiles are structured to fit the illness norms of particular cultures' [p. 116]. The symptom character of dissociative identity and recovered memory disorders, eating disorders (such as anorexia and bulimia), depression, and a host of other illnesses or disorders seem to rest in part on cultural and therein conscious content scaffolds.

Unfortunately the picture of mental illness as that of a broken or diseased brain often is the result of an environmentally de-situated and 'aconscious' or solipsistic caricature of what takes place in mental illness. Existing evidence for the etiology of mental illness sometimes focuses on the impaired brain as the proximate or immediate cause of mental disorder and neglects extra-neural environmental components and the causal explanatory role of conscious representational content in a subject's interaction with those components. De-situated, 'aconscious' neural focus simply buries interesting constituents of the etiology of mental illness in superficial neuroscience and shallow interpretations of the thesis that mental illness is a brain disease. However such explanatorily irreverent dismissals of consciousness, we hope, are not the necessary result of taking the brain seriously in an informed picture of mental illness. Taking the brain seriously means, in part, taking its functional role in conscious environmental interactions seriously — as the engine (as it may be put) of representational consciousness.

Andreasen [1984, 219], who does know better (see below), remarks 'when we talk, think, feel, or dream, ... these mental functions (are) due to electrical impulses passing through ... the human brain.' 'Mental illnesses are due to disruptions in the normal flow of messages through this circuitry.' If one agrees with Andraesen one might be tempted to refer only to disrupted electrical impulses or to some other sort of neural disorder rather than also to conscious representational content to account for panic attacks and compulsive gambling. In the case of gambling, for example, one may wish to refer to 'disrupted' neuronal activity in the ventral tegmental area that is involved in reinforcement processing as part of the explanation for the behavior [Montague and Dayan, 1998]. However we cannot guess why anybody would be tempted to promote such a thin and partial explanatory strategy independent of acknowledging the role of conscious representational content in situation sensitive and culturally scaffolded forms of mental illness. Reinforcement does not occur in a conscious vacuum. Neither does mental illness.

The overall explanatory problem or project facing a brain centered picture of mental illness is how to characterize the distinctive roles both of neural activity in mental illness as well as of conscious representational content in an illness's onset and influence stream. As we see it, no matter future developments in this research project, the following proposition already seems undeniable. Reference to neural impairment or activity without appealing to conscious representational content isn't sufficiently able to account (among other things) for the onset (or course) of mental illness that possesses social sensitivity or situatedness or is culturally scaffolded.[14]

[14]Leslie Brothers [2001], see also [Brothers, 1997] uses the expression 'social neuroscience' to refer to a research program which would consist of developing models of brain activity that is selectively responsive to consciously represented aspects of the social and cultural environment. It's not clear precisely how Brothers believes that social neuroscience is to be conducted or what its environmental interactive data should be. However the phrase 'social neuroscience' is helpful for marking the task facing a friend of MIP as well as of the proposition that mental illness ultimately is a brain disease. The task is to investigate the causal-explanatory roles

9 MENTAL ILLNESS REGIMENTED AND TREATED

The concept of mental illness would be regimented with trumpets if it was decided that MIP is the preferred concept. Suppose, then, that thanks to the efforts of numerous books and articles, and robust clinical and experimental evidence, MIP has been accepted as the preferred concept of mental illness. The DSM of the future, say, uses it. Suppose that we describe many individual cases of illness so that they conform to the prototype. Some conform in all four critical aspects; others conform in most aspects. Here's one instance of illness that we are going to stipulate or assume (for purposes of illustration) conforms in all four aspects.

> The patient, a male professor of mathematics, comes to the clinic seeking antidepressant medication and medicine for sleep disturbance. He reports that his appetite has diminished, and that he lies awake long into the night contemplating his death. He reports that he feels guilty for having squandered a Ph.D. in mathematics from Princeton. He is fatigued and lacks energy. He reports that he feels extreme sadness and emptiness, as well as absence of pleasure in daily activities. He says that he continually shows up for class late and unprepared, and a colleague has reported that he told her that he wished to commit suicide. The professor is barely managing to hold back tears.

The patient described in the above vignette is properly diagnosed, we may suppose, as suffering from major depression. Of course, the distress experienced by the professor is similar in some ways to non-illness distressed mood, but suppose that the patient has been referred to the clinic for depression. This is not his first major depressive episode and, apparently, without treatment, it will not be his last. Suppose we have MIP in hand.

We describe the professor's symptoms as follows. Symptoms with conscious representational content include depressed mood over the course of his life or circumstances, diminished pleasure in activities, recurrent thoughts of death and suicide, and feelings of guilt over a wasted graduate career. Then, suppose we decide that the professor's depressed mood and various related contents are the partial causes of various other symptoms of his condition, including fatigue and loss of energy. We believe, for example, that because he is depressed, the professor arrives late for class and is unprepared. We identify an array of such symptoms and consequences. (Remember symptoms and consequences are two general ways in which representational conscious content critically matters to mental illness.) The list might continue and the vignette expanded, but what of treatment? Thus far we have said nothing in this chapter about treatment.

Suppose we try to identify sources or factors contributing to the professor's depression. The professor's depression impresses us as due, in part and less proxi-

associated with the conscious representational content of environmental interaction at several levels of detail or resolution (including prominently the neural), beginning with facts like the following. Crowded streets may signify loneliness. Tunnels or subways may say that one's personal safety is threatened.

mately, to various attitudes that he acquired during early stages of his professional career.[15] These attitudes may include faulty criteria of personal worth, grossly inappropriate vocational aspirations, over-generalized negative emotional judgments, and other disturbing attitudes. We note that his depression is situation sensitive as well as culturally scaffolded, although, again, details of all this (since the case merely is illustrative) are omitted here. Suffice it to note, now in addition to recognizing the role of conscious representational content in symptoms and consequences of his illness, we develop a picture of the influence of such content in the onset or origin of his depression. Therein we believe ourselves prepared to introduce a treatment regimen.

Some forms of treatment for depression are sensitive to and directly address pathogenic attitudes that are harbored in a person and which surface in conscious symptoms and behavior. Typically these treatments are embedded within a general psychotherapeutic orientation such as (to list names of popular treatments): cognitive therapy, cognitive behavioral therapy, rational emotive therapy, and interpersonal therapy.

The general goals of treatment for depression consist in providing a desirable and believable measure of optimism, addressing accumulated intra-personal conflicts, blocking overgeneralization of negative emotions, and gradually mobilizing a person's skills at recognizing situations that may be responsible for the onset or persistence of depressed mood. For example, blaming one's self for wasted opportunity often requires little real or objective failure. One doesn't have to consistently falter to accuse oneself of wasting a Princeton Ph.D. Questions of whether personal blame truly is deserved are perhaps best put to the test by counter-consideration of particular forms of competence and concrete achievement. In the deflated world of a depressed person subjective standards of success or wisely chosen opportunity tend to be too demanding and unrealistic. Witness the case of Virginia Woolfe. Her standards should have been less elevated and more self-forgiving (on self-forgiveness see [Garrett, 1989].

The general conception of treatment to be embraced in the hypothetical case of the professor, which might be an eclectic mix of various forms of therapy (perhaps even including parts of psychoanalysis), we label 'conscious representational content focused' or 'content focused' for short. Content focused treatment rests on the following three premises, two regarding the nature of an illness and the third regarding responsiveness to treatment or therapy:

(1) *Cognitive competence.* The illness does not prevent the person either from understanding the conscious representational content of the illness as symptoms of the illness or from appreciating factors of personality or situation that may be responsible for its origin or persistence.

[15] His susceptibility to this information may be due, in part, to abnormally high levels of adrenal stress hormones or to some other neuronal impairment or disruption (such as an overactive amydala). We shall not speculate here on the details (see [Hobson and Leonard, 2001, 170f] for discussion).

(2) *Motivational focus.* The illness does not undermine the patient's desire or willingness to improve his or her condition (i.e. to eliminate or ameliorate the illness).

(3) *Therapeutic feedback.* For any given symptom the person is able to understand and appreciate when improvement has or has not been achieved. He or she recognizes positive progress as well as negative outcomes.

One may weaken the set by restating and lowering the demands of each premise. Some illnesses or symptoms, for example, strip a person of insight into the fact that they are ill. In such cases a less cognitively demanding form of cognitive competence than (1), which requires illness and symptom recognition, may permit some version of content focused treatment. We shall not explore such revisions or the conditions that may warrant them (delusions characteristic of schizophrenia are among them; see [Stephens and Graham, 2004; 2006]) or therapies that address them here.

Let's assume that each of the above premises is satisfied in the case of the professor. He knows that he is not well and wants treatment. At the same time, he knows that much of his conscious life is taken up with being depressed. Moreover he is able to recognize when symptoms are being ameliorated or eliminated. He does not go blithely when therapeutic effort fails. He is receptive to alternative therapeutic approaches.

Content focused treatment is the subject of rancorous debate in the therapy literature. Some of this is focused on the comparative merits of different therapeutic orientations. Some is aimed at the merits of combining therapies (including drug and content focused therapies). Some is targeted at the economically overweening character of the therapeutic profession and medical treatment sub-culture. We cannot begin to address various terms of that complex debate here. We just wish to mention the following. Content focused treatment represents the fourth manner in which conscious representational content matters to mental illness, prototypically speaking. Conscious representational content is the direction of content focused treatment. The focus certainly may not be sole or exclusive. It may be combined with therapies, including drug therapies, which although not content focused, address behavioral disturbances associated with illness, such as in the case of depression, sleep disturbance, whose reduction or amelioration contributes to receptivity to content focused treatment. A tired professor is ill suited for rational emotive therapy.

When the therapeutic goal is illness reduction or elimination, and the course is content focused, therapy can proceed by attempting modifications of the 'unhealthy personality' behind the content of the illness. Or it can address, if they exist, situations or sub-cultures that help to provoke or scaffold the illness. Address may take the form of avoiding situations that elicit, say, in the case of the professor, negative moods. Avoiding certain situations may be no easy matter, but the emphasis is more specific than transforming a person's worldview or personality type. The recovering gambler is better equipped to avoid gambling if he associates

with non-gamblers and commutes away from the track. The depressed professor is better able to avoid depressed mood if he cultivates relationships with colleagues of prudent aspiration and refrains from submitting manuscripts to highly selective, severely judgmental journals. The city-dwelling victim of panic attacks may travel by taxi rather than subway. Situated change or address may also be achieved often more prudently or usefully (especially when certain situations cannot be avoided) by various techniques of desensitization that render stressful situations and cultural factors less potent. Circumstances may not need to be avoided if they can be encountered with new conditioning histories or behavioral skills. So, for example, the depressed professor and his therapist together might identify situations that trigger thoughts of wasted career and other misspent opportunities. The professor might be taught to correct those thoughts or substitute more positive self-assessments in their place.

Personality transformation of a pan-situational or cross-cultural sort is a therapeutic exercise without much current guidance as compared with the more specific strategies of desensitization, avoidance, and the like. However let's suppose that by the time, say, DSM-VII appears we have learned a thing or two about how to transform the personalities associated with some forms of mental illness. An influential program of transformation has worked its way into therapeutic practice — perhaps combining elements of philosophical counseling with therapeutic tools from social or preventative medicine central to mental health. These may include imitative modeling (emulating the behavior of healthy people) and understanding and appreciating lifestyle risks and liabilities.

It may also be noted, while on the topic of treatment, that discovery of non-mental origins for illness does not automatically preclude therapeutically addressing such illness, at least in part, with content focused treatment.[16] To support and clarify the point, let us turn abruptly to schizophrenia. Suppose that schizophrenia does not conform along each and every dimension of MIP. Suppose, in particular, that it is non-mental physical (perhaps genetic or neurochemical, since clearly dopamine is critically involved) in etiology. Even so, clinical evidence suggests that content focused cognitive behavioral interventions may sometimes be useful in the diminution or reduction of the delusional experiences that constitute some of the conscious content (including attention deficits) characteristic of schizophrenia.[17] Content focused treatment is impossible in severe delusions, where minimal cognitive competence is compromised. However in less severe cases it may foster diagnostic insight[18] and reduce the resistance to extinction or incorrigibility of delusions.

The possibility of limited application of content focused treatment for schizophrenia contains a general moral about the concept of mental illness that applies to

[16] Partly because we have written of symptoms of schizophrenia elsewhere [Stephens and Graham, 2000] but mainly because of the controversy and complexity surrounding this illness we have as yet not discussed schizophrenia in this chapter. That is about to change, albeit briefly.

[17] See [Fowler *et al.*, 1995; Kingdom and Turkington, 1994; Munro, 2000].

[18] Absence of diagnostic insight is a hallmark of delusional experience. See [Fulford, 1994; Stephens and Graham, 2004; 2006] for discussion.

MIP. Elements of a prototype typically have a weighted or graded structure.[19] Some count more heavily towards application of a concept than others. A fuller investigation of MIP than we can offer in this chapter would propose comparative weightings among the four ways in which conscious representational content matters to mental illness. Whatever the precise details, we believe, the less robust the causal role of conscious representational content in the origin or onset of illness, then the less clear-cut or prototypically mental is the illness in question. So, if an illness such as schizophrenia is non-conscious physical in its conditions of onset or origin, then it lays weak claim to counting as a mental illness. Meanwhile, although successful content focused treatment for an illness surely does not by itself make for a clear-cut case of mental illness, failure of such treatment when combined with non-content origin makes it difficult to conceive of any illness as mental no matter the vividness of its conscious symptoms.

10 CONCLUSION

In this chapter we have tried to tackle a problem — or pair of problems — about mental illness that remain with us despite decades of analysis by philosophers and others. These are the (what may be called) *Two Problems of Mental Illness*. One is the problem of how something mental can be ill (and not just a problem in living) and the other is the problem of how an illness can be something mental (and not just, in some restricted sense, a brain disease). Regarding each problem, and we have focused on the second, disagreement reigns among mental health professionals. So 'the' problem of mental illness still is with us.

It should be mentioned that we surely are not alone in saying that conscious representational content matters to mental illness, although our terminology and emphasis is different from others. The physician and philosopher Karl Jaspers[20] and the tradition of descriptive psychiatric phenomenology that he helped to found vigorously endorse the proposition that understanding the conscious representational experience of victims of mental illness is critical to understanding and treating mental illness.[21] Numerous efforts by both psychiatrists and philosophers have been made to deepen our appreciation of what mental illness is like for those who suffer from it and to deploy such appreciation in characterizing the role of conscious representational content in mental illness.[22] On one construal of the intent of this chapter, we are hoping to reinforce the general wisdom of such attempts. But we are also trying to avoid a promiscuous view of the mental in psychopathology and to argue that conscious representational content has a fundamental place in

[19] See [Hempel, 1965; Laurence and Margolis, 1999].

[20] See [Jaspers, 1997].

[21] For introduction to and discussion of descriptive psychopathology, see [Fulford, Thornton and Graham, 2005] as well as [Doerr-Zegers, 2000].

[22] See, for one of dozens of examples in the philosophical literature, Stephens and Graham [2000]. Some additional references to this literature may be found in [Graham, 2002] as well as, more fully, in [Fulford, Thornton, and Graham, 2006].

mental illness. There is something both self and world immediately appear like to victims of mental illness. To fail to recognize this fact is to fail to be prototypically mindful of mental illness.[23]

BIBLIOGRAPHY

[Adams, 2003] F. Adams. Thoughts and their contents: Naturalized semantics. In S. Stich and T. Warfield, editors, *The Blackwell Guide to the Philosophy of Mind*, pages 143–171. Malden, MA.: Blackwell, 2003.

[Andreasen, 1984] N. C. Andreasen. *The Broken Brain: The Biological Revolution in Psychiatry*. New York: Harper & Row, 1984.

[Andreasen, 2001] N. C. Andreasen. *Brave New Brain: Conquering Mental Illness in the Era of the Genome*. New York: Oxford University Press, 2001.

[Association, 1994] American Psychiatric Association. *Diagnostic and Statistical Manual of Mental Disorders, 4th Edition*. Washington, D. C.: American Psychiatric Association, 1994.

[Ballenger, 2000] J. Ballenger. Panic disorder and agoraphobia. In M. Gelder, J. Lopez-Ibor, and N. Andraesen, editors, *The New Oxford Textbook of Psychiatry*, volume 1, pages 807–822. New York: Oxford University Press, 2000.

[Beam, 2001] A Beam. *Gracefully Insane: The Rise and Fall of Americas Premier Mental Hospital*. New York: Public Affairs, 2001.

[Bechtel *et al.*, 1998] W Bechtel, A. Abrahamsen, and G. Graham. The life of cognitive science. In W. Bechtel and G. Graham, editors, *A Companion to Cognitive Science*, pages 2–104. Medford, MA.: Blackwell Publishers, 1998.

[Brooks, 1986] V. Brooks. *The Neural Basis of Motor Control*. New York: Oxford University Press, 1986.

[Brothers, 1997] L. Brothers. *Fridays Footprint: How Society Shapes the Human Mind*. New York: Oxford University Press, 1997.

[Brothers, 2001] L. Brothers. *Mistaken Identity: The Mind–Brain Problem Reconsidered*. Albany, New York: State University of New York Press, 2001.

[Cockerham, 2003] W. Cockerham. *Sociology of Mental Disorder, 4th edition*. New Jersey: Prentice–Hall, 2003.

[Costin, 1998] C. Costin. Your dieting daughter. In D. Sattler, V. Shabatay, and G. Kramer, editors, *Abnormal Psychology in Context: Voices and Perspectives*, pages 241–246. Boston: Houghton Mifflin, 1998.

[Doerr-Zegers, 2000] O. Doerr-Zegers. Existential and phenomenological approach to psychiatry. In M. Gelder, J. Lopez-Ibor, and N. Andreasen, editors, *New Oxford Textbook of Psychiatry*, volume 1, pages 357–362. New York: Oxford University Press, 2000.

[Dretske, 1981] F. Dretske. *Knowledge and the flow of information*. Cambridge, MA: MIT Press/Bradford Book, 1981.

[Feinberg, 1980] J. Feinberg. Absurd self-fulfillment: An essay on the merciful perversity of the gods. In P. van Inwagen, editor, *Time and Cause: Essays Presented to Richard Taylor*, pages 255–281. Dordrecht: Reidel, 1980.

[Fingarette, 1985] H. Fingarette. *Heavy Drinking: The Myth of Alcoholism as a Disease*. Berkeley: University of California Press, 1985.

[Flanagan, 2000] O. Flanagan. Destructive emotions. *Consciousness and Emotion*, 1:259–281, 2000.

[Fowler and Kuipers, 1995] P. Fowler, D. Garety and L. Kuipers. *Cognitive Behavior Therapy for Psychosis: Theory and Practice*. Chichester: Wiley, 1995.

[23]We received help in writing this paper, not just from individuals (too numerous mention here), who read parts and offered comments and criticisms, but from audiences at various conferences and universities where parts of the paper were delivered by the first named co-author. Special thanks are offered to audiences at Emory University, North Carolina State University, University of Dayton/University of Cincinnati, and Wake Forest University. Comments may be addressed to either or both co-authors.

[Fulford, 1989] K. Fulford. *Moral Theory and Medical Practice*. Cambridge: Cambridge University Press, 1989.

[Fulford, 1994] K. Fulford. Value, illness, and failure of action: Framework for a philosophical psychopathology of delusions. In G. Graham and G. L. Stephens, editors, *Philosophical Psychopathology*, pages 205–233. Cambridge, MA: MIT Press, 1994.

[Fulford and Graham, 2006 (forthcoming)] K. Fulford, T. Thornton and G. Graham. *Oxford Textbook of Philosophy and Psychiatry*. Oxford: Oxford University Press, 2006 (forthcoming).

[Garrett, 1989] R. Garrett. Loves way. In G. Graham and H. LaFollette, editors, *Person to Person*, pages 124–145. Philadelphia, PA: Temple University Press, 1989.

[Garrett, 1994] R. Garrett. The problem of despair. In G. Graham and G. L. Stephens, editors, *Philosophical Psychopathology*, pages 72–89. Cambridge, MA: MIT Press, 1994.

[Gert and Culver, 2004] B. Gert and C. Culver. Mental disorders. In J. Radden, editor, *The Philosophy of Psychiatry*, pages 415–425. Oxford: Oxford University Press, 2004.

[Gorenstein, 1992] E. Gorenstein. *The Science of Mental Illness*. San Diego: Academic Press, 1992.

[Graham and Horgan, 1998] G. Graham and T. Horgan. Southern fundamentalism and the end of philosophy. In M. DePaul and W. Ramsey, editors, *Rethinking Intuition: The Psychology of Intuition and Its Role in Philosophical Inquiry*, pages 271–292. Lanham, Maryland: Roman & Littlefield, 1998.

[Graham *et al.*, 2005] G. Graham, T. Horgan, and J. Tienson. Consciousness and intentionality. In S. Schnieder and M. Velmans, editors, *Companion to Consciousness (forthcoming)*. Malden, MA: Blackwell, 2005.

[Graham, 2002] G. Graham. Recent work in philosophical psychopathology. *American Philosophical Quarterly*, 39:109–134, 2002.

[Grunbaum, 1984] A. Grunbaum. *The Foundations of Psychoanalysis*. Berkeley, CA: University of California Press, 1984.

[Harrison, 1998] B. Harrison. I am not afraid. In D. Sattler, V. Shabatay, and G. Kramer, editors, *Abnormal Psychology in Context: Voices and Perspectives*, pages 3–7. Boston: Houghton Mifflin, 1998.

[Heineman, 1998] M. Heineman. Losing your shirt. In D. Sattler, V. Shabatay, and G. Kramer, editors, *Abnormal Psychology in Context: Voices and Perspectives*, pages 160–164. Boston: Houghton Mifflin, 1998.

[Hempel, 1965] C. Hempel. Fundamentals of taxonomy. In *Aspects of Scientific Explanation and Other Essays in the Philosophy of Science*, pages 137–154. New York: Free Press, 1965.

[Heninger, 1999] G. Heninger. Special challenges in the investigation of the neurobiology of mental illness. In D. Charney, E. Nestler, and B. Bunney, editors, *Neurobiology of Mental Illness*, pages 89–99. New York: Oxford University Press, 1999.

[Hobson and Leonard, 2001] J. A. Hobson and J. Leonard. *Out of Its Mind: Psychiatry in Crisis: A Call for Reform*. Cambridge, MA.: Perseus, 2001.

[Horgan and Graham, 2005] J. Horgan, T. Tienson and G. Graham. Phenomenal intentionality and the brain in the vat. In R. Schantz, editor, *The Externalist Challenge*, pages 297–317. Berlin: De Gruyter, 2005.

[Horgan and Tienson, 2002] T. Horgan and J. Tienson. The intentionality of phenomenology and the phenomenology of intentionality. In D. Chalmers, editor, *Philosophy of Mind: Classical and Contemporary Readings*, pages 520–533. New York: Oxford University Press, 2002.

[Horwitz, 2002] A. Horwitz. *Creating Mental Illness*. Chicago: University of Chicago Press, 2002.

[Jaspers, 1997] K. Jaspers. *General Psychopathology*. Baltimore and London, Johns Hopkins, 1997. trans. Hoenig, J. and Hamilton, W.

[Kingdom and Turkington, 1994] D. Kingdom and D. Turkington. *Cognitive Behavior Therapy of Scizophrenia*. Hove: Earlbaum, 1994.

[Kitcher, 1992] P. Kitcher. *Freuds dream: a complete interdisciplinary science of mind*. Cambridge, MA: MIT Press, 1992.

[Kleinman and Good, 1985] A. Kleinman and B. Good. *Culture and Depression: Studies in the Anthropology and Cross-Cultural Psychiatry of Affect and Disorder*. Berkeley: University of California Press, 1985.

[Laurence and Margolis, 1999] S. Laurence and E. Margolis. Concepts and cognitive science. In E. Margolis and S. Laurence, editors, *Concepts: Core Readings*, pages 3–81. Cambridge, MA: MIT Press, 1999.

[Lovestone, 2000] A. Lovestone. Dementia: Alzheimers disease. In M. Gelder, J. Lopez-Ibor, and N. Andreasen, editors, *New Oxford Textbook of Psychiatry*, volume 1, pages 387–397. New York: Oxford University Press, 2000.

[Marinoff, 1999] L. Marinoff. *Plato, Not Prozac: Applying Philosophy to Everyday Problems*. New York: Harper Collins, 1999.

[Mechanic, 1999] D. Mechanic. Mental health and mental illness: definitions and perspectives'. In A. Horwitz and T. Schied, editors, *A handbook for the study of mental health: social contexts, theories, and systems*, pages 12–28. Cambridge: Cambridge University Press, 1999.

[Montague and Dayan, 1998] P. R. Montague and P. Dayan. Neurobiological modeling. In . W. Bechtel and G. Graham, editors, *A Companion to Cognitive Science*, pages 526–541. Malden, Massachusetts: Blackwell Publishers, 1998.

[Munro, 2000] A. Munro. Persistent delusional symptoms and disorders. In M. Gelder, J. Lopez-Ibor, and A. Andreasen, editors, *The New Oxford Textbook of Psychiatry*, volume 1, pages 651–676. New York: Oxford University Press, 2000.

[Murphy and Stich, 2000] D. Murphy and S. Stich. Darwin in the madhouse: Evolutionary psychology and the classification of mental disorders. In P. Caruthers and A. Chamberlin, editors, *Evolution and Cognition*, pages 62–92. Cambridge: Cambridge University Press, 2000.

[Murphy and Woolfolk, 2001] D. Murphy and R. Woolfolk. The harmful dysfunction analysis of mental disorder. *Philosophy, Psychiatry, and Psychology*, 7:241–252, 2001.

[Poland *et al.*, 1994] J. Poland, B. Von Eckardt, and W. Spaulding. Problems with the dsm approach to classifying psychopathology. In G. Graham and G. L. Stephens, editors, *Philosophical Psychopathology. Cambridge*, pages 235–260. MA.: MIT Press, 1994.

[Raabe, 2001] P. Raabe. *Philosophical Counseling: Theory and Practice*. Westport, Conn.: Praeger, 2001.

[Radden, 2000] J. Radden. From melancholic states to clinical depression. In J. Radden, editor, *The Nature of Melancholy*, pages 3–51. Oxford: Oxford University Press, 2000.

[Ramsay, 1998] W. Ramsay. Prototypes and conceptual analysis. In M. DePaul and W. Ramsay, editors, *Rethinking intuition: the psychology of intuition and its role in philosophical inquiry*, pages 161–177. Lanham, MD.: Rowman and Littlefield, 1998.

[Rappaport, 1998] J. Rappaport. The accident that didnt happen. In D. Sattler, V. Shabatay, and G. Kramer, editors, *Abnormal psychology in context: voices and perspectives*, pages 14–19. Boston: Houghton Mifflin, 1998.

[Russell, 2000] B. Russell. A free mans worship. In E. Klemke, editor, *The Meaning of Life*, pages 71–77. New York: Oxford University Press, 2000.

[Schopenhauer, 2000] A. Schopenhauer. On the vanity of existence. In E. Klemke, editor, *The Meaning of Life*, pages 67–70. New York: Oxford University Press, 2000.

[Schumaker, 2001] J. Schumaker. *The age of insanity: modernity and mental health*. Westport, Connecticut: Prager, 2001.

[Slavney and McHugh, 1987] P. Slavney and P. McHugh. *Psychiatric Polarities: Methodology and Practice*. Baltimore: Johns Hopkins University Press, 1987.

[Stephens and Graham, 2000] G. L. Stephens and G. Graham. *When self-consciousness breaks: alien voices and inserted thoughts*. Cambridge, MA: MIT Press, 2000.

[Stephens and Graham, 2004] G. L. Stephens and G. Graham. Reconceiving delusion. *International Review of Psychiatry*, 16:236–41, 2004.

[Stephens and Graham, 2006] G. L. Stephens and G. Graham. The delusional stance. In K. Fulford M. Chung and G. Graham, editors, *Reconceiving Schizophrenia (forthcoming)*. Oxford: Oxford University Press, 2006.

[Szasz, 1960] T. Szasz. The myth of mental illness. *American Psychologist*, 15:113–118, 1960.

[Szasz, 1961] T. Szasz. *The Myth of Mental Illness: Foundations of a Theory of Personal Conduct, revised edition*. New York: Harper and Row, 1961.

[Taylor, 1970] R. Taylor. *Good and Evil*. London: Methuen, 1970.

[Taylor, 1999] M. Taylor. *The Fundamentals of Clinical Neuropsychiatry*. New York: Oxford University Press, 1999.

[Wakefield, 1992] J. Wakefield. The concept of mental disorder: On the boundary between biological and social values. *American Psychologist*, 47:373–388, 1992.

[Wakefield, 1999] J Wakefield. The measurement of mental disorder. In A. Horwitz and T. Schied, editors, *A handbook for the study of mental health: social contexts, theories, and systems*, pages 29–59. Cambridge: Cambridge University Press, 1999.

[Wheaton, 1999] B. Wheaton. The nature of stressors. In A. Horwitz and T. Schied, editors, *A handbook for the study of mental health: social contexts, theories, and systems*, pages 176–197. Cambridge: Cambridge University Press, 1999.

[Woolfolk, 1999] R. Woolfolk. Malfunction and mental illness. *Monist*, 82:658–670, 1999.

[Zuckerman, 1999] M. Zuckerman. *Vulnerability to Psychopathology: A Biosocial Model*. Washington, D. C.: American Psychological Association, 1999.

THE ADAPTIVE PROGRAMME OF EVOLUTIONARY PSYCHOLOGY

Robert C. Richardson

1 THE PROGRAM OF EVOLUTIONARY PSYCHOLOGY

Human beings, like other organisms, are the products of evolution. Our psychological capacities are evolved traits as much as is our gait and dentition. Human beings, like other organisms, exhibit traits that are the products of natural selection. Our psychological capacities are subject to natural selection as much, again, as is our gait or dentition. In this minimal sense, there can be no reasonable quarrel with evolutionary psychology. It is incontrovertible that evolution is real. It is neither worth defending nor disputing the fact that we are the products of evolution any more than it is worth disputing that our bodies are composed of cells or that the Earth circles about the Sun. It is also clear that natural selection shaped life on Earth and shaped human evolution. It is neither worth defending nor disputing the fact that humans are the products of natural selection any more than it is worth disputing that we are mortal or that we are bound to the Earth by gravity.

Evolutionary psychology claims this much, but more. In part, evolutionary psychology is a psychological program, geared to the reform of psychology. Evolutionary psychology also embraces an aggressive biological program: psychological capacities are adaptations, not to present circumstance, but to our ancestral environment. The fact that they are adaptations, in turn, explains those psychological capacities. Since evolution explains our psychological capacities, it also should offer some insight into the conduct of psychology. Leda Cosmides and John Tooby, two of the most prominent figures in the field, offer this description as the agenda, placing it in an evolutionary perspective:

> The human mind is the most complex natural phenomenon humans have yet encountered, and Darwin's gift to those who wish to understand it is a knowledge of the process that created it and gave it its distinctive organization: evolution. Because we know that the human mind is the product of the evolutionary process, we know something vitally illuminating: that, aside from those properties acquired by chance, the mind consists of a set of adaptations, designed to solve the long-standing adaptive problems humans encountered as hunter-gatherers [Cosmides and Tooby, 1992, 163].

Handbook of the Philosophy of Science. Philosophy of Psychology and Cognitive Science
Volume editor: Paul Thagard
General editors: Dov M. Gabbay, Paul Thagard and John Woods
© 2007 Elsevier B.V. All rights reserved

The program is committed not just to evolution but to adaptation. The mind is a complex phenomenon. Features with complex functional designs must have evolved, and in order to evolve they must have provided a substantial advantage to our ancestors in virtue of their design. It is not just that the mind evolved. It is not just that humans evolved, or that our evolution is subject to natural selection. There must have been selection *for* the complex functional and cognitive designs we observe. Adaptation alone can explain the capacities we exhibit, including the complexities of judgment, language, thought, perception, emotion, and action. Adaptation alone can explain who and what we are. 'Darwin's gift' is the insight that adaptation is sufficient to explain who and what we are. The advantages offered by our cognitive organization must explain its presence and the details of its structure. So, too, our perceptual abilities, and our emotions, reflect our evolutionary history. The mind, from the vantage point of evolutionary psychology, is no different than any other complex feature.

Evolutionary psychology is ambitious in another way. It offers a comprehensive agenda, applying to a broad range of characteristically human behaviors. We are aggressive in defending family and territory; indeed, humans are remarkable for their ferocity, and for our willingness to engage in gratuitous violence even to the point of genocide. Among other things, we are likely to be responsible for the extinction of the striking megafauna characteristic of the Americas toward the end of the last Ice Age, in what is called 'Pleistocene overkill'; and we certainly deserve credit for the demise of less dramatic forms such as the Dodo and the American passenger pigeon. And of course, we happily exterminate our own kind. We have complex sexual relations. There are differences between men and women in terms of what we value and how we behave. We are afraid of strangers, and of heights. There are conflicts between children and parents, and differences between parents. On the more positive side, we engage in complex play. We engage in a variety of cooperative behaviors, and have lasting friendships. Much of the work in evolutionary psychology is devoted to documenting such complex patterns and explaining them.

Assuming the patterns are real, evolutionary psychology offers explanations for such human tendencies in terms of the evolved cognitive and emotional structures. Some human violence, at least, is evidently geared to the acquisition of resources. It is tied to sexual rivalry, to power, to wealth, and to status. In his famous studies of the Yanamamö, Napoleon Chagnon [1968; 1988] illustrates the tendencies in great detail. Men will raid for food, or for women, and they engage in ritualistic fights for status. Mate preferences accordingly are taken to be evolved responses, rather than socially derived. It is a kind of social accident if a Yanamamö male uses an axe on a competitor, since the axe was introduced, but the aggression is natural. The tendency toward aggression is to be explained in evolutionary terms, even if the specific form it takes is socially shaped. In like manner, David Buss [1994; 1999; 2000] traces the differences between the sexes finally to differences in investment between the sexes: in terms of mate preference, Buss tells us females favor mates who are dependable, stable, and high status, while men prefer mates with youth

and health. How status is measured may be socially variable and contingent, but a preoccupation with status is not because it has evolutionary consequences. At least some human fears, but not all, get explanations in evolutionary terms. So a fear of snakes, like our fear of strangers or of heights, supposedly serves to protect us from dangers. Having observed that snakes and spiders are *always* scary, and not only to humans but to other primates, Steven Pinker says 'The common thread is obvious. These are the situations that put our evolutionary ancestors in danger. Spiders and snakes are often venomous, especially in Africa ... Fear is the emotion that motivated our ancestors to cope with the dangers they were likely to face' [Pinker, 1997, 386]; cf. [Nesse, 1990].

1.1 Adaptation and Evolutionary Psychology

One factor these explanations share is explaining the patterns *as adaptations*; that is, the patterns are explained as the products of natural selection acting over generations, molding human behavior to the demands of survival and reproduction. Cosmides and Tooby say this:

Natural selection shapes domain-specific mechanisms so that their structure meshes with the evolutionary stable features of their particular problem domains. Understanding the evolutionary stable feature of problem domains — and what selection favored as a solution under ancestral conditions — illuminates the design of cognitive specializations [1994, 530].

How do we know they are adaptations? Sometimes this is simply assumed. In other cases it is a more serious issue. The most straightforward general argument for a focus on adaptation starts with the complexity of the features. The point is not merely that complex features are evolved. All our features, including the least complex, have evolved. Even the simplest and most basic may be inherited. Even maladaptive traits may be inherited. However simple, and however adaptive, they all have evolutionary explanations. The early development of human beings, for example, has a characteristic radial cleavage that we share with other vertebrates, one that differs from the pattern we see in arthropods. We share this developmental pattern with vertebrates because it was inherited from those ancestors. It evolved as part of the developmental pattern, though alternatives appear to be equally well adapted. Adaptationism is more ambitious than this. This is where complexity matters. One of Darwin's key insights was that it was possible to explain adaptation without Intelligent Design. He also saw that complex features present a special problem for evolution. Darwin wrote in the Introduction to *On the Origin of Species*:

> In considering the Origin of Species, it is quite conceivable that a naturalist, reflecting on the mutual affinities of organic beings, on their embryological relations, their geographical distribution, geological succession, and other such facts, might come to the conclusion that each species had not been independently created, but had descended, like varieties, from other species. Nevertheless, such a conclusion, even if

> well founded, would be unsatisfactory, until it could be shown how the innumerable species inhabiting this world would have been modified, so as to acquire that perfection of structure and coadaptation which most justly excites our admiration (1859, p. 3).

Evolutionary psychologists follow Darwin. The presence of complex features, those exhibiting 'perfection of structure and adaptation,' we are told, *must* be the consequence of evolution by natural selection. Pinker embraces the line, saying 'Natural selection has a special place in science because it alone explains what makes life special. Life fascinates us because of it *adaptive complexity* or *complex design* [Pinker, 1997, 155]; cf. [Tooby and Cosmides, 1992, 49 ff.; Pinker and Bloom, 1992; Grantham and Nichols, 1999]. In *The Extended Phenotype*, Dawkins says 'if we see an animal with a complex organ, or a complex and time-consuming behavior pattern, we would seem to be on strong grounds in guessing that it must have been put together by natural selection. ... The working hypothesis that they must have a Darwinian survival value is overwhelmingly strong' [1982, 43]. The idea is intended by Dawkins — no less than by Tooby, Cosmides, and Pinker — to be a straightforwardly Darwinian one: natural selection is the most reasonable, or perhaps most reasonable, explanation for the evolution or presence of complex functional designs. Other evolutionary factors such as mutation or drift, by contrast, may be of evolutionary consequence but will not tend to lead systematically to such complex features, and if there are constraints on the evolutionary process, those will tend to limit rather than facilitate adaptation. In this Darwinian sense, these alternative evolutionary mechanisms can be regarded as 'chance' factors, uncorrelated with evolutionary advantage. Tooby and Cosmides say 'It would be a coincidence of miraculous degree if a series of these function-blind events, brought about by drift, by-products, hitchhiking, and so on, just happened to throw together a structure as complexly and interdependently functional as an eye' [1992, 57]; cf. [1992]. In an even more dramatic fashion, Buss acknowledges that though some features evolve as 'by-products' or 'spandrels,' and though some features are prevalent by chance, it is natural selection that occupies center stage for evolutionary biology:

> Despite scientific quibbles about the relative size of the three categories of evolutionary products, all evolutionary scientists agree on one fundamental point: adaptations are the primary product of evolution by natural selection ... Those characteristics that pass through the selective sieve generation after generation for hundreds, thousands, and even millions of years, are those that helped to solve the problems of survival and reproduction [1999, 39].

He concludes that 'the core of all animal natures, including humans, consists of a large collection of adaptations' (loc. cit.). Evolutionary psychologists need not deny that there are other evolutionary factors at work. They need not deny the workings of 'chance.' They do characteristically focus on natural selection. This

is because they are convinced that in doing psychology we are faced with features so complex that they demand explanation as adaptations.

The central programmatic goal of evolutionary psychology, correspondingly, is to provide evolutionary explanations of our natural psychological capacities in terms of natural selection. To take the most prominent psychological example, Cosmides and Tooby claim that human reasoning consists of a set of mechanisms organized around social exchange. They describe the program this way:

> According to the evolutionary psychological approach to social cognition ... the mind should contain organized systems of inference that are specialized for solving various families of problems, such as social exchange, threat, coalitional relations and mate choice. ... Each cognitive specialization is expected to contain design features targeted to mesh with the recurrent structure of its characteristic problem type, as encountered under Pleistocene conditions. Consequently, one expects cognitive adaptations specialized for reasoning about social exchange to have some design features that are particular and appropriate for social exchange, but that are not activated by or applied to other content domains [1992, 166].

On this view, the evolutionary function of human reasoning involves facilitating and monitoring social exchange and social relations. Human reasoning then would be a 'cognitive adaptation' to social conditions encountered by our evolutionary ancestors, established and maintained by natural selection. In this case, the social context of hominid life shapes the behavior. The goal of evolutionary psychology, thus, is explaining psychological processes as biological adaptations to 'Pleistocene' conditions, and hominid social structures, shaped and maintained by natural selection. It does not follow, of course, that human reasoning is adapted to our current conditions — with increased crowding, larger social groups, or overwhelming amounts of information. Human reasoning may now fail to be adaptive, but how we reason is supposed to be explained by our history. Once that history is exposed, we may even find that we can redesign our environment to better suit our natural ways of thinking (cf. [Gigerenzer, 1998]). According to evolutionary psychologists, this is true not only for social cognition but for a wide array of psychological mechanisms and social behaviors. Similar explanations, as I have said, have been offered by evolutionary psychologists for family structure, parental care, marital jealousy, sex roles, sexual preferences, familial affection, personality, and the moral sentiments, to mention just a few.

1.2 Adaptation and Psychological Design

The evolutionary explanations offered are a means to an end. Most advocates of evolutionary psychology are themselves psychologists or anthropologists rather than biologists. The problem psychologists saw in their own field was a kind of intellectual malaise following on the lack of a definite vision. Evolutionary biology

offers a vision which, they claim, can transform *psychological* research (cf. [Davies, 1996; 1999]). The synthesis of psychology and biology that evolutionary psychology offers is meant in the end to reform psychology. The contributions on offer come at a number of distinctive levels. Tooby and Cosmides are fundamentally interested in the potential of evolutionary biology to grounding work in psychology and the social sciences. 'Modern biology,' they say, 'constitutes, in effect, an 'organism design theory" [1992, 53]. It can be used as a lever both to undercut the 'standard' models in social science, and to 'guide the construction' of an improved psychology. Donald Symons [1992] similarly tells us that Chagnon's [1988] work on the use of terms for kinship among the Yanamamö was inspired by a selectionist vision. No doubt, this is so. Kinship is here a social matter. The particular category to which an individual is assigned affects their eligibility for marriage. That in turn affects the fitness of individuals. Chagnon noticed that people do not simply accept the existing classification if they can benefit by changing it. So, for example, he describes one Yanamamö man who called a young woman by a name that enhanced her sexual eligibility for his son. Symons takes this to be a defiance of social imperative. Chagnon's conclusions are more modest than Symons', and he did not here offer anything like an 'organism design theory.' Symons' point is nonetheless correct: the goal is a contribution from biological to social theory, in this case to anthropology.

There are other obvious contributions evolution could make to the understanding of human psychology. If biology offers an 'organism design theory,' evolutionary psychology could, with little overstatement, be said to offer a science of 'human nature.' It begins with a description of our adaptation to ancestral environments, and moves to explaining our capacities, including our psychological capacities. To take one example, many evolutionary psychologists are prepared to explain how and why the modularity of mind is a natural consequence of evolution (cf. [Fodor, 1985; Baron-Cohen, 1995]). The thought is that that the mind is an amalgam of special-purpose mechanisms, rather than a general purpose machine. So as phrenologists portrayed the mind as a mosaic of faculties two centuries ago, we still find the mind portrayed as a mosaic of relatively independent modules, though of course the emphasis on the shape of the skull has been replaced with PET and fMRI, allowing us to glimpse the functioning of the brain in action. That a modular organization would have facilitated the evolution of mind is often taken as obvious. Symons [1992] dismisses any alternative out of hand with the observation that there is no such thing as a general problem solver, since there is no such thing as a general problem. Cosmides and Tooby [1992] are equally insistent, but not so dogmatic. Here is one passage:

> A basic engineering principle is that the same machine is rarely capable of solving two different problems equally well. ... Our body is divided into organs such as the heart and the liver, for exactly this principle. Pumping blood throughout the body and detoxifying poisons are two very different problems. Consequently, the body has a different machine for solving each problem [1992, 80].

Functional specialization, they assume, is good engineering design. Similarly, if the mind is a machine selected for its ability to cope with environmental demands, then it should also 'consist of a large number of circuits that are *functionally specialized*' (loc. cit.). One key thought evolutionary psychologists want to enforce is that, if we assume an evolutionary perspective, then the mind should consist of a set of specialized systems, each 'designed' to solve some relatively specific family of problems, such as problems of social exchange, mate choice, threat, and the like. The alternative is to assume that our reasoning consists in the application of relatively general procedures which are not content specific. This is integral to what they call the 'standard social science model.' General procedures are unlikely, they claim, to yield adequate solutions to the problems needing to be solved. There are several distinct questions one might want to answer. First, *is* our cognitive organization modular? Second, if it is modular, what sort of content specialization should we expect? Third, what are the implications of adaptive thinking for such questions? Even at the first level, we should be cautious. The idea that there is modular organization in the brain, with specialized and relatively independent subsystems, is controversial (cf. [Uttal, 2001; Buller, 2005, ch. 4]). This is an issue that would take us too far afield (but see [Bechtel and Richardson, 1993]). Evolutionary psychology likely has more implications for the second question than the first. If we know the evolutionary 'problems' our ancestors faced, and we assume a modular organization, then we arguably should be able to project the kinds of specialized 'problem solving' abilities humans exhibit. Our ancestors surely did not need the capacity to play chess, or solve differential equations, though no doubt the ability to coordinate among themselves was desirable. They did not need the capacity to navigate at the speed of sound, though no doubt a sensitivity to flying objects was desirable. The third question is the one that will occupy us. This is where the methodological implications of evolutionary psychology should be most apparent. Nonetheless, the questions are not independent. The implications depend, like the Devil, on the details.

1.3 Psychology and Social Science

These details depend, among other things, on the connection between psychological data and the interpretation of it by evolutionary psychologists. The first stage aims to tie the psychological data to an evolutionary interpretation relatively directly. So we are offered a contrast between two pictures. One is what Tooby and Cosmides call the 'Standard Social Science Model,' according to which humans come equipped with only general purpose rules for learning. The other is the 'Evolutionary Psychological Model,' according to which, in addition to whatever general purpose rules there might be, there are also special purpose, content specific, rules. Cosmides and Tooby suggest, specifically, that there are specialized cognitive procedures — what they call Darwinian algorithms — for social exchange. Here is one way they tie the psychological data to its biological interpretation. If there were only general purpose learning rules, they reason, then

all the specific content would be attributable to 'cultural' influences, or to local differences. The rules, if general, could provide no specific content. The social environments into which we are born would provide all the needed 'content' for social judgment. If, on the other hand, there are specialized rules for social exchange, these can form 'the building blocks out of which cultures themselves are manufactured,' and these rules would constitute a kind of 'architecture of the human mind' undergirding a social exchange psychology. Every human being would come with the same basic cognitive equipment, and the traditions, rituals and institutions characteristic of human life at most would supply the 'specifics' [1992, 208]. We might be geared cognitively, for example, to discern cheaters, though the specific opportunities for cheating vary substantially.

Cosmides and Tooby draw two conclusions, both favoring evolutionary psychology. First of all, if there is a 'universal evolved architecture of the human mind,' then there should be some features of social exchange which are common across individuals and across cultures. They tell us that this is exactly what we find. The implication is, supposedly, that we should *not* find this if we assume only generalized learning mechanisms. This latter implication is plainly not true. Even with only generalized learning rules, we would expect *some* features to be common across cultures, and across individuals. An inborn content specific mechanism is not needed to recognize that humans all have navels; neither is an inborn content specific mechanism for social reasoning needed to explain why humans typically trade foods, or seek what is necessary for warmth. An inborn content also would not be needed to recognize a need for communication. There would undoubtedly be *some* common features even with only generalized learning rules. Presumably this is not what Cosmides and Tooby, or other evolutionary psychologists, intend to deny. It must be *some specific features* common across individuals or across cultures that tell in favor of specialized rules for social exchange. The important issues must revolve around the specific form of the features we find and whether these are sufficiently general.

Cosmides and Tooby offer a second argument for modularity and specialization. It is worth quoting at some length because it is echoed in the work of other evolutionary psychologists:

> ... the Standard Model would have to predict that wherever social exchange is found to exist, it would have to be taught or communicated from the ground up. Because nothing about social exchange is initially present in the psychology of the learner, every structural feature of social exchange must be specified by the social environment as against the infinity of logically alternative branchings that could exist. It is telling that it is just this explicitness that is usually lacking in social life. Few individuals are able to articulate the assumptions that structure their own cultural forms [1992, 208].

Cosmides and Tooby suggest that it is shared assumptions that make the communication of specific cultural information possible, just as it is the shared as-

sumptions that make the learning of human languages possible. These shared assumptions are what constitute the content behind social exchange. They are the background against which social exchange makes sense, or is possible. Without these assumptions, there would be no way to learn about our social life. The general line of reasoning should be relatively familiar. It is a line of reasoning with an impressive pedigree, including Immanuel Kant and Noam Chomsky. The problem with this line of thought is reasonably clear. The fact that we have at our disposal cognitive mechanisms which facilitate reasoning about social exchange is something everyone must concede. Were that not so, we certainly would not have international banking systems or international trade. We would not even have markets in which we can buy food or household goods. Even bartering would be impossible. We can assume the psychological evidence settles that issue, though we need not turn to the *Psychological Bulletin* for definitive evidence. That we engage in such exchange is not sufficient to show that there must be inherent, shared assumptions, concerning social exchange which are responsible for the acquisition of those rules. It is possible, so far as that evidence goes, that early learning establishes specialized social rules which we deploy in reasoning about social contracts. Nothing offered bears directly on how the cognitive mechanisms are acquired. Consider an analogy. We almost certainly have at our disposal cognitive mechanisms which facilitate reasoning about mathematics. It may be true that there are inherent cognitive schemata which make it possible for us to learn mathematics. Kant certainly thought this was so and thought it something that could be established *a priori*, just as Descartes thought geometrical ideas are innate. Nonetheless, it is not simply the fact that we can reason about mathematics that, in Kant's view, establishes there are such schemata. Kant attempted to ground the necessity for these prior schemata in peculiar features of mathematical knowledge. Descartes focused on the idea of the infinite, and the fact that this cannot be constructed from any series. Cosmides and Tooby move seamlessly from the idea that we have some specialized cognitive mechanisms for social reasoning to the conclusion that the acquisition of these cognitive mechanisms depends on the prior existence of foundational social schemata. Where Kant thought an argument was needed, Cosmides and Tooby think none is necessary. I do not doubt Cosmides and Tooby's conclusion could be true. We certainly *do* engage in social exchange. Perhaps there are specialized and innate cognitive mechanims responsible for this. That does not change the fact that it is not rooted in any persuasive argument. The simpler observation give us no reason to believe it is true, or even that it is more likely than not.

1.4 Evolutionary Psychology and Sociobiology

Evolutionary psychology recapitulates some of the issues surrounding sociobiology following the publication of E. O. Wilson's seminal *Sociobiology: The New Synthesis* [1975]. Richard Dawkins supposedly characterized evolutionary psychology as 'rebranded sociobiology' (cf. [Rose, 2000]). For those who have read Robert

Ardrey's *The Territorial Imperative* [1966], Konrad Lorenz's *On Agression* [1966] or Desmond Morris' *The Naked Ape* [1967], the Hobbesian vision they share seems clear enough. There is certainly some continuity of vision, but there is not quite an identity of vision (cf. [Downes, 2001]). So, for example, Charles Lumsden and E.O. Wilson [1981] wrote that 'The central tenet of human sociobiology is that social behaviors are shaped by natural selection' [1981, 99]. This tenet of 'human sociobiology' is certainly shared with evolutionary psychology, as it is with some of the most severe biological critics. It was also embraced by Herbert Spencer, with the assent of Charles Darwin. Still, even though even some of the players are the same, and even though there are views common between them, there are significant differences.

To begin with, evolutionary psychologists appeal fundamentally to computational mechanisms rather than social structures to explain current behavior. Typically, they draw more from cognitive psychology, and from computational psychology, than from ethology, both in terms of method and theory (cf. [Griffiths and Sterelny, 1999]). Methodologically, evolutionary psychologists approach their psychology with cognitive techniques. Cosmides and Tooby, for example, focus on performance in reasoning tasks; and what is taken to be informative is our ability to solve reasoning problems. Gerd Gigerenzer similarly focuses on probabilistic reasoning; and it is performance on these problems that is informative. Theoretically, the psychological mechanisms turn out to be the key explanatory factors, or at least the key proximate factors. The thought seems to be that evolutionary causes will reveal the proximate causes. Sociobiology focuses on the evolutionary causes, with little attention to the detailed mechanisms.

More importantly, sociobiology is an enterprise much broader in ambition than evolutionary psychology. Many of the most striking successes of sociobiology focused on evolutionary models for animal social behavior — that is, for non-human animal behavior. There is, naturally, no reason biologists should neglect the explanation of social behavior. And they do not. In the decades following Wilson's *Synthesis*, comparative biology has flourished, and benefited from being placed in a more rigorous evolutionary context. Studies of animal behavior have likewise flourished and benefited from being placed in a more rigorous evolutionary context. The range of this work is considerable, from the social behavior of spiders to the mating patterns of hyenas, and the capacities of our primate cousins. Evolutionary psychology has been largely unconcerned with animal behavior, or with comparative issues. They could benefit from the larger perspective.

Human sociobiology, on the other hand, bears substantial resemblance to evolutionary psychology. Evolutionary psychology is thus geared especially to the human case. It is fundamentally a piece of human sociobiology, updated to recognize the importance of cognitive mechanisms to behavior. Philip Kitcher [1985] derisively called this 'pop sociobiology.' Stephen Jay Gould called it 'pop ethology.' Incest was a centerpiece for much of human sociobiology. Here is the way the reasoning went at the time (see [Ruse, 1979; van den Berghe, 1980; 1983]). Nearly all cultures treat incest as taboo. Incest taboos are a 'cultural universal.'

It is important to be careful in assessing what this means. Kinship relations are crucial for all societies, defining the core structures in nearly all social negotiation; more specifically, kin are limited in their sexual relations. These prohibitions are incest taboos. We can, apparently, explain the prevalence of those taboos in terms of evolutionary advantage. Inbreeding has a variety of adverse effects, many of which can be understood genetically. When close relatives interbreed (siblings or first cousins, say) there is a (proportionally) greater likelihood that recessive and deleterious genes will be expressed that would be masked in unions among more distant relatives. In fact, that is something like what we see. Mortality rates are markedly higher among children that result from inbreeding, as are debilitating illnesses. So we find what we should expect. Given the evolutionary advantages taboos on such unions would offer, the reasoning goes, there must be some evolved mechanism that decreases sexual attraction under circumstances which would be conducive to inbreeding. Since the human family — including both the 'nuclear family' and the extended family group — involve close contact, it appears likely that the close contact is the 'trigger' reducing attractiveness. Again, there is some evidence, though largely anecdotal, in support of the result.

The appearance of an easy connection between incest and inbreeding is illusory, even though the genetics is straightforward. Questions concerning the 'data' are many (cf. [Kitcher, 1985, 169 ff.]). Estimates of incest levels vary a great deal, as do the very definitions of what counts as incest. Generally, getting reliable estimates of sexual behavior is a difficult matter, and given the severe social sanctions it would not be surprising if we systematically underestimate the frequency of incest. It follows immediately that we overestimate the effectiveness of the taboos. With serious social sanctions, admitting to incest is something we would not be inclined to do. Sometimes the social prohibitions preclude a wide variety of intimate contact. Others limit the scope to heterosexual intercourse. Incest taboos, though, are not limited to those genetically related. Relations among step siblings are within the scope, typically, as much as relations between full siblings. Here's Kitcher's thoughtful conclusion:

> How strong is the human propensity to avoid copulating with those known intimately from childhood? How strong is it in contemporary members of our species? How strong has it been in the past? We simply do not know. ... Any serious study of incest should recognize the extent of our ignorance rather than rushing to pronounce on the behavioral rule that people must be following [1985, 274-4].

There are also questions concerning how we could use the data, even if we were convinced of its robustness and reliability. For example, what is known about the relatedness of offspring when there are not deleterious effects? The simplest question is the most obvious: why are there taboos for something that we 'naturally' avoid?[1] Sociobiology made a good deal out of the pattern we find

[1]The point was made by [Lewontin, Rose and Kamin, 1984]. That it is obvious does not make

among contemporary cultures, and of the broader pattern shared among animals. That pattern is complicated to say the least. Kitcher is right to emphasize the extent of our ignorance. The most striking problem derives from the clear fact that social kinship does not generally reflect biological kinship. In matrilineal societies, the biological father belongs to a different social unit than his biological offspring; so though there are limits on sexual and marital relations among clan members, a male may not be prohibited from these relations with his biological offspring — who are, after all, not his 'daughters' (cf. [Sahlins, 1976; Benton, 2000]). The evolutionary explanation of the taboos, however, is supposed to be grounded in the genetic problems resulting from inbreeding. Socially defined relationships are not biological. The categories are mismatched.

Evolutionary psychology sidetracks many questions of this sort, by attempting to focus on what was present in our ancestral groups, rather than the confusing mosaic we find in contemporary cultures. This introduces some complications to the questions, but they look like reasonable complications. If we want to explain some pattern of human behavior, or some human social arrangements, then we need to step back to see what that pattern is. If we want to know whether selection would favor incest taboos, what is relevant is not the current array of varied social arrangements, but the consequences of incest in the past. If incest taboos extend to adopted children who are not biological relatives at all, that may be because contemporary families do not reflect ancestral conditions. *Assuming* ancestral groups were essentially extended families, *then* sexual relations within those groups *would* be expected to have adverse biological consequences. Of course, all this is done in spite of the fact that we do not know the social organization of ancestral groups, and in spite of the fact that extant 'primitive' groups do not respect the condition. In the case of the !Kung, apparently, bands are also not organized around biological patterns. Marriages between close biological relatives is not at all uncommon. Incest taboos remain in force. Kitcher's observation that we are largely ignorant of the conditions stands firm. Whether these taboos now maximize fitness under current social arrangements is not relevant. What matters is whether they maximized fitness under those ancestral conditions. We do not know the answer to this, and little evidence bears directly on the question.

Evolutionary psychology is, at least superficially, more modest in its pretensions than is Sociobiology concerning what biology implies for social theory, if not for psychology. To be sure, we are the products of our evolutionary history. We are also the products of our social upbringing. Cosmides and Tooby are content to suggest that psychological theories provide the 'foundations' for 'theories of culture,' and at least in more moderate moments do not claim overtly that they 'constitute' alternative theories of culture. In this they appear to be less ambitious

it less important. Lumsden and Wilson [1981] attempt to answer the point, but unsuccessfully (cf. [Kitcher, 1985, 347]). Humans are generally averse to consuming putrid foods. We are not averse, or at least many are not averse, to eating pork. The former may be due to biological features, though it is important to ask, for example, why dogs do not share our aversions even though they are not all that distant biologically. The aversion to pork needs a decidedly different form of explanation. It alone is a subject for taboos.

than their sociobiological forebears, who would have dispensed with both psychology and social science in favor of Biology. With the caveat in place, Tooby and Cosmides continue:

> Nevertheless, increasing knowledge about our evolved psychological architecture places increasing constraints on admissible theories of culture. Although our knowledge is still very rudimentary, it is already clear that future theories of culture will differ significantly in a series of ways from Standard Social Science Model theories. Most fundamentally, if each human embodies an evolved psychological architecture that comes richly equipped with content-imparting mechanisms, then the traditional concept of culture itself must be completely rethought [1992, 115].

The picture turns out difficult to articulate cleanly. As they continue to lay out the project, Tooby and Cosmides say they do not intend to abandon 'the classic concept of culture.' Instead, they say, they are 'attempting to explain what evolved psychological mechanisms cause it to exist' [1992, 118]. There is evidently some ambiguity, but in at least some moments, evolutionary psychology is not the aggressive eliminativist program human sociobiology seemed to be.

1.5 Setting the Standard for Evolutionary Psychology

E.O. Wilson [1975], in reflecting on the understanding of human nature, mused that 'in the free spirit of natural history' humans are, like other organisms, the products of natural history, with their own characteristic suite of adaptations. In Wilson's hands, these included a variety of features. Humans are aggressive, xenophobic, deceitful, and 'absurdly easy to indoctrinate.' In part Wilson pleads that an informed biology should be capable of reforming social science, and psychology. This sentiment is certainly shared with contemporary evolutionary psychologists. Wilson's critics, as is well known, were vocal and uncompromising. Wilson's defenders saw his critics as dogmatic ideologues, pursuing a political agenda at the cost of scientific objectivity. A similar view is often voiced concerning critics of evolutionary psychology. On both sides, the common thought is that the other side has deserted respectable scientific standards. I think we should stick to the standards, both biological and psychological.

We should not demand the impossible. It would be unreasonable to set standards of evidence for evolutionary psychology that evolutionary biologists would not, or could not, meet. The goal is after all to reach a reasonable assessment of evolutionary psychology, not to peddle some fatuous form of philosophical skepticism. If we impose excessive standards, we can surely rule out evolutionary psychology. It could also rule out evolutionary biology. To rule out evolutionary biology would be to rule out the background in which the discussion takes place. It would be to take the discussion outside the realm of science. It also would be unreasonable. On the other side, it also would be unreasonable to set

the standard at a level that is insufficiently demanding. We should not expect evolutionary psychologists to meet standards of evidence evolutionary biologists do not meet, though we can expect them to meet similarly demanding standards of evidence. Evolutionary psychology offers proposals *within* the broad framework of evolutionary theory. This gives us one central charge relevant to evaluating the claims of evolutionary psychology. The simple fact is that evolutionary hypotheses are subject to a variety of empirical tests, though rarely the kind of test that would warm the heart of more narrow experimentalists. Generally, the task involved in assessing an evolutionary explanation — whether of human psychology, or of clutch size in birds or of flower structure in orchids — is one of assessing the historical antecedents and context given both the contemporary forms and the relevant historical record. These antecedents in turn provide evidence essential for a defensible evolutionary explanation. The assumption I begin with is that the sorts of conjectures defended by evolutionary psychologists should be held to the same standards that properly are demanded of evolutionary explanations. If we are offered evidence that would not suffice as evidence for an evolutionary explanation of clutch size in birds or flower structure in orchids, then we should reject the explanation. If the evidence is of the sort that would be taken to support an evolutionary explanation in other domains, then it should be embraced here as well.

2 ENGINEERING DESIGN AND EP

The study of evolutionary adaptation within an evolutionary framework inevitably involves inferring historical process from contemporary products. The aim is to understand the historical sequence and causal antecedents, revealing prior conditions as determinants of contemporary patterns. We infer cause from effect. This is a difficult task, but certainly not one that is impossible. Direct evidence concerning ancestral environments, variation, social structure, and other relevant features are often not available. This makes the task difficult, but, again, not impossible. Some biologists, especially evolutionary ecologists and behavioral ecologists, focus instead on questions of current form and function, abstracting initially from the historical paths that produced what we see. Like archaeologists, they begin with the available artifacts and use it to unravel the history. In using reverse engineering, the evidence for evolutionary psychologists is not initially historical, but concerns, rather, the extent to which a trait optimizes fitness among a specified set of variants and within a specific environment. The standard for fitness in all these cases becomes optimality of design, and in reverse engineering this is gauged by current utility.[2]

[2] A caveat is in order. We could evaluate design relative to an ancestral environment, as advocates of evolutionary psychology often, and correctly, insist. That is no less engineering design. If we also inferred the historical process from that design, it would be reverse engineering. So, for example, ammonoids are coiled cephalopod mollusks from the Devonian. They were able predators, mostly extinct by the end of the Permian. The design evaluation of their predatory

2.1 Reverse Engineering and Adaptation

Reverse Engineering thus is a matter of inferring adaptive function from structure. It aims to infer the historical causes from observed organic form. That is, the goal is to infer the historical function, and then to use this in order to explain current form. If *adaptive thinking* begins with the ecological 'problems' an organism confronts and explains or infers the likely 'solutions' based on the problems, reverse engineering turns the reasoning around, beginning with the 'solution' and inferring what the ecological 'problem' must have been if those traits were to evolve (cf. [Griffiths, 1996]). John Beatty [1980] offers an elegant example of reverse engineering, whose goal is to explain the 'morphological and behavioral design of organisms' [Beatty, 1980, 533]. Beatty draws on a case describing the predatory behavior of preying mantids. Mantids have a forelimb with a claw, with which they grasp prey. Beatty explains that Holling [1964] examined the maximum size prey which could be efficiently grasped and retained by the mantids. Assuming that larger prey is a more efficient food source, and that efficiency in feeding will maximize fitness, Holling's models implied that this should be the favored size for a food source. Notice that we begin with biological form and infer the adaptive function. We do *not* begin with an adaptive problem — say a menu of available prey of various sizes — and infer what would be an optimal form to utilize the prey that are available. This would be adaptive thinking. Instead, we look first to the mantid's claw structure, and infer what its function must be, or must have been. One virtue of these sorts of models, and optimality models generally, is that it allows us to sidestep the relative paucity of evidence we sometimes find when we want to understand the historical processes, and to try to infer something about the process from the outcomes. When it turns out that *Mantis crassa* does preferentially attack prey of the predicted size, Beatty seems willing to infer that the structure and the behavior are the consequence of natural selection [1980, 545]. We infer the historical function from current form. This is reverse engineering (cf. [Richardson, 2003a; 2003b]).

The general strategy should be familiar. In *Darwin's Dangerous Idea,* D.C. Dennett applauds reverse engineering as 'the crucial lever in all attempts to reconstruct the biological past' [1995, 233], and claims this is *the* central 'feature of the Darwinian Revolution: the marriage, after Darwin, of biology and engineering' [1995, 186]. Here's one example he uses:

> Did *Archaeopteryx*, the extinct birdlike creature that some have called a winged dinosaur, ever really get off the ground? Nothing could be more ephemeral, less likely to leave a fossil trace, than a flight through the air, but if you do an engineering analysis of its claws, they turn out to be excellent adaptations for *perching on branches*, not for *running.* An analysis of the claw curvature, supplemented by aerodynamic analysis of archaeopteryx wing structure, makes it quite plain that the creature was *well designed for flight* [1995, 233].

capabilities could lead to evaluating their evolutionary origins. In fact, it has.

As a merely historical matter, reverse engineering is surely not the central feature of the Darwinian Revolution; moreover, I think it is less than clear that *Archaeopteryx* was a full blown flyer. Let's bypass these questions, however interesting they may be. Reverse engineering may nonetheless be an important approach to the understanding of adaptation, as Dennett claims it is. In fact, it *sometimes* is, though it is hardly the pancea Dennett imagines it to be. Engineering design is designed to take us *from* structure *to* function and then *from* function *to* history. We are supposed to infer that mantid claw structure is designed for specific prey sizes, and therefore that it is an adaptation for capturing selected prey. Likewise, we are supposed to infer that *Archaeopteryx* was airborne from the structure of the foot, and therefore that *Archaeopteryx* was adapted for flight.

Tooby and Cosmides, similarly, claim that the evolutionary process is analogous to the intentional construction of devices, and so can be understood in terms of reverse engineering. The problems 'solved' by evolutionary design are limited by the fact that no design is successful unless it either provides a solution to a problem that enhances an individual's fitness, or offers a solution to a problem that increases the fitness of kin. Here is what Tooby and Cosmides say:

> At its core, the discovery of the design of human psychology and physiology is a problem in reverse engineering: We have working exemplars of the design in front of us, but we need to organize our sea of observations about these exemplars into a map of the causal structure that accounts for the behavior of the system [1992, 55].

Applied as a methodological tool for evolutionary biology, then, we begin with a behavioral pattern, and infer 'the causal structure' responsible for the behavior we see; for an evolutionary psychologist, seeing *how* we behave informs us what the function was. It helps us 'to organize our sea of observations.' In the case of evolutionary psychology, all that changes is that *we* are the organism of interest, and that *our own* behavior is what we seek to explain.

Here is a favored example. M. Profet [1992] suggests that 'pregnancy sickness' is an adaptation designed to limit maternal ingestion of substances that might harm a developing fetus. The key thought is that plants contain an array of toxins to repel herbivores. In some cases, the function is to regulate when fruits are eaten. (Unripe fruits are less desirable, while ripe fruits are advertized by color.) Likewise, leaves are typically bitter, which would be a design to repel herbivores. Yet again, since seed predators are a problem, plants invest in toxins to deter them. The Kentucky coffee tree, *Gymnocladus dioicus*, is laced with toxins to prevent ingestion by predators. Caffine and nicotine also are toxins, which is why eating raw coffee seeds is not a good idea, and why cigarettes are an equally bad idea. Apple seeds contain arsenic to deter seed predators, though wrapped in a fruit to encourage animals which will disperse the seeds without crushing them. Profet supposes that humans must have evolved in response to plant toxins. Assuming these toxins are more of a problem for a developing fetus, he concludes that it would be natural for a pregnant female to develop an aversion to the toxins in

order to protect the fetus. Since 'pregnancy sickness' does limit what women can eat, all we need assume is that our bodies are inclined to reject what is unhealthy. In this case, what is unhealthy is what would be harmful to a developing fetus. Plants toxins dissuade animals from eating them, and so the aversion to eating might appear to be protection. So spicy foods, which are often mild toxins, are not favorites early in pregnancy.

2.2 Spurious Design Arguments

Reverse engineering is certainly not unproblematic as an explanatory strategy. Gould and Lewontin famously, or infamously, charged that the use of optimality — including reverse engineering — resulted only in poorly confirmed 'just so' stories. Michael Ghiselin echoes Gould and Lewontin:

> Panglossianism is bad because it asks the wrong question, namely, What is good?... The alternative is to reject such teleology altogether. Instead of asking, What is good? ask What has happened? The new question does everything we could expect the old one to do, and a lot more besides [1983, 363].

How one understands appeals to engineering design has distinctive consequences for the understanding of adaptation. If engineering design is a good method for uncovering adaptations, that favors Dennett, Cosmides and Tooby, and Buss. If it is not, then that is support for Gould and Lewontin, and Gheselin.

Steven Vogel recognizes that the designs of nature are imperfect, and incomplete. He nonetheless finds the assumption that there is a 'decent fit between organism and habitat a useful working hypothesis' [1988, 10]. He has a number of fascinating cases. One vivid illustration of physics at work in the biological world derives from Bernoulli's principle (after Daniel Bernoulli 1700–1782). This principle explains why airplane wings give lift. The pressure from a fluid decreases as the rate of flow of the fluid increases. So consider a simple plane surface with a hole and fluid flowing horizontally (it could be liquid or gas, but it must be a fluid with relatively low viscosity). As the fluid moves more rapidly across the opening, the pressure from the fluid will decrease. If we start at equilibrium above and below the surface, the result would be to draw from below to above at a rate proportional to the difference in rate of flow. As it turns out, we find the principle at work in Prairie dog burrows. It is also what explains why the draw on a fireplace is better with a raised chimney, or a breeze. Burrows have more than one opening. Changing the shape and height of the openings will bring Bernoulli's principle into play. Air flow will be more rapid for raised openings, and that will create a pressure differential with air motion through the burrows. A slight breeze will then draw air through the burrow toward the elevated opening. Showing exactly how this works depends on burrow lengths, depths, and the distance between openings. Prairie dogs adeptly work to maintain a difference between the heights of openings in such a way as to maintain circulation. This has the consequence

that air is replenished, and the animals are not asphyxiated in the bottom of the burrows. Of course, prairie dogs did not read Bernoulli, but they 'understood' the principle. It does not take much brain power to use the principle, even if it takes a genius to articulate it.

These are cases of reverse engineering, moving from organismic design (or behavior) to adaptive explanation. They are similar to cases we have already considered. However, they differ in at least one important way. In the cases from Vogel, there are systematic and independent reasons to impose the constraints on design that the explanations depend upon. The issue is whether the constraints defining what counts as optimal form (or behavior) are prior to the assessment of fit between form and function. In assessing the function of prairie dog behavior, information concerning physical constraints can be incorporated into evolutionary models as prior constraints, independently established, on the available range of biological form. It is useful to think of this as an *a priori* structuring of the 'problem.' Environmental information gives a structure to the problem, constraining the alternatives systematically, based on independent physical principles. It is, unfortunately, difficult to use this approach generally, even if it is illuminating when we can. George Lauder observes:

> The claim that relevant design criteria can be specified *a priori* to allow the analysis of biological design amounts to a claim that we can specify in advance the problem or problems that the design is supposed to solve. Although it is almost always possible to specify *some* design criterion, the more complex the design, the less likely it is that we will be able to determine what the relevant performance and mechanical functions are that any given structure needs to solve. And furthermore the less likely it is that we will be able to meaningfully weigh alternative performance goals [1996, 71].

The difficulty is simply that *a priori* design constraints are often not specified, and in many cases cannot be specified before the fact. The general point is that interaction among features, or design constraints, increases the dimensionality of the design space, so that 'solutions' become increasingly hard to come by.

When we lack principled reasons to impose design constraints, we *can* introduce design constraints *a posteriori*, using the performance of the phenotype to determine the environmental factors which shaped that design. This is exemplified in the case I've already described above from Beatty [1980], and is Dennett's favored approach. To take another example, G.C. Williams claims the human hand is a consequence of selection, while recognizing that this depends on the sort of variation that was historically present and that there are alternative 'designs' which might work equally well for grasping. This is a variant on the venerable problem of the vertebrate limb. Selection *could* explain, as he says, the presence of five digits (even if six or four might result in a better-engineered tool for grasping), as a result of frequency-dependent selection. It *could* also explain quantitative characters, such as digit length. It *might* even explain why, when the functional demands

change as they did with the horse, or the bird, there was a loss of digits. Williams does not claim any of these *actually were* selected *for*, as far as I can see, and may intend them primarily as illustrative. Lauder takes the example to task, as one for which there is too much freedom in defining the constraints on design. The hand is certainly a complex device, with 26 separate bones, in addition to nerves, blood vessels, and tendons. There are five metacarpals , fourteen phalangeal elements which constitute the fingers, and seven carpal bones. The hand is very complex structurally. Evidently it is designed for grasping. But what of these structures are adaptations for grasping? Why, for that matter, are four fingers and thumb better for grasping than three or five fingers and a thumb? And of course, there are many creatures with five digits that do not have our manipulative abilities. It is at least clear that much of the structure reflects ancient features common to vertebrate forelimbs. The basic structure is surely independent of the need for manipulation, from an evolutionary point of view. The pentadactyl hand is not an adaptation, or at least not an adaptation for grasping in humans. As Lauder says,

> ... it is clear that the possession of independently mobile jointed elements ('fingers') is not a design component that could be linked to any specific function that is unique to the human hand: fingers are an ancient design feature of the vertebrate forelimb ... and occur in many animals that do not have the manipulative abilities of the human hand [1996, 75].

Of course, the hand *is* used for manipulation. The ability to grasp and manipulate tools doubtless was influenced by natural selection. Saying how this is so in any informative way is just not that simple, and acknowledging this does not suffice to explain the complex structure exhibited by the hand. The five-fingered structure supports human dexterity. It also allows for the ability to type. Beginning here leaves us woefully short of anything resembling a *principled* analysis of design. We lack the needed constraints on design.

When we have only *a posteriori* constraints, adaptation (or adaptiveness) is not tested as much as it is assumed. When we have only *a posteriori* constraints imposed after the fact on models of organismic structure, as in the case of the hand, again there seems little substantial and independent support for the adaptationist conclusions. The point applies to the agendas of Dennett, of Buss, and of Cosmides and Tooby. They do not begin with the environmental constraints, independently specified. They claim to have evidence concerning design. This is what would give them traction in the goal of reforming psychology. Having assumed the 'fact' of design, the complex structures and behaviors we see are 'explained' as the consequence of natural selection, often without independent evidence. They do not argue *for* design but *from* form to function and then again *from* function *to* form. Without the traction offered by prior constraints, design will not, and cannot, support the aspirations of evolutionary psychology.

3 DYNAMIC MODELS OF NATURAL SELECTION

What is evidently needed is an independent method for assessing historical function. The point is, after all, to understand the function of, say, reasoning or language for our Pleistocene ancestors; our understanding of the architecture and functioning of reasoning or language, in turn, is supposed to be informed by recognizing how reason functioned *for our ancestors*. Thus, as evolutionary psychologists commonly urge, to claim that something is an adaptation is to advance an historical claim. An *adaptation* is the product of evolution by natural selection (cf. [Brandon, 1978; 1990; Burian, 1983; Burian and Richardson, 1992; Richardson, 1996]). Traits may be present because of chance, or drift; they may be genetically linked to traits that are selected for; or they may be pleiotropic effects. In any of these cases, a trait may be beneficial and improve fitness, but that does not make it an adaptation. A trait is an adaptation only if its prevalence is due to the fact that it conferred a greater fitness. The task of explaining some trait as an adaptation thus depends on understanding the evolutionary history which produced it. This is the root evolutionary problem which evolutionary psychologists face: they claim that human psychological traits are adaptations, and seek to explain them as adaptations to ancestral environments. The task is then one of reconstructing evolutionary history from a contemporary record, together with some modest historical information. This is not an impossible task, but it is a difficult one.

If this can be done, it answers perfectly to the ambitions of an evolutionary psychology. Here again is the problem, according to Cosmides and Tooby:

> Like a key in a lock, the functional organization of each cognitive adaptation should match the evolutionarily recurrent structural features of its particular problem domain ... Because the enduring structure of ancestral environments *caused* the design of psychological adaptations, the careful empirical investigation of the structure of environments from a perspective that focuses on adaptive problems and outcomes can provide powerful guidance in the exploration of our cognitive mechanisms [1994, 534]

This turns the question around, treating it as one of adaptive thinking rather than reverse engineering: we look first to determine evolutionary function, and infer form. The question, from this perspective, is what kind of information we need in order to construct a reliable evolutionary explanation for an adaptation, beginning with function rather than form.

4 DYNAMIC MODELS OF DESIGN

In *Adaptation and Environment* [1990, ch. 5], Robert Brandon, lays out five conditions on what he calls 'adaptation explanations.' If fully specified, we would have a sufficient model for evolutionary dynamics. These are conditions which

are only infrequently met perfectly. They provide, in Brandon's view, something like an ideal type for explaining adaptation, in the sense that a full and complete evolutionary explanation of an adaptation would draw on evidence to answer all of the conditions, though we often have cases in which we have only partial answers. Most evolutionary explanations fall short of the ideal in one way or another, but the four conditions nonetheless provide a reasonable yardstick to consider in assessing the adequacy of specific evolutionary proposals. These are the five conditions:

1. There must be evidence that selection has occurred: we need to distinguish selection from other evolutionary causes, and to reveal something about the character and strength of selection, given that it is a factor; this will require that we have some information about the character and extent of variation in the ancestral forms.

2. There must be an ecological basis for selection: there must be ecological factors which are responsible for the strength of selection, whether this is some factor in the external environment such as an environmental toxin, or some factor in the social environment such as sexual selection.

3. Differences between individuals must be heritable: there must be a correlation in the phenotypic traits of parents and offspring which is greater than would be expected by chance.

4. There must be information concerning the environment, the population structure, and gene flow: the effective size of populations, population structure, gene flow and interbreeding, as well as mutation rates do affect the rate of evolution; we need to be able to fix these parameters if we are to distinguish various evolutionary scenarios empirically.

5. We must distinguish primitive and derived traits: we need to know whether a trait is ancestral within a clade, or whether its presence is the result of convergence among lineages; in general, we need to have an independently established phylogeny.

The example Brandon relies on for support and illustration is the evolution of heavy metal tolerance in plants, which is a clear and well documented case of adaptation (summarized in [Antonovics, Bradshaw and Turner, 1971; Jain and Bradshaw, 1966]). Here is how the example works. Grasses that can grow on soil contaminated with heavy metals, such as the tailings from mining sites, evolve a tolerance for the high level of metals in the soil. These locations not only have an overabundance of heavy metals, but also are low in nutrients. There are a surprising number of grasses that can adapt to the severe conditions these sites provide. The work illustrates Brandon's five conditions. (1) Selection can be measured directly, by survivorship. In the case of heavy metal tolerance, the intensity of the selection is very high, sometimes apparently as high as 95%. Since the selection is intense, and continual, it can be observed relatively directly. (2) It

is possible to identify the ecological factor driving selection. In many cases, it is possible to show that selection occurs, but very difficult to provide a convincing mechanism. In some cases, it is clear what the ecological factor is. In the case of heavy metal tolerance, just as intolerant types perform poorly on contaminated soils, tolerant types perform poorly on better ground. The physiological cause for this difference in performance is not exactly clear, but lower nutrient levels are likely to be important, as is the need to sequester the heavy metals. (3) Estimates of heritability are difficult to obtain with any precision. Heritability is a statistical measure of the similarity between parents and offspring in a common environment, and is critical for evolutionary models because it determines the response to selection. In the case of heavy metal tolerance, the genetic basis for tolerance is not well understood, but it is clear that heritability values are relatively high [MacNair, 1979]. (4) It is also important to know a good deal about population structure. For example, it is important to know the relevant scale so that sampling does not occur on a larger scale than an effective population. It is important to know immigration and emigration values, so that the rate of change due to selection is not confounded with other sources of change. In at least one of the grasses, *Anthoxanthum*, there are different flowering times for the different strains, which have the effect of sexually isolating intolerant types from tolerant forms because they are wind pollinated. This tends to promote evolutionary divergence. (5) It is critical to know trait polarity. This is central to Lauder's point, discussed above. Since pentadactylity is an ancestral trait, present in many organisms where there is no grasping, it is unreasonable to infer that pentadactylity is an adaptation for grasping. In the case of heavy metal resistance, because the sites are of relatively recent origin, on mine tailings, it is unproblematic to identify the ancestral state. The ancestors were not resistant, or were less resistant, than the derived forms.

These five conditions offer no simple test for adaptations. They can be thought of as illustrating how a complete explanation in terms of natural selection would be constructed. Certainly, some explanations will fall short on some of the conditions, but may still count as reasonably good evolutionary explanations. An ideally complete explanation would do well on all five conditions. An explanation which falls short on one or more is correspondingly deficient as an explanation.

4.1 Natural Language and Natural Selection

Let's look to see how this these criteria work in a concrete example within evolutionary psychology. Steven Pinker and Paul Bloom [1992] defend the idea that human language is the product of Darwinian natural selection. The principle argument offered by Pinker and Bloom is simple and elegant. Human language certainly does have the mark of an adaptation. Human language involves a complex capacity, governed by rules and representations which are not taught in any straightforward way. There are rules governing phrase structure which provide information about the underlying structure and meaning; grammars depend on

lexical and phrasal categories that simultaneously limit what can be said and give language its remarkable expressive power. Children acquire language with remarkable proficiency, without formal instruction, a fact that is most naturally explained by the presence of substantive prior constraints on languages which facilitate learnability. Its universality among human societies alone suggests that it is part of our evolutionary heritage. It is, moreover, the kind of complex capacity which is plausibly the consequence of natural selection. Pinker and Bloom exploit these general features in concluding that language shows 'signs of design' [1992, 459]. Language does seem to be a natural candidate for explanation in terms of evolution by natural selection. It *does* show 'signs of design.' Do we have anything more substantive than this general suggestion? Do we have anything approaching a dynamic model for the evolution of language capacities? Let's look at how well scenarios for the evolution of human language fit the ideal type for adaptation explanations.

1. The general case which Pinker and Bloom offer does support the claim that there has been selection for linguistic capacities. We know nothing about the strength of the selection, or its character. It is not clear from the scant evidence we have whether language arose with modern humans, or is older (cf. [Lieberman, 1992]). If the origins of language are relatively ancient — following perhaps the enlarged brains of *Homo habilis* or of *Homo erectus* — then the complex suite of adaptations for speech would be intelligible, though they would no longer be derived traits within *Homo sapiens*. The kind of information this gives us is woefully inadequate to the task set for it by evolutionary psychology. Selection requires contrasting groups, or contrasting individuals, differing in trait values. One group might have a trait, and another lack it. One individual might have it while others lack it. If the selection is within a single population, individuals would have to possess a trait to varying degrees. Selection, after all, requires variation. In either case, selection depends on prior variation and an explanation in terms of selection depends on knowing something of the character of that variation. We have no information concerning the kind of variation, or the effect it would have on fitness. Both would be necessary to know the strength of selection. We have nothing substantive concerning the kind or character of the selection pressures influencing the emergence of language.

2. What ecological basis do we have for selection on linguistic abilities, assuming there is such selection? The answer we are offered is that it facilitates communication. This is no doubt true. It is also nearly wholly inadequate. Communication can, after all, take many forms. Dolphins utilize a distinctive set of whistles. Birds use a variety of visual displays in mating. Bees convey the location of pollen sources using ornate dances. Ants communicate by laying down pheromone trails. No doubt the use of all these abilities were favored because they facilitated communication. In order to explain why our linguistic abilities were favored by natural selection, we need a much more

articulated hypothesis, explaining the feature of human language that make is so striking.

Palaeoecology has given us a good deal of information concerning the environment relevant to hominid evolution, reconstructed from such things as the distribution of extinct flora and fauna. We know, for example, that during the Pleistocene (4-1 million years ago) australopithecines were widely distributed in Africa, occupying both woodland and savanna; and there is reason to think that climatic changes were significant in the emergence and diversification of hominids. The various species of *Homo* that subsequently developed likewise occupied a variety of habitats, and in spite of the considerable variation among them their appearance is generally associated with the presence of tools. Such information offers scant help in understanding the factors which shape the evolution of language. What factors in our ancestral environment favored verbal communication? Doubtless, they are largely social factors. How were these factors different from those which shaped the communicative abilities of our prelinguistic and protolinguistic forebears? How did the linguistic abilities of humans differ from the corresponding abilities in related groups in the genus *Homo*? Why did any of this favor a system of rules and representations? These are hard questions, not yet addressed in any significant way. It is in fact hard to see what evidence would be relevant to resolving them. It is difficult but not impossible.

3. It might seem easy to answer the question of heritability. It is part and parcel of contemporary linguistic theory to emphasize the underlying unity of human languages, and to downplay their diversity. The underlying similarities which characterize human language are both universal and deep, providing explanations of both the similarities of existing languages and of the ease of language learning. Given constraints that are both universal and contingent, the only explanation would seem to be that they are somehow innate. As evolutionary psychologists emphasize, what is universal must be inherited. Unfortunately, for more or less technical reasons, this does little to address the issue of heritability. Heritability is a measure of the significance of genetic variance. It requires that there be significant variance to begin with, and if the variance is zero then heritability is undefined, not settled. What we need is a measure of variance for linguistic competence, and measure of variances within and between different varieties for ancestral populations, in order to assign heritability values. This is simply unknown. Given how little of human variance is genetic, it is difficult to see how such values could be obtained, limited as we are to contemporary populations.

4. What do we know concerning the structure of the environment and population relevant to gene flow? For language, as for human reasoning and for most cognitive and emotional factors, the relevant factor is surely social organization. There is, of course, a great deal known about primate social organization. There is something known, too, about early hominids;

for example, they were habitually bipedal, and the importance of tools and hunting with the earliest forms of *Homo* as well as their *Australopithicene* ancestors, indicates a level of cooperative behavior. Sociality was doubtless important. Beyond this we have speculation and extrapolation from contemporary groups and primates. We are left with very little firm knowledge of the social structure of early hominid societies.

5. Human language is a derived trait. This much at least seems a fixed point, though it may not help in the long run. The obvious problem is that we have no information that is helpful in reconstructing the ancestral traits, since none specifically helpful to the evolution of language are present in contemporary primate cousins. The basic problem is located nicely by Terrence Deacon:

> Why did language first evolve? In hindsight it is easy to justify as a useful means of communication, but the fact that it evolved only once gives us no comparative perspective. Hominid ancestors were the first to achieve the necessary mental conditions for language, but we do not know what these were [1992b, 133].

Understanding how a complex trait evolves depends on knowing the primitive traits in ancestral groups which make its evolution possible. Again, we are ignorant of what we would need to know in order to offer an explanation.

The overall assessment is dismal. It is plausible to think that our linguistic abilities are the result of a history involving selection for those abilities. However, in the face of the request for a substantive explanation of the character of the selective forces, their ecological basis, information about the heritability of the relevant traits, the genetic structure of the population, or the relevant ancestral traits, we are given nothing, save the facile suggestion that the evolutionary function of language is to facilitate communication. This is almost certainly true, but is so general as to be uninformative. There are many systems of communication, of which human language is only one. The complex suite of characteristics which make it so plausible that human language is an adaptation, moreover, are not explained by this very general characterization of the evolutionary function of language. Human language does facilitate communication. How are we to explain the lexical and phrasal categories incorporated in human language? Dominance rules and precedence rules governing linear order are also part of human languages. How do we explain them? These are the kind of features Pinker and Bloom appeal to in their 'argument for design in language' [1992, 460 ff.] Yet for these features, we are offered no explanation. In the absence of any substantive constraints, we are no better off here than in the case of the pentadactyl hand. We have a device that is evidently evolutionarily constrained, and which was certainly subject to natural selection, but very little that sheds substantive light on the evolutionary factors that shaped it.

5 PHYLOGENY AND ADAPTATION

A trait, behavior, or character is *adaptive* only if that trait confers an advantage relative to other organisms which lack that trait. A trait, behavior, or character is an *adaptation* only if that trait is present because in the past it did confer an advantage relative to other organisms that lacked that trait. Adaptations are thus doubly relative. A trait is adaptive relative to alternatives present. An adaptation then must have been adaptive relative to historically available alternatives. Moreover, a trait is adaptive relative to a given environment, or a specific selective pressure; what is adaptive in one setting may be maladaptive in another. An adaptation then must have been adaptive in some ancestral environment, and thus relative to its previous ecological setting. Current use or advantage is not enough to show that a trait is an adaptation. To use an example from Darwin [1859], skull sutures are certainly adaptive in humans, since they facilitate passage through the birth canal. They have an important, and even essential, current use. Skull sutures are nonetheless not an adaptation for parturition, since birds and reptiles also have skull sutures, though they hatch rather than being live born. As Darwin says, even though these skull sutures 'no doubt facilitate, and may even be indispensable for' mammalian birth, they are not adaptations for that function. Useful features are not necessarily adaptations, and if they are the consequence of natural selection, they are not necessarily adaptations for their current function. Whether a trait is an adaptation depends on historical antecedents and conditions. We need to know those conditions if we are to offer an acceptable evolutionary explanation. Dynamic models hinge on more or less direct access to those conditions. It is our lack of such access that makes reliance on dynamic approaches problematic. Reverse engineering likewise lacks the kind of evidence — in this case, concerning the environmental parameters — we would want to make it a reasonable approach for human cognition.

5.1 The Comparative Method

There are explicitly historical methods for articulating evolutionary explanations. One is the 'comparative method,' an approach that has recently gained considerable currency among evolutionary biologists in studying macroevolution as well as microevolution. The comparative method requires that we compare a trait to phylogenetically related ancestors and conditions. These constitute the evolutionarily relevant alternatives within related evolutionary lineages. If it turns out that a trait would not enhance performance or survival relative to the actual variants among its ancestors, then that would undercut the claim that that trait is in fact an adaptation. Similarly, if a trait originated in a lineage for which the current function would be irrelevant, then that would undercut the claim that that trait is in fact an adaptation, whether or not it is currently adaptive.[3] To draw again from

[3]It is often thought, mistakenly, that this would preclude traits maintained by stabilizing selection. That it does not is easily seen. If a trait is present due to stabilizing selection, then it

Darwin's *Origin*, there are structures, hooks, on the branches of climbing variants of bamboo. Darwin was not oblivious to the adaptive functions they serve for these climbing species, but nonetheless he rejected them as adaptations. The reason is simply that those same structures are present in other, nonclimbing, relatives. The hooks, Darwin concluded, are ancestral, present before climbing was part of the behavior. Since the structures precede their 'function,' that function cannot explain their presence.

The comparative method makes evolutionary history the centerpiece for the analysis of adaptation, encompassing the dual relativity of adaptation. Phylogenetic analysis reveals an historical perspective on the pattern of acquisition of traits within lineages, allowing us to infer the adaptive character of traits from the order of appearance of those traits. The pattern of acquisition or loss of a putative adaptation can then be mapped onto an independently established phylogeny. Coupled with knowledge of the selective regimes associated with ancestral forms, we compare alternative characters for their adaptive potential. The method is powerful when it is applicable, but it is not easy to apply: we need comparative data on related taxa, developmental information, information concerning the character of the environment, the trait family under consideration, and the relative adaptiveness of the traits characteristic of the several taxa (see [Baum and Larson, 1991; Brooks and McLennan, 1991; Coddington, 1988; Harvey and Pagel, 1991]).

The basic idea of phylogenetic systematics is straightforward: given a group of related organisms, it is possible to recover the phylogenetic relationships within that group on the basis of shared similarities among extant or historical species. A phylogenetic tree is typically given as a branching diagram, representing a series of speciation events; such a tree is essentially an hypothesis concerning the genealogical relationships among a set of taxa. If we begin with a monophyletic group (or clade) consisting of an ancestral species and all its descendants, the problem is to determine which of the possible phylogenetic trees is the actual one. With even a small number of species, the number of possible trees can be quite large. In order to determine the correct tree, we consider the traits characteristic of the species within the clade. In theory, similar characters could be due to common descent (homologies), or to parallel or convergent evolution (homoplasies). Cladists, following Willi Hennig [1966], use the patterns among shared derived characters to construct a phylogeny representing the pattern of ancestry.[4] They start with the conviction that genealogical influences should be sufficiently pervasive that homologous characters should predominate over homoplasous ones. Convergent and parallel evolution, that is, should be less common than similarity due to common descent. Whenever possible, then, cladists assume that similar characters are

must have been maintained in a lineage by selection. That too is a comparative issue, depending on whether, in phylogenetic ancestors, there was variation or whether, in phylogenetic ancestors, that variation would have been maintained.

[4]Cladistics is not the only approach to phylogenetic reconstruction. It is actually simpler, and more forgiving, than some other approaches. I do not think a shift to other methods would make the task easier for evolutionary psychologists.

due to common descent rather than convergent or parallel evolution. The correct tree can then be inferred by examining character similarities, with the best tree minimizing homoplasies.

Systematics certainly has had its share of controversies. Even among cladists, there are a number of currently unresolved issues within phylogenetic systematics. These depend, for example, on how to establish character polarities and what counts as a single character (for a good discussion, see [Maddison, *et al.*, 1984]. A phylogenetic analysis depends critically on defining an appropriate outgroup; that is, a closely related species outside the clade used for comparative purposes. Selecting an outgroup is critical because the goal is one of determining what the ancestral, or plesiomorphic, character state is. Ancestral features shared within a clade do not give information concerning the structure of the phylogenetic tree. The insistence on using monophyletic groups is, however, widely accepted even if the particular means for determining which groups are monophyletic is not.[5] The net result, if the data is sufficient to reach a definitive resolution, is that the resulting phylogeny can be used to reconstruct patterns of character diversification; and this in turn can be used as a check on adaptationist explanations.

Adaptations then become traits, or combinations of traits, correlated with environmental variables, which have arisen within a lineage as an adaptive response to common selective pressures. In evaluating the significance of a putative adaptation, the comparative method requires information concerning phylogenetic relationships, the ancestral environment, and the factors affecting its evolution. The comparative method requires assessing the character under study to alternatives, including in particular the alternatives present within the ancestral environment. If the selective regime offers a selective advantage to one variation within the lineage, that is consistent with an hypothesis of adaptation.

Comparisons across species are useful in evaluating adaptation, and must take into account their respective environmental settings. With similar environments, similar traits will reflect common ancestry or stabilizing selection if the traits are homologous; similarity will indicate convergent evolution if the traits are homoplasious. If the environments are not similar, then similarity of traits suggests phylogenetic constraints are critical or that there is convergence that is not environmentally driven. One quick example may serve to illustrate the usefulness of the comparative method for evaluating adaptationist hypotheses, and to suggest some morals that might be derived from comparative studies. Anoles lizards are a very diverse group in the Iguanid family (see [Losos and Miles, 2002]). Anoline lizards have proliferated on Cuba, Hispaniola, Jamaica and Puerto Rico and oc-

[5] Once we accept that some characters may evolve more than once, there are a number of procedures for selecting the best tree (see [Felsenstein, 1982] for an extensive discussion). Allowing derived characters to evolve more than once, the most parsimonious tree might be the one with the fewest character state derivations. If we hold that complex characters cannot evolve more than once but can be lost repeatedly, the most parsimonious tree would be that with the fewest losses of derived characters. We might hold simply that the most parsimonious tree is the one with the fewest character state transitions. The various criteria result in significantly different choices, but such disputes do not centrally affect the points I will be making here.

cupy a diverse range of habitats. On each island, the range of habitats is similar and the lizards have developed similar morphologies. Some arboreal forms have toepads that allow them to hold on to smooth surfaces. (Others that live on the ground lack this particular feature.) Functional studies have confirmed that these toepads are important for their ability to cling, and phylogenetic studies support the view that arboreality evolved before the toepads. Moreover, though there are morphologically similar types on the four islands, similar morphological types have evolved independently; cladistic relatedness is uncorrelated with morphological similarity. This supports the conclusion that toepads are an adaptation for the arboreal lifestyle of the lizards (cf. [Larson and Losos, 1996]), and not merely that they are maintained within the group because of their selective advantage. Generally, the recurrent evolution of similar morphological types in these habitats suggests that adaptation rather than some other process — for example, colonization or developmental constraints — is the cause of the evolutionary changes (see [Losos *et al.*, 1988; Glossip and Losos, 1997; Irschick *et. al.*, 1997; Jackman *et. al.*, 1999]).

Conversely, it is possible to see what kinds of considerations can make a case more ambiguous. The genus *Montanoa*, or so-called 'daisy trees' because of the structure of their florets, includes some thirty taxa. Twenty-one of these species are shrubs, five are trees, and four are vines. The shrub form is apparently ancestral, and the prevalence of this form suggests significant conservatism within the group. Roughly 15% of the species are associated with a change in habitat, and have given rise to tree and vine forms in roughly equal proportion. A study of convergent adaptation might be appropriate for these forms, following the lines of the previous example; and, of course, an adaptive explanation would be misplaced when applied to the shrub form, which is plesiomorphic. The tree forms are clearly derived, and so are candidates as adaptations. Among the tree forms, there are a number of morphological changes present, and though they belong to different clades, these characters allow them to live in higher elevations. As far as this goes, adaptation is a plausible hypothesis. We have a derived form, within several lineages; and it is correlated with significant environmental differences. It is in fact not clear whether an adaptive explanation is right. Diploidy is the ancestral condition among the *Montanoa*, but the trees and only the trees are polyploids. Polyploids are typically larger than diploids within the group, and polyploids often arise from diploids in the group. The distinction between trees and shrubs may in fact be a simple by-product of changes in ploidy level. It may, that is, be a spandrel (for more discussion of the case, see [Brooks and McLennan, 1991, 159 ff.]).

Phylogenetic methods thus provide a useful and important independent method for framing and evaluating evolutionary explanations. It can be used to rule out some traits as adaptations. It can also lend varying degrees of support to the hypothesis that a trait is an adaptation. The comparative method offers no simple litmus test for adaptation. It may support the claim that a trait is an adaptation, or it may undercut it; but there are a number of specific morals. First, unless a trait is derived within a lineage, it cannot be an adaptation. Second, the support

for the claim that a trait is an adaptation is strongest if there are more than one independent evolutionary origins associated with comparable selective regimes. If a character evolves within a specific selective regime more often than would be expected by chance, then that supports the claim that it is an adaptation within that group. Toepads, for example, have evolved in at least three lineages of lizards. This can be subjected in some cases to a statistical test, evaluating whether an evolutionary transition takes place within a common environment more often that would be expected by chance. Larson and Losos [1996] estimate that the probability of three origins for toepads is roughly 0.028, which supports the claim that it is indeed an adaptation to an arboreal habitat. Third, since historical variables concerning rates of evolution, ancestral environment, and branching of lineages, are more available for some taxa than others, there is no guarantee that a phylogenetic study will yield a definitive answer; nonetheless, adaptation does indicate a number of results which would not generally be expected under alternative evolutionary scenarios.

5.2 Phylogeny and Adaptation in Human Evolution

How might this apply to the human case? There is a remarkably good fossil record relevant to the advent of modern humans, and there is a great deal known about our primate kin. This might seem a reason to be optimistic over the prospect of using comparative data for the purposes of evolutionary psychology. We do, after all, have independent information about the relevant phylogenies based on the fossil record and on contemporary species. There are, moreover, some reasons for viewing human psychological capacities as adaptations that are largely consonant with comparative methodology. The grounds for optimism turn out to be illusory insofar as they are geared toward elaborating evolutionary explanations. A look at the state of the art in hominid evolution should suffice to make the point. Impressive as it is, what is known of hominid phylogeny and what is known of the evolution of these traits is not enough for the program of evolutionary psychology. We are far short of an ideal subject for the comparative method.

Within the primates, apes form a well-defined group. They are a clade; unfortunately, it contains only five living genera and of those only two have more than one living species. The fossil record for the apes is disappointing, with few fossil remains. The great apes probably originated in the middle Miocene (11-16 mya). The most famous specimen is probably *Proconsul*, though there are a number of genera representative of ancestral hominoids (see [Ward, 1999]). Some apes are more closely related to humans than to other apes, and from a cladistic perspective we need to be included among the apes.[6] It is clear that the Orangutan (*Pongo*) and Gibbons (*Hylobates*) are relatively distant relations of the gorillas

[6]This does not mean that humans evolved from any of the modern apes. Humans and apes evolved from a common ancestor. There are some fossil remains from the Miocene which are plausible candidates for a common ancestor, and it is the record in the late Miocene and early Pliocene which is especially lacking. This is evidently the critical period for the divergence between apes and hominids.

(*Gorilla*), common chimpanzees (*Pan troglodytes*), pigmy chimpanzees (*Pan paniscus*), and humans (*Homo sapiens*). The four latter species form an evolutionary clade, though there is less than unanimity concerning the pattern of relationship among them; that is, there is a question whether, within the hominoid primates, the African apes are more closely related than any is to their human sister group. The problem is that the branching events which distinguished gorillas, chimps and humans were nearly simultaneous; as a result, there is a great deal of similarity among all of us from genetic, morphological, and behavioral standpoints. Morphological evidence tends to favor a grouping of chimpanzees with gorillas, some of which are are plausibly construed as shared derived characters [Kluge, 1983]. Molecular phylogenies favor a grouping of chimpanzees with humans [Sibley and Ahlquist, 1984]. Insofar as there is a consensus, it is that the chimps and hominids form a clade, with gorillas branching off first (see figure 5.2).

Within the hominids, there are a number of fossil forms of australopithecines and several species within the genus *Homo*. All are characterized by a bipedal gait. Like other apes, the hominids are sexually dimorphic, probably with long gestation periods and extended maternal care. Australopithecines have relatively large teeth with small brains. The species of *Homo* are broadly distinguished by an enhanced cranial capacity by comparison with the australopithecines. The relationships among these various hominid forms are uncertain and subject to considerable divergence of opinion (cf. [Chamberlain and Wood, 1987; Strait *et al.*, 1997; Fleagle, 1999]). The several species of *Homo* do at least form a clade distinct from the australopithecines, though some treat *A. afarensis* as a sister group, while others treat *A. africanus* as the sister group. Within the australopithecines, relations among *A. bosei, A. robustus, and A. aethiopicus*, and, more recently, *A. garhi*, are especially controversial, though these species appear to be more closely related to one another than any are to *africanus* or *afarensis*.[7] There are at least four forms of *Homo* commonly recognized, and often more than four (see figure 5.2). It is generally conceded that *H. habilis* is distinct from the remaining groups within the genus. The relationships among *H. erectus*, *H. rudolfensis*, *H. sapiens*, and *H. neanderthalis* are less clear; on some accounts the Neandertals are a subspecies of human and on others it is a distinct species. Insofar as any distinguishing features are behavioral, the fossil record is not decisive. As a result, hominid phylogenies are very much matters of controversy. Figure 5.2 presents a plausible hominid phylogeny, though not the only one defended in recent literature.

Within the genus *Homo*, intelligence has increased. Brain size has certainly evolved, and the pattern of its evolution supports the conclusion that natural selection was important. The brains of australopithecines fall within the range to be expected within the great apes, but by roughly 1.5 million years ago the brains

[7]These robust australopithecines share a number of cranial features, probably related to diet. However, these craniodental features are evidently not independent developmentally and so cladistic analyses which treat them as independent may overweight them. For interesting discussions from a cladistic perspective, see [Skelton and McHenry, 1992; Lieberman, 1995; Lieberman *et al.*, 1996; Skelton and McHenry, 1997; McCollum, 1999].

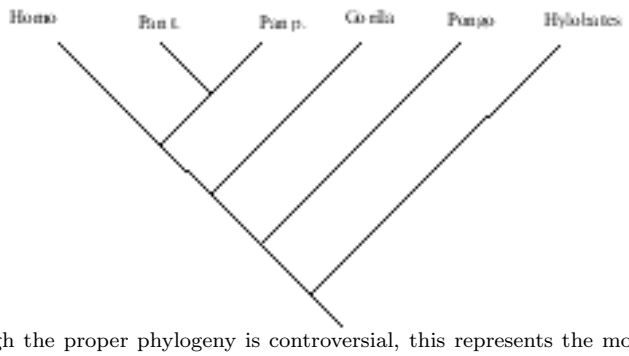

Though the proper phylogeny is controversial, this represents the most commonly held phylogeny among living apes. It is uncontroversial that humans, chimps and gorillas are more closely related to one another than to either Orangs or Gibbons. On the basis of genetic similarity, humans and chimpanzees would form a clade, with gorillas branching off earlier. There are alternative phylogenies, according to which chimpanzees and gorillas form a clade, with the human line branching off earlier.

Figure 1. A phylogeny for living apes

Species	**Height**	**Cranial Capacity**	**Known Dates**
H. habilis	1-1.5 m.	500-800 ml.	2-1.6 mya
H. erectus	1.3-1.5 m.	750-1250 ml.	1.8-0.3 mya
H. neanderthalis	1.5-1.7 m.	1200-1750 ml.	130,000-60,000 years
H. sapiens	1.6-1.9 m.	1200-1700 ml.	130,000 years

The most marked characteristic of the clade is a larger cranial capacity; many other traits, including a bipedal gait are shared with australopithecines. Times are measures in terms of million years ago (mya)

Figure 2. Features of four species within the genus *Homo*

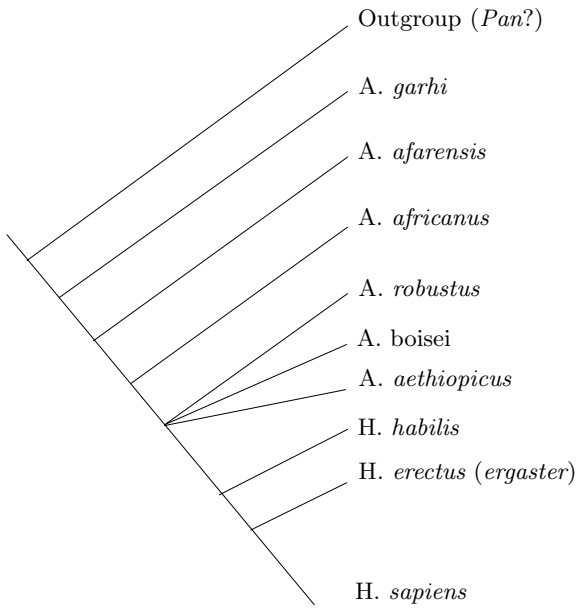

A number of phylogenies have been proposed between the various hominid forms; the resolution of some relationships is not clear. The number of species changes frequently. Even the number and rank of differences is contested. It is generally agreed that *H. erectus (ergaster)* is closely related to *H. sapiens*, that *habilis* is more primitive than either, and that *A. robustus*, *A. boisei* and *A. aethiopicus* are closely related. Some assign these latter three to a distinct genus, though there is controversy over even whether they form a clade. I therefore have depicted this as an unresolved node. There is no clear consensus on whether *A. africanus* or *A. afarensis* is more closely related to the robust forms of the australopithecines, What is given above thus is a plausible phylogeny, though emphatically only one of many defended in recent discussions.

Figure 3. A plausible, but incomplete, hominid phylogeny

of *Homo erectus* averaged more than twice that of the australopithecines. Gross brain size by itself is unreliable as a measure of intelligence. Allometric scaling — a measure of brain size relative to body size (cf. [Gould, 1966]) — further supports the conclusion that the relative cranial capacity of *Homo sapiens* is a derived character. Within mammals generally, brain size can be reliably predicted from body size as a power function. Though there is uncertainty concerning such comparisons, it is generally agreed that at any stage of development primates have a larger brain size for a given body size than would be expected from the mammalian index alone, by roughly a factor of two. Within primates, the growth curve for human infants follows the primate norm until birth; in most primates, the relative growth rate of the body then increases after birth while in humans the brain grows at an elevated rate postnatally. As a result, adult humans also have a brain size that is markedly higher than would be expected given the primate norms (see [Deacon, 1992a; 1992b]). Primates are therefore exceptional within mammals. Humans are exceptional even for primates. Such comparisons are suggestive, and are consistent with selection affecting brain size, but they must be regarded with caution. The comparisons among primates are less problematic than are comparisons across broader taxonomic groups. Structurally, primate brains are very similar. They have the same gross cortical and subcortical structures, with a similar architecture. However, the similarity is not perfect. Aside from its larger size, the human brain is more asymmetrical, and olfactory centers are reduced. Endocranial values alone do not reflect such structural differences. None of the differences in overall size, in any case, do anything to explain the marked cognitive differences between humans and other primates, including, most obviously, our specialized linguistic abilities. If we take seriously the idea that our cognitive organization is modular, then overall differences would not directly reflect the evolution of human cognitive capacities. Organizational and functional differences *are* surely more important than sheer bulk in assessing human intelligence and our communicative facility, and on this the fossil record is largely silent. It is not clear how even the allometrically corrected measures would explain such differences in ability. Realistically, encephalisation quotients are crude measures of intelligence at best. They reflect overall size, but gloss over the kinds of specializations that are likely to explain major differences among species and that are likely to be important to human cognitive capacities. Nonetheless, this enhanced cranial capacity is at least a derived feature within the genus. That is at least consistent with selection for enhanced intelligence. This does nothing to settle what our intelligence might be an adaptation *for*.

Language provides a similar case. Again, the neurophysiological changes are the clearest. They clearly support the involvement of the (left) cortical structures in language. Humans have a larger prefrontal cortex than would be expected. Cortical regions adjacent to the Sylvian fissure are directly involved in language use. Studies of patients with trauma, as well as stimulation studies involving these regions show that they are specialized for language use. Brain imaging supports the general picture though language functions rely on a diverse range of brain

systems especially in the left hemisphere. There are functional and morphological asymmetries between hemispheres in other species as well, even if the extent of the asymmetry in humans is more marked than in other extant primates. (Moreover, other primate calls are controlled in the midbrain and the forebrain rather than in the motor cortex.) As we turn to comparisons with other hominids, the picture is still less clear, largely because we must rely on indirect evidence. Fossil remains do allow us to discern some of the gross features of hominid brains since the surface features of the brain leave some imprints on the skull. Endocasts of fossil skulls display these features, though changes in size and other changes resulting from upright posture make the interpretation of these features difficult. The Sylvian fissure does show some asymmetry in archaic humans, including both *H. habilis* and*H. erectus*. That would support the view that there is an enlarged left hemisphere, and an asymmetry similar to that common in humans. Some researchers conclude on the basis of such evidence that these other species of hominid had specializations associated with language. If so, linguistic capacity would not be a derived character specifically within *Homo sapiens*, but a characteristic of other hominids as well. It is, though, a derived trait within the hominid line, and an adaptation characteristic of *Homo* to the exclusion of the australopithecines. The point is that we do not know the appropriate homonid group to focus on in explaining language. Absent knowledge of the taxonomic rank, we do not know where to place the trait on a tree. Comparative methods are then silent.

Thus, it is plausible to hold that language use and human intelligence are derived traits, somewhere in the homonid line; but this provides only marginal comfort to evolutionary psychology. Our knowledge of the cladistic structure of the relevant species is imperfect, even though it is very good. There is good reason to hold that both human intelligence and human language are apomorphic; that is, they are both derived from and differ from the ancestral conditions. This is at least consistent with the claim that they are adaptations.[8] One thing that is strikingly absent is a knowledge of the relevant homologous structures or behaviors in ancestral groups. In assessing evolutionary adaptation this provides a critical feature in understanding the variation, and thus the relative advantage that structure or behavior provides. It is not clear, furthermore, whether language use is a derived trait within *Homo sapiens*, or a more general feature of the hominids. A parallel point applies to the evolution of a greater cranial capacity. This makes a comparative assessment particularly problematic. The ecological settings for the several species, insofar as we can reconstruct them from paleoecological data, are varied. During the Pliocene and early Pleistocene (roughly 4 to 1 mya), australopithecines were widely distributed in southern and eastern Africa. These habitats included wooded regions, open savanna and grassland. *Homo habilis* overlapped temporally with australopithecines, but is evidently associated more with wooded habitats. With *Homo erectus*, we find more use of open habitats, perhaps indicating increased reliance on hunting. By roughly one million years ago, toward

[8]Cladists limit adaptations to derived, or apomorphic, characters, though of course not all apomorphic characters are adaptations.

the middle of the Pleistocene, the range of hominids had spread out into Europe and Asia. It is anybody's guess what physical environment is associated with the trend toward increased cranial capacity. The ecological information we do have is varied, indicating a wide range of environments, and little information concerning what ecological associations would be involved in the evolution of increased intelligence or the use of language. We are left considering traits that may be ancestral to *Homo sapiens*, or may be specialized, with an uncertain ecological connection. The *result* is clear enough. *Homo sapiens* is an accomplished generalist, adapting to varied environments. The *causes* are less clear. We are hardly an ideal case for a comparative analysis; the evidence provides little substance supporting any particular adaptive explanation.

Convergent evolution, resulting in a polyphyletic group displaying analogous features, is suggestive of natural selection, though even convergence can be the result of alternative mechanisms (see [Leroi *et al.*, 1994]). Under the best of circumstances, comparative studies may in fact not be able to disentangle the various features affecting evolution, including selection, drift, mutation and migration. Under the best of circumstances, it is a difficult task to disentangle one evolutionary explanation from another. Hominids and primates simply do not provide the ideal group for comparative analysis. We are a small group with few extant species, and the fossil record does not offer the sort of evidence we need. We know precious little of the specific social structures in early hominid societies which no doubt molded our cognitive capacities. We do not even know the appropriate taxonomic level, so that we could sensibly address the question. Moreover, much of what we are most interested in are unique to the hominid line. Much as we might care about our origins, much as we might want to understand the origins of sociality, intelligence, or language, the case we provide to evolutionary biology is not an easy one. The record is simply not sufficient for the questions posed. This may be as good as it gets.

6 CONCLUSION: EVOLUTIONARY PSYCHOLOGY AND EVOLUTION

The program of evolutionary psychology is ambitious. The goal is not just to defend the idea that humans evolved, but to offer evolutionary explanations of specifically human features in terms of natural selection. These explanations are intended to reflect evolutionary origins. Human language and cognition are certainly complex capacities. As with other capacities, we may reasonably inquire into their evolutionary origins. Evolutionary psychologists offer us explanations of psychological mechanisms in terms of the specific environmental demands which shaped them in our evolutionary past. An understanding of these mechanisms, and of their evolutionary functions, is intended in turn illuminate the functions of these psychological mechanisms. Evolutionary explanation thus motivates a reform of psychological practice.

These issues affect many features in our cognitive and emotional lives. Let's return to one paradigm. When we turn to language, the very features which

make it plausible to treat it as an adaptation — its uniqueness relative to other forms of communication — make it an especially poor subject for evolutionary analysis. Linguists recognize that the problem with treating language acquisition as a matter of learning — as some sort of social conditioning or generalization — is that the sort of evidence which is available to the child could not result in what we see in terms of the child's capacities. The human child develops language of remarkable complexity, with an ease and facility that is stunning. Developmental linguists emphasize the remarkable strides children make in learning language at an age at which other cognitive skills are not especially well developed, suggesting that language acquisition is mediated by distinctive mechanisms. Human language acquisition is so unlike what we think of as learning in other contexts that many linguists are reluctant to count the acquisition of language as *learning* at all. The uniqueness of what happens in language acquisition may be overstated, but not in comparison to the kinds of mechanisms simple models of learning can offer. We lack the kinds of incremental stages that would explain its acquisition in these terms. This much actually counts in favor of evolutionary psychology.

As Cosmides and Tooby like to emphasize, general purpose learning mechanisms are inadequate as models of human learning. Accounts of the evolutionary origins of language are in an analogous position, for many of the same reasons. The very lack of homology between human language and other forms of communication — even with the forms of communication which are present in other primates — leaves us unable to correlate it in any interesting way with any environmental variables, or to compare it with other forms of communication.[9] Social organization is surely involved *somehow* in the evolutionary origins of language. Our primate cousins also participate in complex social arrangements; nonetheless, there is no real analogue of human language in the rest of the extant primates. We do not know what forms of communication were present in ancestral hominid species, or even whether they had linguistic capacities, for the simple reason that the sort of information available in the fossil record is indecisive. So we do not know what the specific social structures were, or how they might have impacted the evolution of language. The uncertainty concerning the relevant taxonomic level for the emergence of language only exacerbates the problem. Thus, just as environmentally-driven accounts of language learning fail to explain our substantial accomplishments as individuals, so too environmentally-driven accounts of the evolution of language within the species fail to explain our accomplishments as a species. The same moral applies to human cognition, with the additional difficulty that the problems are even less well defined.

The mere fact that social coordination is facilitated by the use of human language or that it is promoted by human cognition provides no comfort; for though coordination and social interaction favors some form of communication, there is no basis for thinking that it would favor the sorts of complex features that are

[9]By contrast, there is no reason to think that the human brain includes structures wholly lacking in other primates, though there undoubtedly are modifications in architecture specific to humans.

integral to human language. Social coordination is surely facilitated by human reasoning, and that is often offered as grounds for thinking human cognition is an adaptation. Perhaps this is true. We have no evolutionary evidence bearing on the question directly. If there is any generalization concerning primate social organization, it is that it is highly variable. Gibbons are monogamous, orangutans are largely solitary, gorillas are polygamous, and chimpanzees exhibit a fluid and variable social structure. Baboons are highly social, though being matriarchal it is very different from human sociality. The sociality of hominids is certainly not at issue. The particular character of hominid social organization is difficult to discern from the mosaic offered among primates. This makes the extrapolation of ancestral hominid social organization very problematic. There is virtually nothing which would help in understanding the evolution of language. If we were seriously asked how social organization could favor, for example, the particular lexical and phrasal categories of human language, or how the rules underlying meaning and structure would contribute to enhanced fitness, it is hard to imagine how to answer the question. In any case, all the hominids, as well as primates, are social creatures; little is known concerning the specific social organization of ancestral hominids. As Deacon says,

> A full evolutionary account cannot stop with a formal description of what is missing or a scenario of how selection might have favored the evolution of innate grammatical knowledge. It must also provide a functional account of why its particular organization was favored, how incremental and partial versions were also functional, and how structures present in nonhuman brains were modified to provide this ability [1997, 38].

No doubt tool use, hunting, social coordination, and communication were significant features in the shaping of human capacities. Beyond such vague observations we are offered little that is concrete in explaining the specific features of human cognition or language use. If the goal is one of explaining the specific organizational patterns of human cognition, such an explanation is lacking in evolutionary psychology. Pinker, Tooby, and Cosmides emphasize the modular character of cognition; but no explanation of its modular structure is forthcoming. Pinker, Bloom and Dennett are fond of emphasizing the need for communication, and extol the virtues of social coordination; but this suggests no explanation whatsoever of the novel and unusual features characteristic of human language use, or of the remarkable features of human language.

We should not think the problems derive from a lack of evidence, so that with time the evidence will settle the issues. The record of hominid evolution assembled by anthropologists is remarkable for its quality. It leaves no serious doubt that humans are the result of evolution. Though that record sheds a tremendous amount of light on our ancestry, it does not favor the sort of speculations concerning the evolution of human intelligence and human language offered within evolutionary psychology. In the ideal cases for the comparative method there is a sufficiently

diverse taxonomic group in which we can see the same traits evolving in several lineages under similar environmental contingencies. It is not that simple increments in evidence will help. The problem is that the character of the evidence is of a sort that is not probative for the problems posed by evolutionary psychology.

Evolutionary psychology embraces an ambitious reform of psychology, based on evolutionary principles. It also embraces an ambitious evolutionary programme, based on the centrality of adaptation to evolutionary thinking. Together, these result in a proposed reform of psychology, based on evolutionary and adaptive psychological explanations. The programme of evolutionary psychology therefore depends on convincing adaptive explanations. There are three methods for confirming adaptive explanations prevalent in evolutionary biology: one depends on reverse engineering, another on dynamic models (or adaptive thinking), and the third on the comparative method. Each of these can be readily exhibited using well known and well regarded examples within evolutionary biology. When the claims of evolutionary psychologists are evaluated in terms of the standards which would be expected within evolutionary biology, the evidence offered for the adaptive explanations is deficient. Since the reform of psychology depends on independently defensible adaptive explanations, the credentials of evolutionary psychology as a reform program are yet to be established. The interesting question is whether we should expect the programme to succeed in the future. That depends on whether the sort of evidence that is so far lacking would be expected to surface in the future. I am skeptical.

BIBLIOGRAPHY

[Antonovics and Bradshaw, 1970] J. Antonovics and Bradshaw. Evolution in closely adjacent plant populations. vii. clinal patterns at a mine boundary. *Heredity*, 25:349–362, 1970.

[Antonovics *et al.*, 1971] J. Antonovics, A. D. Bradshaw, and R. G. Turner. Heavy metal tolerance in plants. *Advances in Ecological Research*, 7:1–85, 1971.

[Ardrey, 1966] R. Ardrey. *The Territorial Imperative*. New York: Aneneum, 1966.

[Barkow *et al.*, 1992] J. Barkow, L. Cosmides, and J. Tooby. *The Adapted Mind: Evolutionary psychology and the generation of culture*. NY: Oxford University Press, 1992.

[Baron-Cohen, 1995] S. Baron-Cohen. *Mindblindness: An Essay on Autism and Theory of Mind*. Cambridge: M.I.T. Press, 1995.

[Baum and Larson, 1991] D.A. Baum and A. Larson. Adaptation reviewed: A phylogenetic methodology for studying character macroevolution. *Systematic Zoology*, 40:1–18, 1991.

[Beatty, 1980] J. Beatty. Optimal-design models and the strategy of model building in evolutionary biology. *Philosophy of Science*, 47:532–561, 1980.

[Bechtel and Richardson, 1993] W. Bechtel and R.C. Richardson. *Discovering Complexity: Decomposition and Localization as Strategies in Scientific Research*. Princeton: Princeton University Press, 1993.

[Benton, 2000] T. Benton. Social causes and natural relations. In Rose and Rose, editors, *Cognition and Categorization*, pages 249–271. New York: Harmony Books, 2000.

[Brandon, 1978] R. Brandon. Adaptation and evolutionary theory. *Studies in History and Philosophy of Biology*, 9:181–206, 1978.

[Brandon, 1990] R. Brandon. *Adaptation and Environment*. Princeton: Princeton University Press, 1990.

[Brooks and McLennan, 1991] D. R. Brooks and D.A. McLennan. *Phylogeny, Ecology and Behavior: A Research Program in Comparative Biology*. Chicago: University of Chicago Press, 1991.

[Buller, 2005] D. J. Buller. *Adapting Minds: Evolutionary Psychology and the Persistent Quest for Human Nature*. Cambridge: MIT Univeristy Press, 2005.

[Burian, 1983] R. Burian. Adaptation. In M. Grene, editor, *Dimensions of Darwinism*, pages 287–314. Cambridge: Cambridge University Press, 1983.

[Buss *et al.*, 1996] D. M. Buss, R. J. Larsen, and D. Westen. Sex differences in jealousy: Not gone, not forgotten, and not easily explained by alternative hypotheses. *Psychological Science*, 7:373–375, 1996.

[Buss, 1994] D.M. Buss. *The Evolution of Desire: Strategies of Human Mating*. New York: Basic Books, 1994.

[Buss, 1995] D.M. Buss. Evolutionary psychology: A new paradigm for psychological science. *Psychological Inquiry*, 6:1–30, 1995.

[Buss, 1999] D.M. Buss. *Evolutionary Psychology: The New Science of Mind*. Boston: Allyn and Bacon, 1999.

[Buss, 2000] D.M. Buss. *The Dangerous Passion: Why Jealousy is as Necessary as Love and Sex*. New York: The Free Press, 2000.

[Chagnon, 1968] N. A. Chagnon. *Yanomamo: The Fierce People*. New Yourk: Holt, Rinehart and Winston, 1968.

[Chagnon, 1988] N. A. Chagnon. Male yanomamo manipulations of kinship classifications of female kin for reproductive advantage. In L. Betzig, M. Borgerhoff Mulder, and P. Turke, editors, *Human Reproductive Behavior: A Darwinian Perspective*, pages 23–48. New York: Cambridge University Press, 1988.

[Chamberlain and Wood, 1987] A. T. Chamberlain and B. A. Wood. Early hominid phylogeny. *Journal of Human Evolution*, 16(1):119–133, 1987.

[Coddington, 1988] J.A. Coddington. Cladistic tests of adaptationist hypotheses. *Cladistics*, 4:5–22, 1988.

[Cosmides and Tooby, 1987] L. Cosmides and J. Tooby. From evolution to behavior: Evolutionary psychology as the missing link. In J. Dupré, editor, *The Latest on the Best: Essays on Evolution and Optimality*, pages 277–306. Cambridge: M.I.T. Press, 1987.

[Cosmides and Tooby, 1992] L. Cosmides and J. Tooby. Cognitive adaptations for social exchange. In Barkow, Cosmides, and Tooby, editors, *The Adapted Mind*, pages 163–228. NY: Oxford University Press, 1992.

[Cosmides and Tooby, 1994] L. Cosmides and J. Tooby. Origins of domain specificity: The evolution of functional organization. In L.A. Hirschfeld and S.A. Gelman, editors, *Mapping the Mind: Domain Specificity in Cognition and Culture*. Cambridge: Cambridge University Press, 1994. Reprinted in: Cummins, R. and Cummins, D.D., editors, *Minds, Brains, and Computers*, Blackwell, pp. 523–543.

[Cosmides and Tooby, 1997] L. Cosmides and J. Tooby. The modular nature of human intelligence. In A. B. Scheibel and J.W. Schopf, editors, *The Origin and Evolution of Intelligence*, pages 71–101. Cambridge: Cambridge University Press, 1997.

[Daly *et al.*, 1982] M. Daly, M. Wilson, and S. J. Weghorst. Male sexual jealousy. *Ethology and Sociobiology*, 3:11–27, 1982.

[Darwin, 1859] C. Darwin. *On the Origin of Species*. Cambridge: Harvard University Press, 1859. A Facsimile of the First Edition, edited with an introduction by Ernst Mayr.

[Darwin, 1871] C. Darwin. *The Descent of Man, and Selection in Relation to Sex*. London: Murray, 1871. 2 volumes.

[Davies, 1996] P.S. Davies. Discovering the functional mesh: On the methods of evolutionary psychology. *Minds and Machines*, 6:559–585, 1996.

[Davies, 1999] P.S. Davies. The conflict of evolutionary psychology. In V.G. Hardcastle, editor, *Where Biology Meets Psychology*, pages 67–81. Cambridge: MIT Press, 1999.

[Dawkins, 1976] R. Dawkins. *The Selfish Gene*. Oxford: Oxford University Press, 1976.

[Dawkins, 1982] R. Dawkins. *The Extended Phenotype*. San Francisco: W.H. Freeman, 1982.

[Dawkins, 1986] R. Dawkins. *The Blind Watchmaker*. Oxford: Oxford University Press, 1986.

[Dawkins, 1997] R. Dawkins. *The Symbolic Species: The Co-evolution of Language and the Brain*. New York: W. W. Norton, 1997.

[Deacon, 1992a] T. W. Deacon. Biological aspects of language. In Jones, Martin, and Pilbeam, editors, *The Cambridge Encyclopedia of Human Evolution*, pages 128–133. Cambridge: Cambridge University Press, 1992.

[Deacon, 1992b] T. W. Deacon. The human brain. In Jones, Martin, and Pilbeam, editors, *The Cambridge Encyclopedia of Human Evolution*, pages 115–123. Cambridge: Cambridge University Press, 1992.

[Dennett, 1995] D. C. Dennett. *Darwin's Dangerous Idea: Evolution and the Meanings of Life*. New York: Simon and Schuster, 1995.

[Downes, 2001] S.M. Downes. Some recent developments in evolutionary approaches to the study of human cognition and behavior. *Biology and Philosophy*, 16:575–595, 2001.

[Dupré, 1987] J. Dupré, editor. *The Latest on the Best*. Cambridge: MIT Press, 1987.

[Felsenstein, 1982] J. Felsenstein. Numerical methods for inferring evolutionary trees. *Quarterly Review of Biology*, 57:397–404, 1982.

[Fodor, 1985] J. Fodor. *Numerical Methods for Inferring Evolutionary Trees*. Cambridge: MIT Press/Bradford Books, 1985.

[Ghiselin, 1983] M. Ghiselin. Lloyd morgan's canon in evolutionary context. *Behavioral and Brain Sciences*, 6:362–363, 1983. on D.C. Dennett.

[Gigerenzer, 1998] G. Gigerenzer. Ecological intelligence: An adaptation for frequencies. In Cummins and Allen, editors, *The evolution of mind*, pages 9–29. New York: Oxford University Press, 1998.

[Gould and Lewontin, 1979] S.J. Gould and R.C. Lewontin. The spandrels of san marco and the panglossian paradigm: A critique of the adaptationist programme. *Proceedings of the Royal Society of London*, B205:581–598, 1979.

[Gould, 1966] S. J. Gould. Allometry and size in ontogeny and phylogeny. *Biological Reviews*, 41:587–640, 1966.

[Gould, 1977] S. J. Gould. *Ontogeny and Phylogeny*. Cambridge: Harvard University Press, 1977.

[Grantham and S., 1999] T. Grantham and Nichols S. Evolutionary psychology: Ultimate explanations and panglossian preditions. In V.G. Hardcastle, editor, *Where Biology Meets Psychology*, pages 47–66. Cambridge: MIT Press, 1999.

[Griffiths, 1996] P. Griffiths. The historical turn in the study of adaptation. *British Journal for the Philosophy of Science*, 47:511–532, 1996.

[Hardcastle, 1999] V. G. Hardcastle, editor. *Where Biology Meets Psychology: Philosophical Essays*. Cambridge: MIT Press, 1999.

[Harvey and Pagel, 1991] P.H. Harvey and M.D. Pagel. *The Comparative Method in Evolutionary Biology*. Oxford: Oxford University Press, 1991.

[Hennig, 1966] W. Hennig. *Phylogenetic Systematics*. Urbana: University of Illinois Press, 1966.

[Holling, 1964] P.S. Holling. The analysis of complex population processes. *Canadian Entomologist*, 96:345–347, 1964.

[Jain and Bradshaw, 1966] S.K. Jain and A.D. Bradshaw. Evolutionary divergence among adjacent plant populations. i. the evidence and its theoretical analysis. *Heredity*, 21:407–441, 1966.

[Jones *et al.*, 1992] S. Jones, R. Martin, and D. Pilbeam, editors. *The Cambridge Encyclopedia of Human Evolution*. Cambridge: Cambridge University Press, 1992.

[Ketelaar and Ellis, 2000] T. Ketelaar and B. J. Ellis. Are evolutionary explanations unfalsifiable? evolutionary psychology and the lakatosian philosophy of science. *Psychological Inquiry*, 11:1–21, 2000.

[Kitcher, 1982] P. Kitcher, editor. *Abusing Science: The Case Against Creationism*. Cambridge: MIT Press, 1982.

[Kitcher, 1985] P. Kitcher, editor. *Vaulting Ambition: Sociobiology and the Quest for Human Nature*. Cambridge: MIT Press, 1985.

[Larson and Losos, 1996] A. Larson and J. B. Losos. Phylogenetic systematics of adaptation. In Rose and Lauder, editors, *Adaptation*. San Diego: Academic Press, 1996.

[Lauder, 1996] G.V. Lauder. The argument from design. In Rose and Lauder, editors, *Adaptation*, pages 55–91. San Diego: Academic Press, 1996.

[Leroi *et al.*, 1994] A. M. Leroi, M. R. Rose, and Lauder G. V. What does the comparative method reveal about adaptation? *American Naturalist*, 143:381–402, 1994.

[Lewontin *et al.*, 1984] R. C. Lewontin, Rose S., and L. Kamin. *Not in Our Genes?* New York: Pantheon, 1984.

[Lieberman, 1992] P. Lieberman. Human speech and language. In Jones, Martin, and Pilbeam, editors, *The Cambridge Encyclopedia of Human Evolution*, pages 134–137. Cambridge: Cambridge University Press, 1992.

[Lorenz, 1966] K. Lorenz. *On Agression*. New York: Harcourt, Brace, Jovanovich, 1966.
[Losos, 1998] J.G. Losos. Contingency and determinism in replicated adaptive radiations of island lizards. *Science*, 279:2115–2118, 1998.
[Lumsden and Wilson, 1981] C. Lumsden and E. O. Wilson. *Genes, Minds and Culture*. Cambridge: Harvard University Press, 1981.
[MacNair, 1979] M. MacNair. The genetics of copper tolerance in the yellow monkey flower, *mimulus guttatus*. i. crosses to nontolerants. *Genetics*, 91:553–563, 1979.
[Maddison *et al.*, 1984] W. P. Maddison, M. J. Donaghue, and D. R. Maddison. Outgroup analysis and parsimony. *Systematic Zoology*, 33:83–103, 1984.
[Morris, 1967] D. Morris. *The Naked Ape*. New York: McGraw-Hill, 1967.
[Nesse, 1990] R. M. Nesse. Evolutionary explanations of emotions. *Human Nature*, 1:261–289, 1990.
[Nozick, 1993] R. Nozick. *The Nature of Rationality*. Princeton: Princeton University Press, 1993.
[Nozick, 1994] R. Nozick. *The Language Instinct: How the Mind Creates Language*. New York: William Morrow & Co, 1994.
[Nozick, 1997] R. Nozick. *How the Mind Works*. Cambridge: MIT Press, 1997.
[Orzack and E., 1994a] S. Orzack and Sober E. How (not) to test an optimality model. *Trends in Ecology and Evolution*, 9:265–267, 1994.
[Orzack and E., 1994b] S. Orzack and Sober E. How to formulate and test adaptationism. *American Naturalist*, 148:202–210, 1994.
[Orzack and E., 1994c] S. Orzack and Sober E. Optimal models and adaptationism. *The American Naturalist*, 143:361–380, 1994.
[Pinker and Bloom, 1990] S. Pinker and P. Bloom. Natural language and natural selection. *Behavioral and Brain Sciences*, 13:707–784, 1990.
[Pinker and Bloom, 1992] S. Pinker and P. Bloom. Natural language and natural selection. In Barkow, Cosmides, and Tooby, editors, *The Adapted Mind*, pages 451–493. NY: Oxford University Press, 1992.
[Profet, 1992] M. Profet. Pregnancy sickness as an adaptation: A deterrent to maternal ingestion of teratogens. In Barkow, Cosmides, and Tooby, editors, *The Adapted Mind*. NY: Oxford University Press, 1992.
[Richardson and Burian, 1992] R. C Richardson and R. M. Burian. A defense of propensity interpretations of fitness. In M. Forbes and D. Hull, editors, *PSA 1992 Vol. 1*, pages 349–62. East Lansing: Philosophy of Science Association, 1992.
[Richardson, 1996] R. C. Richardson. Critical review of robert brandon *adaptation and environment*. *Philosophy of Science*, 63:122–136, 1996.
[Richardson, 2000] R. C. Richardson. The prospects for an evolutionary psychology: Human language and human reasoning. *Minds and Machines*, 6:541–557, 2000.
[Richardson, 2001] R. C. Richardson. Conceptual challenges in evolutionary psychology: Innovative research strategies. In H. R. Holcomb, editor, *The Adapted Mind*, pages 327–373. Kluwer Academic Publishers, 2001.
[Richardson, 2003a] R. C. Richardson. Engineering design and adaptation. *Philosophy of Science*, 70:1–12, 2003.
[Richardson, 2003b] R. C. Richardson. Epicycles and explanations in evolutionary psychology. *Psychological Inquiry*, 11:46–49, 2003.
[Rose and Lauder, 1996] H. R. Rose and G. V. Lauder, editors. *Adaptation*. New York: Academic Press, 1996.
[Rose and Rose, 2000] H. R. Rose and S. Rose, editors. *Alas, Poor Darwin: Arguments Against Evolutionary Psychology*. New York: Harmony Books, 2000.
[Ruse, 1979] M. Ruse. *Sociobiology: Sense or Nonsense?* Dordrecht: D. Reidel, 1979.
[Sahlins, 1976] M. Sahlins. *The Use and Abuse of Biology*. Ann Arbor: University of Michigan., 1976.
[Sterelny and Griffiths, 1979a] K. Sterelny and P. E. Griffiths. *The Evolution of Human Sexuality*. Cambridge: Harvard University Press, 1979.
[Sterelny and Griffiths, 1979b] K. Sterelny and P. E. Griffiths. *On the Use and Misuse of Darwinism in the Study of Human Behavior*. Cambridge: Harvard University Press, 1979.
[Sterelny and Griffiths, 1999] K. Sterelny and P. E. Griffiths. *Sex and Death: An Introduction to Philosophy of Biology*. Chicago: Chicago University Press, 1999.

[Symons, 1992] D. Symons. On the use and misuse of darwinism in the study of human behavior. In Barkow, Cosmides, and Tooby, editors, *The Adapted Mind*, pages 137–159. NY: Oxford University Press, 1992.

[Tooby and Cosmides, 1992] J. Tooby and L. Cosmides. The psychological foundations of culture. In Barkow, Cosmides, and Tooby, editors, *The Adapted Mind*, pages 19–136. NY: Oxford University Press, 1992.

[Tooby and de Vore, 1987] J. Tooby and I. de Vore. The reconstruction of hominid evolution through strategic modeling. In W. G. Kinzey, editor, *The Evolution of Human Behavior: Primate Models*, pages 19–136. Albany: SUNY Press, 1987.

[Uttal, 2001] W. R. Uttal. *The New Phrenology: The Limits of Localizing Cognitive Processes in the Brain*. Cambridge: MIT. Press., 2001.

[van den Berghe, 1980] P. van den Berghe. Incest and exogamy: A sociobiological reconsideration. *The Behavioral and Brain Sciences*, 5:27, 1980.

[van den Berghe, 1983] P. van den Berghe. Human inbreeding avoidance: Culture in nature. *The Behavioral and Brain Sciences*, 6:91–123, 1983.

[Vogel, 1988] S. Vogel. *Life's Devices: The Physical World of Animals and Plants*. Princeton: Princeton University Press, 1988.

[Vogel, 1998] S. Vogel. *Cat's Paws and Catapults: Mechanical Worlds of Nature and People*. New York: W.W. Norton, 1998.

[Williams, 1966] G. C. Williams. *Adaptation and Natural Selection*. Princeton: Princeton University Press, 1966.

[Williams, 1992] G. C. Williams. *Natural Selection: Domains, Levels and Challenges*. Princeton: Princeton University Press, 1992.

[Wilson, 1975] E. O. Wilson. *Sociobiology: the New Synthesis*. Cambridge: Harvard University Press, 1975.

[Wilson, 1976] E. O. Wilson. Academic vigilantism and the political significance of sociobiology. *The New York Review of Books*, 1976. reprinted in Caplan, 1978, pp. 291–303.

[Wilson, 1978] E. O. Wilson. *On Human Nature*. Cambridge: Harvard University Press, 1978. reprinted in Caplan, 1978, pp. 291–303.

SITUATED COGNITION

Miriam Solomon

1 INTRODUCTION

Cognition is always situated. It is always concretely instantiated in one way or another. There are no disembodied cognitive achievements. The situated cognition literature details the ways in which cognition can be instantiated (or situated), and, instead of abstracting what is in common to all cognition — the traditional philosophical project —explores the epistemic significance of particular routes to cognitive accomplishment.

Over the last twenty years, phenomena of 'situated cognition' have been described in several disciplines. Cognitive scientists such as Lave (1988), Suchman (1987), Hutchins (1995), Norman (1988), Kirsh (1991, 1995a, 1995b), Brooks (1991), Agre (1997), Lakoff and Johnson (1999), Barwise and Perry (1985) have described ways in which representation of the world, learning, memory, planning, action and linguistic meaning are embedded in the environment, tools, social arrangements and configurations of the human body. Andy Clark (1997, 2001) and William Clancey (1997) have given excellent overviews of this work. Feminist epistemologists and philosophers of science such as Harding (1991), Haraway (1989, 1991), Keller (1985) and Longino (1990, 1993) have argued that ideologies that inform scientific creativity and scientific decision-making derive, in large part, from social variables including family psychodynamics, political orientation and societal position. Science studies researchers such as Latour (1987), Galison (1997), Hacking (1983), Knorr-Cetina (1999) and Pickering (1995) have argued that scientists have situated knowledge practices that are constituted around local experimental successes and are dependent on particular tools, domains, historical contexts and forms of social organization.

Situated cognition approaches have in common a *rejection* of the ideas that cognition is individualistic (accomplished by each human individually), general (true of all individual humans and applicable in all situations), abstract, symbolic, explicit, language based and located in the brain as mediator between sensory input and action output. These ideas were dominant in twentieth century thinking in cognitive psychology, artificial intelligence and analytic epistemology, and still have more than their share of advocates. Situated cognition marks a departure from one or more of these traditional assumptions, and acceptance of the view that cognition can be social, particular, concrete, implicit, non-linguistic and distributed. The phrase 'situated cognition' marks the recognition that the

Handbook of the Philosophy of Science. Philosophy of Psychology and Cognitive Science
Volume editor: Paul Thagard
General editors: Dov M. Gabbay, Paul Thagard and John Woods
© 2007 Elsevier B.V. All rights reserved

approaches of traditional cognitive psychology, traditional epistemology and traditional AI are inadequate. This does not mean that these traditional approaches should be discarded completely; rather that they should be either supplemented or reframed in most cases.

The phrase 'situated cognition' thus has the most purchase on those who are addressing the deficiencies of these assumptions about cognition. It is worth mentioning that outside the fields of cognitive psychology, artificial intelligence and analytic philosophy, there are a number of intellectual traditions that have been fruitful sources of ideas about situatedness. Philosophy in the Continental tradition (especially the work of Heidegger and Merleau-Ponty), philosophy in the tradition of American Pragmatism (especially Dewey and Mead) and the psychological ideas of Lev Vygotsky and J.J. Gibson are each important theoretical resources. In these intellectual traditions, 'situated cognition' was with us throughout the twentieth century; however, it was 'discovered' in cognitive psychology, artificial intelligence and analytic philosophy in the 1980s.

The term 'situated' was used by the American Pragmatist George Herbert Mead, and by the sociologist C.W. Mills in the 1930s and 1940s (see Clancey 1997, p. 23). It was not until the 1980s that it entered wide usage, in the work of Barwise and Perry, Suchman, Lave, Haraway, Hutchins and others (interestingly, these users of the term were all in the San Francisco Bay area). The work of Barwise and Perry may have provided useful terminology — although not a broad enough theoretical framework-to cover a variety of epistemic phenomena that have been more specifically called 'positional,' 'local,' 'standpoint,' 'centered' and 'embodied'.[1]

One purpose of this paper is to provide a structured overview of work on situated cognition. This overview will, it is hoped, facilitate interdisciplinary cross-fertilization of ideas. The main fields in which situated cognition is studied-cognitive science, feminist epistemology and science studies-are unnecessarily isolated from one another.

A second purpose is to exhibit to naturalistic epistemologists and philosophers of mind the full range of phenomena of situated cognition. While a few naturalistic epistemologists (e.g. Goldman (1999, 2002), Kitcher (1993), Solomon (2001), Thagard (1993)) have worked with the ideas of distributed cognition, and some (e.g. Giere (1988, 1999), Nersessian (1992)) write about non-linguistic representation, for the most part situated cognition has not been incorporated into either descriptive or normative discussions. Normative concepts like 'justified' and 'knows' remain largely individualistic, linguistic, general and explicit. It is time for analytic philosophy to catch up with the rest of cognitive science.

[1] In July 31 and August 1 1996 e-mail communications with Ed Hutchins, he wrote, 'I don't really know exactly where the term 'situated' came from in my work. I think I would say 'The San Francisco Bay area'....My work has always focused on situated cognition, even long before I heard the word'. '...I think Barwise and Perry influenced all of us...I remember being frustrated by their theory of semantics, but I'm sure their work encouraged my thinking about the notion of situation.'

2 EXAMPLES

Here are some specific ways in which cognitive phenomena are situated. I divide the examples into seven categories: environment, goals, social distribution, political location, tools (including language, knowledge artifacts and complementary strategies), historical context and individual embodiments. This has the desirable effect of mingling examples from different academic fields. Note that *particular cognitive practices are typically situated in more than one way.*

At this stage of the paper, I am treating together a range of cognitive phenomena: learning, discursive knowledge, language, know-how (practice), planning, memory, teaching, creative strategies and representation. I am also treating descriptive and normative epistemology together. Necessary distinctions will be made later.

2.1 Environment

Many human activities are *situationally determined* rather than controlled by internal representation and conceptualization. For example, walking, juggling, tying shoelaces are all accomplished without following an internal script (see e.g. Kirsh 1991 and Brooks 1991). Some roboticists (e.g. Brooks) are designing their machines to work in a particular environment without having an internal representation or description of that environment. Even where there is some internal representation or prior planning, there are often situational determinants of the cognitive activity. Situation semantics (Barwise and Perry 1985) captures the dependence of linguistic interpretation on knowledge of the specific circumstances of the speaker and the utterance, especially including the local environment. Some richer examples are canoeing through rapids (Suchman 1987, p. 52), midwifery (Dalmiya and Alcoff 1993), flying airplanes and cooking. Intentional behavior during these activities is shaped by both environment and inner cognitive processes. And, for example, Evelyn Fox Keller describes Barbara McClintock's method of 'listen to the material' and 'let the experiment tell you what to do' in genetics (Keller 1985, p. 162). McClintock criticized other geneticists for over-planning their experiments and prejudging the acceptable range of scientific conclusions. One computer scientist (Resnick, 1994) reminds us of the role of the environment with the aphorism, 'The hills are alive'.

Knowledge is also, often, *domain specific.* For example, Hutchins (1995) describes Micronesian methods of navigation, which, in their domain (the islands of Micronesia) are as successful as Western navigational methods. Micronesian methods accomplish the cognitive task by supposing that the ship is stationary and that islands–both real and imaginary–are passing by. In oceans where islands are further apart, Micronesian methods fail. Jean Lave (1988), in her work on the Adult Math Project, has argued for the domain specificity of everyday mathematical knowledge: subjects can excel at mathematical knowledge in particular domains (e.g. supermarket shopping) without excelling at abstract mathematical

tasks, or mathematical tasks in unfamiliar domains. She writes (1988, p. 3), 'The same people differ in their arithmetic activities in different settings in ways that challenge theoretical boundaries between activity and its settings'. Denise Cummins (1996) has discovered that children first develop competence at conditional reasoning in domains involving permission and obligation schemas; only later do they develop competence in other real-world domains, and even later they may develop competence at abstract conditional reasoning. Finally, scientific methods can vary in successfulness depending on the domain in which they are employed. For example, Bechtel and Richardson (1993) argue that decomposition and localization, a two-stage heuristic for building models of complex biological phenomena, will have different success rates in the various domains in which it is used.

2.2 Goals

Reasoning practices are typically governed by epistemic, pragmatic, or other goals. Some goals may be obvious; other goals of reasoning may take investigation or insight to discern.

Analytic epistemologists focus on the epistemic goal of truth;[2] the examples discussed here show that there is a wide range of goals in reasoning practices, and that truth is not (and should not be) the only normative goal, nor does it underlie (in any subtle or approximate way[3]) all normative goals.

Pragmatic goals are often the easiest to discern. For example, an important goal of Micronesian navigation is to navigate successfully among the islands of Micronesia (truth of the theory used to do this is not necessary). An important goal of building bridges is to build a sturdy structure that will withstand wind and other natural forces (truth of Newtonian mechanics not necessary).

Reasoning can have 'unofficial' goals. For example, Ziva Kunda (1990) has shown that people frequently perform 'motivated reasoning', in which the preference for particular outcomes skews decisions. Evidence still plays a role, and motivated reasoning rarely overcomes extensive contrary evidence, but it skews many decisions (typically, contrary evidence is not extensive). The desire to continue a coffee drinking habit, for instance, leads coffee drinking subjects to undervalue research suggesting the negative health benefits of coffee (Kunda 1990). It is preferable for the dedicated coffee drinker to come to the conclusion that pleasurable coffee drinking habits do not need to be changed. In such cases, the goal of reasoning is not only truth (or some other epistemic or pragmatic goal), but also the desire to avoid coming to a difficult, unpleasant or dissonant conclusion. In general, humans prefer conclusions that cohere with, or support, prior beliefs. This preference extends to a preference for theories, including scientific theories,

[2]Stephen Stich (1990) is a notable exception. Stich argues that truth is both an unattainable and an undesirable goal, and gives some programmatic suggestions about other epistemic goals. He has not taken these ideas further, or tried to incorporate the sort of data I am covering in these examples.

[3]For example, Philip Kitcher (1993) writes about *significant* truth; and Ronald Giere (1988) writes about theories that accurately *model* (correspond to, represent) the world.

which cohere with prior training, ideology and values. Examples of such cases have been documented by Kuhn (1962), Longino (1990) and many others.

Motivated reasoning is not always conservative. Sometimes humans can be motivated to change their beliefs and practices. For example, the influence of peer pressure often makes it less uncomfortable to change than to stay the same. By the mid-1960s, peer pressure was a major force converting geologists to drift theories. Or, for example, one can be motivated to embrace an idea or theory that challenges authority, as was Galileo[4] Some theories have had appeal because they support and further political agendas–for example, racist and sexist views of human nature based on morphological studies with weak data had a large following in the late nineteenth and early twentieth century (Stepan 1986). Donna Haraway (1989) has shown that more recent work in primatology has such a political agenda. Science studies, and especially work on gender and science, is replete with examples of motivated reasoning in science. Religion, politics and fashion are other domains with similar patterns of motivated reasoning.

Some reasoning may have biological or evolutionary goals: Denise Cummins has suggested that deontic reasoning (reasoning involving permissions and obligations) has particular significance because it is reasoning that is involved in the maintenance of the hierarchical social organizations that constitute primitive human societies (Cummins 1996). (If Cummins is right, deontic reasoning is situated in two ways: environment and goals.) Evolutionary psychologists (the successors to sociobiologists), in general, seek evolutionary goals for behavior, including reasoning behavior. The evolutionary context can also be viewed as a ('deep history') form of historical situatedness (see category 6 below).

Philosophers have traditionally taken on the task of normative evaluation of reasoning. When reasoning is evaluated *instrumentally*,[5] it can be evaluated for attainment of *actual* goals, for a *subset* of actual goals, or for some other goals that are of interest. Some of these evaluations will be of more practical interest than others. (More normative remarks will come later.)

2.3 Social distribution

Inquiry can be socially distributed in several ways. Members of a community can have different areas of expertise, each of which contribute to a joint project. Hutchins (1995) describes such a case for navigation on board a naval vessel, where sailors involved in a 'fix cycle' each have different roles and different expertise: some have the skill to use particular instruments, others transmit and compile data. Many creative projects–from scientific research to advertising to sports coaching–benefit from the collaboration of a variety of experts. Hutchins remarks that

[4]Frank Sulloway (1996) argues that younger siblings are more likely to challenge authority than older siblings. This phenomenon also belongs (under various descriptions) in the categories of social distribution, social location or individual embodiments.

[5]This is not the only way to evaluate reasoning. Reasoning can also be evaluated intrinsically, in terms of its correspondence to some preferred form (e.g. syllogistic, Bayesian, deductive validity etc).

'overlapping expertise', where no single individual is indispensable, is the most robust form of social organization.

Inquiry can also be socially distributed when members of a community adopt a deliberate or fortuitous 'divide and conquer' strategy. Individuals, or subgroups of a community, can pursue *different* projects. This increases the probability of overall success for the community by increasing the search space.[6] Kitcher (1993) and Hull (1988) have argued that individual incentives such as the desire for fame and success can adequately distribute cognitive efforts among viable competing scientific theories. Longino (1990) and Keller (1985) dispute this; they claim that research decisions are influenced by ideology and values, and that these do not typically result in an adequate distribution of cognitive effort. According to them, adequate distribution of cognitive effort is expected only in scientific communities which are especially democratic. I will say more about this in the next section.

Finally, inquiry can be socially distributed for decision making. This may be because decisions are made in a group; Hutchins models such a case (1995, Chap 4), and shows how groups can sometimes make better navigational decisions than individuals, and sometimes worse. Or it may be epistemologically appropriate to consider aggregate decisions even when the members of that aggregate do not come to joint decisions. I have argued (Solomon 2001) that a scientific consensus can be rational, even though each scientist has a different reason or cause for joining the consensus. Consensus on plate tectonics theory during the 1960s is such a case. In my view, rationality can be a social accomplishment rather than an achievement of individual minds.

Distributed artificial intelligence (DAI) can model each kind of social distribution of inquiry, by using several processors that are networked in parallel with one another. The connections between processors model the relations (e.g. of transmission of information, peer pressure and authority) between individuals. Thagard (1993) has begun such a model for scientific decision making.

2.4 Political Location

Haraway (1991, p. 188) has written, 'feminist objectivity means quite simply *situated knowledges*'. Situated knowledge, for Haraway, is embodied and politically located knowledge. Individuals each have 'partial perspective', and objectivity is reached through acknowledging this fully, and adopting a critical perspective. Haraway, along with other feminist epistemologists who make use of Marxist and socialist-feminist ideas (starting with Hartsock 1983, Smith 1990, Harding 1986), especially trusts the standpoints of those individuals who are both politically subjugated and politically aware (without false consciousness). Knowledge from such a political standpoint is more likely to be inclusive and critical, since the politically subjugated can come to understand their oppressors as well as themselves.

[6] It is usually supposed that scientists are aiming to eliminate all but one successful theory. This supposition has been challenged, notably by Longino (1990, 1993) and by Solomon (2001) who argue for pluralism. This reinforces the claim that knowledge is socially distributed.

As Haraway puts it, the view is best 'from below' and the vision is best when partial views are joined into a 'collective subject position' (p. 196). Haraway is not only writing about gender ideology, but also about differences due to other kinds of socio-political positioning, such as class, race, disability, ethnicity and age.

Haraway introduced 'situated knowledge' terminology to feminist philosophy of science, and is the only one to favor this terminology over the pervasive 'standpoint epistemology' terminology. Haraway means to include not only the situatedness that comes from political position, but also other forms of situatedness (with respect to tools, environment and individual embodiments, in particular). Situatedness is described from a point of view that is influenced by postmodernism and poststructuralism. Others (especially Harding, but also Hartsock, Jagger and Smith) describe their positions as 'standpoint epistemologies': they make similar claims about political situatedness, with the same origins in Marxist and socialist-feminist ideas. Their resistance to 'situated knowledge' terminology marks only their distance from Haraway's postmodern/poststructuralist style; their understanding is more material (in both the Marxist and the everyday senses of the word 'material') than it is textual.[7] I regard them all as making 'situated cognition' claims in the sense that is relevant for this paper: all claim that political standpoint is epistemologically relevant.

Not all feminist epistemologists claim that the view is better from below; some claim that it is just *different*, and advocate pluralism. Keller (1983) and Longino (1990), for example, show that different values and ideology produce different — and equally valuable[8] — scientific ideas. They claim that science is done best when the scientific community has a variety of background assumptions, and that this typically happens when the scientific community is socially diverse.[9] (This might be called dual or multiple standpoint theory, depending on how many standpoints are possible.)

It is not enough for an individual mainstream scientist to *try* to be objective and critical.[10] It is like a physician trying (and claiming) to remain neutral in her prescribing practices while receiving gifts from pharmaceutical companies: it fails too often to be trustworthy. Efforts at pure thought are nice, but not sufficiently effective. Actual political locatedness of scientists in subjugated positions is empirically necessary for improving science.

[7] In e-mail correspondence on July 9, 1996, with Sandra Harding, Harding writes that she uses both 'situated' and 'standpoint' terminology, but prefers 'standpoint' terminology since it is less likely to be dismissed by those (including many feminist philosophers) who are skeptical of poststructuralist/postmodern approaches. (Haraway, in all fairness, also distances herself from pure poststructuralism/postmodernism.)

[8] Equally valuable for science, that is (they may not be morally or politically on a par). This is my, admittedly controversial, reading of Longino and Keller. I take it that their support of ideas that have more 'feminist' ideology in the background is an attempt to redress the balance.

[9] 'Diversity' can be realized in different ways. Keller is insistently anti-essentialist, claiming that what matters is that scientists have different ideologies (not that they are men or women). Ideologies, Keller claims, are profoundly influenced by childhood experiences which are responsible for gender identity (she uses a psychodynamic account).

[10] Contra Louise Antony's suggestions in (1993) and (1995).

2.5 Tools

A large group of objects and phenomena fall into this category. I define tools broadly: as artifacts,[11] or material constructs,[12] that aid humans with various goals, including knowledge goals. Tools typically (although not always, especially for simple tools) store information. Sometimes this is information about how to use the tool itself. Sometimes it is information about the world stored implicitly or explicitly. Some tools also produce information.

Donald Norman, in *The Design of Everyday Things* (1989) (formerly, and more aptly, titled *The Psychology of Everyday Things* (1988)) discusses the difficulty of designing tools well enough that we know, easily enough, how to use them. Typically, it takes five or six full attempts to design everyday objects such as stovetops, clock radios and automobile controls. The better the design, the more we would say (in computer parlance) that the tool is user-friendly. A tool is user-friendly when the information about how to use it is obvious from looking at it or from using it.[13] For example, when stove controls are arranged in the same pattern as the burners (rather than, e.g. in a straight line) there is no confusion about which control goes with which burner. When a door handle can only be pushed, pulled or turned in the appropriate direction, there is no confusion about what to do to open or close the door. When an arrow together with an unambiguous placement of the word 'hot' or 'cold' is engraved on a shower control, there is less danger of being singed. Humans need to bring their own general knowledge–of arrows, spatial representations and words–to the tool, but much of the information required to use the tool is in the tool, not the head (or even the instruction manual). Norman goes so far as to say 'much of our everyday knowledge resides in the world, not in the head' (1989, p. xiii). 'Obviousness' of usage of a tool is of course relative to culture and prior knowledge, as anyone who has tried to take a shower in a foreign country knows well.

Tools can store information about the world. Some information storage is fairly explicit: for example, the information contained in printed materials (journals, newspapers, books) and on the World Wide Web.[14] One has to know how to read, and how to access and sort these materials (not always trivial), and then the information is fully accessible. Some information storage is much more implicit. For example, David Kirsh (1995a) writes about 'complementary strategies' in which subjects actively rearrange the world (creating an artifactual arrangement) so as to store information necessary for carrying out a task. Everyday cases of this are: arranging the ingredients for a recipe on the countertop in order for use (an

[11] Not in the pejorative sense of experimental artifacts (undesirable effects created by the experimenter).

[12] This definition leaves it unclear whether or not language is a tool. I think that it is, but this matter does not need to be addressed in this paper.

[13] Tools have 'affordances': ways in which they may be used. A tool is 'user-friendly' when its operation takes advantage of its affordances in an unambiguous way (of course, relative to culture and prior knowledge).

[14] These are tools, on my definition of the word. I have found no definition of 'tool' that excludes knowledge artifacts (printed materials, the World Wide Web) in a non-arbitrary way.

assembly line), arranging groceries at a supermarket checkout to pack them intelligently (those packed last are pushed farthest away from the bagging station) and arranging piles of papers on one's desk in a manner that classifies projects and topics (if someone 'tidies up' the desk, information is lost).[15] As Kirsh puts it, we modify the world so that it is external memory for us. We 'recruit external elements to reduce internal loads'.

Some tools produce information. Hutchins (1995) shows that the cognitive abilities of a tool plus a human differ from the cognitive abilities of a human alone. He discusses a variety of tools: alidade (used to measure the angle between the ship's longitudinal axis and to some target), gyrocompass, radar, maps etc., showing how a tool can both produce data and 'black box' necessary calculations, thereby enhancing cognitive abilities while reducing individual cognitive loads. We are especially familiar with this use of tools in scientific research, where experimental apparatus (the tools) can produce data or desirable technology. The new attention given to experimentation in science studies research has shown how experimental knowledge can often be more robust than theoretical interpretations.[16] Donna Haraway's 'cyborg manifesto' (1991, Chapter 8) is a useful rhetorical ploy which captures the idea that a hybrid of organism and machine is an epistemic unit.

2.6 Historical context

Knowledge is often relative to temporal, historical context. Sometimes this is because conventions have already been set for a practice. For example, it is a convention that the faucet on the left is hot and the faucet on the right is cold–at least, in the United States. (Some of historical context is due to cultural context.) It is a convention that the y axis is the dependent variable and the x axis the independent variable. Conventions are regarded as arbitrary: if they had been otherwise, it would not make a significant different to the practice.

Often the relativity to historical context pervades a practice more thoroughly. Hutchins (1995) writes about historical contingencies in the development of tools for navigation. He points out (p. 113), for example, that the modern plastic star finder is not a contemporary invention, but a descendent of the astrolabe. Also (p. 106), Hutchins shows that advances in navigation were dependent on advances in digital calculation, particularly on the development of logarithms in the early seventeenth century. The data that came from the chip log, astrolabe, quadrant, sextant, and instruments to measure bearing and depth could be computed with logarithms or a slide rule. Relativity to historical context includes relativity to the state of other knowledge practices at the time.

[15] An well-known example here is the case that earned the slogan 'the intelligence is in the cottage cheese'. Someone is asked to measure 3/4 of 1/2 a cup of cottage cheese. Instead of calculating it by doing the fractions and using a measuring cup, they can fill a cup, remove half, smoosh it into a circle, and cut away a 1/4 of the circle. I heard this case from David Kirsh.

[16] See, for example, the work of Ian Hacking (1983), Bruno Latour (1987), Peter Galison (1997) and Andy Pickering (1995).

Relativity to historical context includes the influence of historical contingencies other than related knowledge practices. Often, for example, religious climate, political convictions or other aspects of culture affect scientific work, as recent historians of science, from Kuhn on, have documented. This kind of situatedness overlaps with the discussion of political standpoint, above.

Relativity to historical context can be interpreted so broadly as to include evolutionary context (really 'pre-history'). In this controversial area of cognitive psychology (Leda Cosmides and John Tooby are particularly prominent proponents: see Barkow, Cosmides and Tooby 1995), it is argued that human cognitive processes are adaptations to ancient selective pressures.

2.7 Individual embodiments

Human reasoning has particular characteristics: it is reasoning that is situated in bodies, that involves particular strategies and heuristics, depends on prior knowledge and training and is shaped by individual motivations.

The brain (and especially the parts that may do symbolic processing) is not the sole locus of cognition, even in individual human beings. Kirsh (1995b) writes about 'why we use our hands when we think' showing that, for example, hands can be used to point in order to fix attention, and hands can be used to move objects (e.g. right side up) so as to perceive them better. Hutchins (1995, p. 363) complains that 'the hands, the eyes, the ears, the nose, the mouth, and the emotions all fell away when the brain was replaced by a computer'. He insists on a reintegration, through acknowledging the material (rather than the symbolic) aspects of cognition. Earlier in the book (Chapter 3), he shows how processes such as hand-eye co-ordination are necessary for the cognition that goes into navigation. The body as well as the brain is involved in cognition. Marcel Kinsbourne, in his discussions of anasognosia,[17] has suggested that bodily awareness is the basis for general spatial awareness.

Much human reasoning involves the use of heuristics and strategies that are particular to humans, and do not correspond to abstract general accounts of reasoning such as Bayesian logic, confirmation theory or truth functional logic. Tversky, Kahneman and others have described heuristics of salience, availability and representativeness (e.g. Kahneman, Slovic and Tversky 1982 and Nisbett and Ross 1980). Holyoak and Thagard (1995) say more about the representativeness heuristic, showing how both representativeness and analogical reasoning depend on judgments of similarity. Both representativeness and analogical reasoning can be flawed because the judgments of similarity can be based on irrelevant (but salient) traits. Heuristics can be described in general terms, but specific application is dependent on particular prior states: what is salient and available to a subject, and what a subject takes to be a significant similarity. Prior knowledge, experience and training, as well as motivational factors, are the important variables. For example, geologists who worked in the Southern Hemisphere in the first

[17] Talk, Society for Philosophy and Psychology Meetings, June 1996.

part of the century were much more likely to support Wegener's theory of continental drift–because the data in support of drift were more salient and available to them, from their direct observations. Darwin's knowledge of artificial selection was extended, analogically, to propose natural selection.

Motivated reasoning is another way in which human reasoning is situated. It was discussed under the heading of 'goals', above. And finally, many argue that human reasoning is shaped by evolutionary needs in pre-history (discussed above).

3 EPISTEMOLOGICAL REFLECTIONS

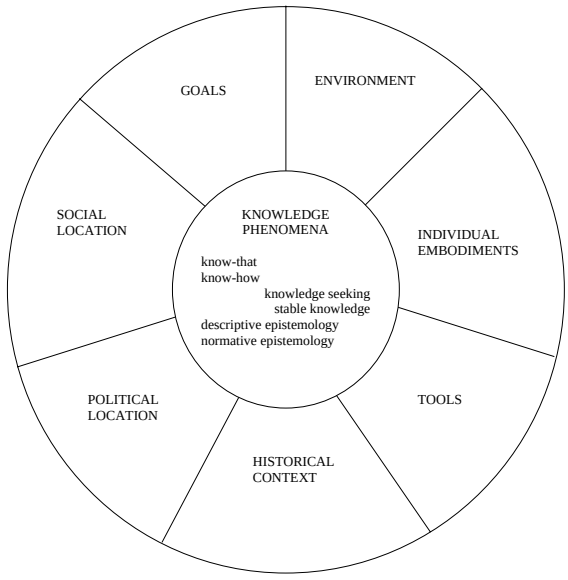

Figure 1.

Figure 1 summarizes the ways in which cognition is situated. Diagrams are themselves situated knowledge tools; they make information more salient and available than they are from text.[18] The order of kinds of situatedness is not significant, so I arranged them in a circle. The order around the circle is not significant either. Types of cognitive phenomena are not arranged or ranked, either. There are no hierarchies here, no privileging of either individual or society, and no privileging of kinds of cognitive phenomena. It is easy to represent cognitive practices that are situated in more than one way, as inkblots that cover some or

[18]This comes from the Tversky, Kahneman et al literature. A more recent, and related study, by Whittaker and Schwarz (1995) shows the benefits of material, public diagrams over electronic co-ordination tools.

all kinds of situatedness and some or all kinds of knowledge phenomena. Two examples of this are given — traditional Micronesian navigation and traditional Naval navigation — in Figures 2 and 3.

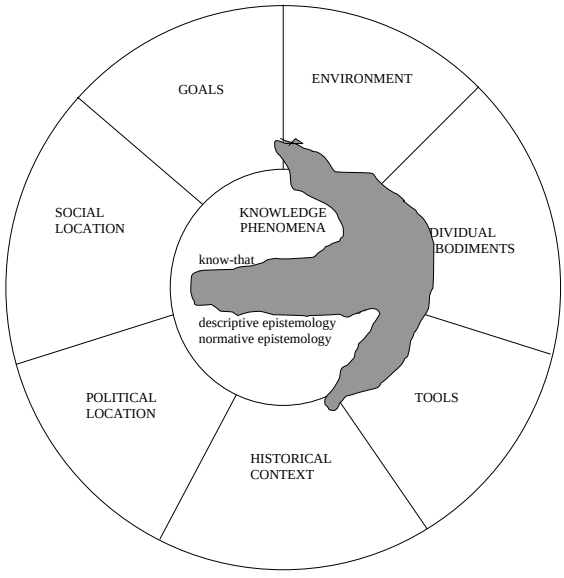

Figure 2. Traditional Micronesian navigation. Practiced by individuals in a particular environment (archipelago) according to a particular historical tradition, passed down as know-how rather than as explicit knowledge.

It is clear that naturalized epistemology ought to be highly interdisciplinary. We have come a long way from Quine's (1969, p. 82) pronouncement that 'epistemology...falls into place as a chapter of psychology'. It is not only cognitive psychology that is relevant to epistemology (as Quine and other early naturalized epistemologists thought), but also social psychology, feminist critiques, sociology, distributed artificial intelligence, situated planning, history, and the study of tools.

It is hoped that this diagram can increase awareness of the many ways in which cognition is situated among those who work on some aspect of situatedness, and thereby enhance their accounts. For example, those cognitive scientists who work on situated planning and action could incorporate political standpoint and social distribution into their analyses. Feminist critics and sociologists of science could fill in some gaps in their accounts by considering individual embodiments (not as traditional, individual rationality, but as a location for creativity and decision making). Those who work on distributed cognition could model motivated, as well as non-motivated (or 'truth-motivated') reasoning.

A natural reaction to this discussion is to say that 'situated cognition' is so inclusive a concept that it has lost all content. I am sympathetic to this reaction; it

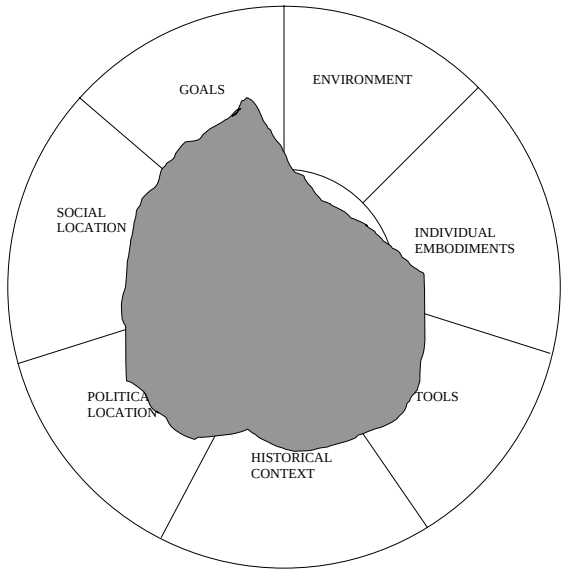

Figure 3. Traditional Naval navigation. Practiced by groups of sailors, working together with astronomical tools according to the roles appropriate to their rank, works anywhere on Earth.

arises because 'situated cognition' is really defined by exclusion, as was mentioned at the beginning of the paper. Here, concretely speaking, is what is *not* situated cognition: knowledge of universal grammar (if there is such), knowledge of Bayes' theorem, knowledge of deductive logic, general discussion of reliabilism, discussion of Marr's level of computation, Thagard's ECHO simulation of scientific decision making. Functional, abstract, over-general, symbolic, a priori and idealized accounts are eschewed by those who take situated cognition seriously, or at least highly supplemented. (For example, Goldman's reliabilism is a skeleton theory that would need to be fleshed out by specifying process, subject and domain of inquiry.[19]) If, in the future, researchers generally move to situated approaches, the term 'situated cognition' will become redundant.

The diagram is a classification of ways in which cognition is situated. Are there further classifications to make? For example, is it worth making something out of traditional distinctions between 'knowledge how' and 'knowledge that', between knowledge-seeking and stable knowledge and between descriptions and normative judgments? It is my view that these distinctions can still be made, but not in such a way that yields the traditional privileging of 'knowledge that', stable knowledge and normative judgments. And, in particular, we need all these categories in order to do epistemology of any domain of inquiry. Recent work in science studies has

[19]Goldman has done this more recently in (1999) and (2002).

shown that practice, in addition to theory (know-how in addition to know-that) describes the state of a science at any particular time. Other domains are surely no less anchored in practices.

My final conclusion is that epistemology is complex (there are many variables) and specific (lacking many global generalizations). Any normative recommendations and judgments will be for particular kinds of situation (domain, social organization, etc) and highly dependent on descriptive understanding of the situation. The move to situated cognition is thus a move away from generality in epistemology. Cartwright (1983) and Dupre (1993) have argued that generality in ontology (the 'laws of nature') is a myth, and that the most that scientists find is local regularities. I am arguing for a similar position in epistemology.

ACKNOWLEDGEMENTS

I am grateful to Gary Ebbs, Ron Giere, Cliff Hooker, Donna Keren, David Kirsh, Nancy McHugh, Paul Thagard and Sydney White for discussions about the content of this paper. An early version of this paper was presented at a symposium on situated knowledge at the APA Eastern Division Meetings, December 1996.

BIBLIOGRAPHY

[Agre, 1997] Philip Agre. *Computation and Human Experience*. Cambridge University Press, 1997.

[Antony, 1993] L. Antony. Quine as feminist: The radical import of naturalized epistemology. In *A Mind of One's Own: Feminist Essays on Reason and Objectivity*. Boulder, CO: Westview Press, 1993.

[Antony, 1995] L. Antony. Sisters, please, i'd rather do it myself: a defense of individualism in epistemology. *Special issue of Philosophical Topics: 'Feminist Perspectives on Language, Knowledge, and Reality'*, 23(2):59–94, 1995.

[Barkow *et al.*, 1995] J. Barkow, L. Cosmides, and J. Tooby. *The Adapted Mind: Evolutionary psychology and the generation of culture*. NY: Oxford University Press, 1995.

[Barwise and Perry, 1985] J. Barwise and J. Perry. *Situations and Attitude*. Cambridge, MA: MIT Press, 1985.

[Bechtel and Richardson, 1993] W. Bechtel and R.. Richardson. *Discovering Complexity*. Princeton, NJ: Princeton University Press, 1993.

[Brooks, 1991] R... Brooks. Intelligence without representation. *Artificial Intelligence 47*, pages 139–159, 1991.

[Cartwright, 1983] N.. Cartwright. *How the Laws of Physics Lie*. Oxford: Oxford University Press, 1983.

[Clancey, 1997] W.. Clancey. *Situated Cognition: On Human Knowledge and Computer Representations*. Cambridge: Cambridge University Press, 1997.

[Clarke, 1997] A. Clarke. *Being There: Putting Brain, Body and World Together Again*. Cambridge, MA: MIT Press, 1997.

[Clarke, 2001] A. Clarke. *Mindware: An Introduction to the Philosophy of Cognitive Science*. New York and Oxford: Oxford University Press, 2001.

[Cummins, 1996] D. Cummins. Evidence for the innateness of deontic reasoning. *Mind and Language*, 11(2):160–190, 1996.

[Dalmiya and Alcoff, 1993] V. Dalmiya and L. Alcoff. Are old wives tales justified? In Linda Alcoff and Elizabeth Potter, editors, *Feminist Epistemologies*, pages 217–244. New York, NY: Routledge Press, 1993.

[Dupré, 1993] J. Dupré. *The Disorder of Things: Metaphysical Foundations of the Unity of Science*. Cambridge, MA: Harvard University Press, 1993.

[Galison, 1997] P. Galison. *Image and Logic: A Material Culture of Microphysics*. Chicago: University of Chicago Press, 1997.

[Giere, 1988a] R. Giere. *Explaining Science: A Cognitive Approach*. Chicago, IL: University of Chicago Press, 1988.

[Giere, 1988b] R. Giere. *Science Without Laws*. Chicago, IL: University of Chicago Press, 1988.

[Goldman, 1999] A. Goldman. *Knowledge in a Social World*. Oxford: Oxford University Press, 1999.

[Goldman, 2002] A. Goldman. *Pathways to Knowledge: Public and Private*. Oxford and New York: Oxford University Press, 2002.

[Hacking, 1983] I. Hacking. *Representing and Intervening*. Cambridge: Cambridge University Press, 1983.

[Haraway, 1989] D. Haraway. *Primate Visions: Gender Race and Nature in the World of Modern Science*. New York, NY: Routledge Press, 1989.

[Haraway, 1991] D. Haraway. *Simians, Cyborgs and Women: The Reinvention of Nature*. New York, NY: Routledge Press, 1991.

[Harding, 1986] S. Harding. *The Science Question in Feminism*. Ithaca, NY: Cornell University Press, 1986.

[Harding, 1991] S. Harding. *Whose Science? Whose Knowledge? Thinking from Women's Lives*. Ithaca, NY: Cornell University Press, 1991.

[Hartsock, 1983] N. Hartsock. *Money, Sex and Power: Towards a Feminist Historical Materialism*. Northeastern University Press, 1983.

[Holyoak and Thagard, 1995] K. Holyoak and P. Thagard. *Mental Leaps: Analogy in Creative Thought*. Cambridge, MA: MIT Press, 1995.

[Hull, 1988] D. Hull. *Science as a Process: An Evolutionary Account of the Social and Conceptual Development of Science*. Chicago, IL: University of Chicago Press, 1988.

[Hutchins, 1995] E. Hutchins. *Cognition in the Wild*. Cambridge, MA: Cambridge University Press, 1995.

[Kahneman *et al.*, 1995] D. Kahneman, P. Slovic, and A. Tversky, editors. *Cognition in the Wild*. Cambridge, MA: Cambridge University Press, 1995.

[Keller, 1985] E. Keller. *Reflections on Gender and Science*. New Haven, CT: Yale University Press, 1985.

[Kirsh, 1991] D. Kirsh. Today the earwig, tomorrow man? *Artificial Intelligence*, 47:161–184, 1991.

[Kirsh, 1995a] D. Kirsh. Complementary strategies: why we use our hands when we think. In J. D. Moore and J. M. Fain Lehman, editors, *Proceedings of the Seventeenth Annual Conference of the Cognitive Science Society*, pages 212–217. New York, NY: Routledge Press, 1995.

[Kirsh, 1995b] D. Kirsh. The intelligent use of space. *Artificial Intelligence*, 73:31–68, 1995.

[Kitcher, 1993] P. Kitcher. *The Advancement of Science*. Oxford: Oxford University Press, 1993.

[Knorr Cetina, 1999] K. Knorr Cetina. *Epistemic Cultures: How the Sciences Make Knowledge*. Cambridge Mass and London UK: Harvard University Press, 1999.

[Kuhn, 1962 [1970]] T. Kuhn. *The Structure of Scientific Revolutions*. Chicago: University of Chicago Press, 1962 [1970].

[Kunda, 1990] Z. Kunda. The case for motivated reasoning. *Psychological Bulletin*, 108(3):480–498, 1990.

[Lakoff and Johnson, 1999] G. Lakoff and M. Johnson. *Philosophy in the Flesh: The Embodied Mind and Its Challenge to Western Thought*. Harper Collins Publishers, 1999.

[Latour, 1987] B. Latour. *Science in Action*. Cambridge, MA: Harvard University Press, 1987.

[Lave, 1988] J. Lave. *Cognition in Practice*. Cambridge: Cambridge University Press, 1988.

[Longino, 1990] H. Longino. *Science as Social Knowledge*. Princeton, NJ: Princeton University Press, 1990.

[Longino, 1993] H. Longino. Subjects, power and knowledge: Description and prescription in feminist philosophies of science. In Linda Alcoff and Elizabeth Potter, editors, *Feminist Epistemologies*, pages 101–120. New York, NY: Routledge Press, 1993.

[Nersessian, 1992] N. Nersessian. How do scientists think? capturing the dynamics of conceptual change in science. In Ronald Giere, editor, *Cognitive Models of Science*, pages 3–44. Minneapolis: University of Minnesota Press, 1992.

[Nisbett and Ross, 1980] R. Nisbett and L. Ross. *Human Inference: Strategies and Shortcomings of Social Judgment*. Englewood Cliffs, NJ: Prentice-Hall, 1980.

[Norman, 1988] D. Norman. *The Psychology of Everyday Things*. New York: Basic Book, 1988.

[Pickering, 1995a] A. Pickering. *The Mangle of Practice: Time, Agency and Science*. Chicago: University of Chicago Press, 1995.

[Pickering, 1995b] A. Pickering. *The Mangle of Practice: Time, Agency and Science*. Chicago: University of Chicago Press, 1995.

[Quine, 1969] W.V. Quine. Epistemology naturalized. In *Ontological Relativity and Other Essays*, pages 69–90. New York, NY: Columbia University Press, 1969.

[Resnick, 1994] M. Resnick. *Turtles, Termites and Traffic Jams: Explorations in Massively Parallel Microworlds*. Cambridge, MA: MIT Press, 1994.

[Smith, 1990] D. Smith. *The Conceptual Practices of Power: A Feminist Sociology of Power*. Northeastern University Press, 1990.

[Solomon, 2001] M. Solomon. *Social Empiricism*. Cambridge, MA: MIT Press, 2001.

[Stepan, 1986] N. Stepan. Race and gender: The role of analogy in science. In Sandra Harding, editor, *The Racial Economy of Science: Toward a Democratic Future*, pages 359–376. Bloomington and Indianapolis: Indiana University Press, 1986.

[Stich, 1990] S. Stich. *The Fragmentation of Reason: Preface to a Pragmatic Theory of Cognitive Evaluation*. Cambridge, MA: MIT Press, 1990.

[Sulloway, 1996] F. Sulloway. *Born to Rebel: Birth Order, Family Dynamics and Creative Lives*. New York: Pantheon Books, 1996.

[Thagard, 1993] P. Thagard. Studies in history and philosophy of science. *New York: Pantheon Books*, 24:49–67, 1993.

[Whittaker and Schwarz, 1995] S. Whittaker and H. Schwarz. Back to the future: pen and paper technology supports complex group coordination. In *Proceedings of CHI'95 Conference on Computer Human Interaction*, pages 495–450. NY: ACM Press, 1995.

ARTIFICIAL INTELLIGENCE: HISTORY, FOUNDATIONS, AND PHILOSOPHICAL ISSUES

B. Jack Copeland, Diane Proudfoot

Part One describes the history of AI, including its origins in Turing's work during World War II, and outlines the fundamental developments in the field to the present day. Topics covered include search, learning, board games, evolutionary computing, the micro-world approach to AI, robotics, expert systems, common-sense reasoning, neuron-like computing, and situated (or 'nouvelle') AI. Part Two investigates a number of philosophical issues in AI, including the Turing test and its motivation, John Searle's Chinese room argument, AI and hypercomputation, and the 'Gödel Objection' to AI (made famous by Roger Penrose).

Part One
The History and Development of Artificial Intelligence

1 THE ORIGINS OF AI

The earliest substantial work in Artificial Intelligence was carried out by the British logician and computer pioneer Alan Turing. In 1936 Turing dreamed up an abstract digital computing machine — the 'universal Turing machine' [Turing, 1936]. (His machine is universal in the sense that it can be programmed to carry out any calculation that can be performed by an idealized 'human computer', a clerk who works systematically and who has unlimited time and an endless supply of paper and pencils.) Turing's abstract machine consists of a scanner and a limitless memory-tape that moves back and forward past the scanner; the scanner reads the symbols on the tape and writes further symbols. The tape is both the memory and the vehicle for initial data ('input') and output. The machine has a small repertoire of basic operations; complexity of operation is achieved by chaining together basic operations. (Turing referred to his abstract machines simply as 'computing machines' — the American logician Alonzo Church dubbed them 'Turing machines' [Church, 1937, 43].)

The universal machine has a single, fixed table of instructions built into it. These 'hard-wired' instructions enable the machine to read and execute (encoded) instructions inscribed on its tape. By inscribing different programs on the tape, the machine can be made to carry out different tasks. This, then, was Turing's 'stored

Handbook of the Philosophy of Science. Philosophy of Psychology and Cognitive Science
Volume editor: Paul Thagard
General editors: Dov M. Gabbay, Paul Thagard and John Woods
© 2007 Elsevier B.V. All rights reserved

program' concept, the idea fundamental to computer science: a single machine of fixed structure that, by making use of coded instructions stored in memory, could change itself, chameleon-like, from a machine dedicated to one task into a machine dedicated to another. All modern (stored-program digital) computers are in essence universal Turing machines — with, of course, real-world limitations on memory, durability, etc.

During the Second World War Turing was a leading cryptanalyst at the Government Code and Cypher School (GC & CS), Bletchley Park, where a large proportion of the Wehrmacht's enciphered radio communications were decoded. Turing was thinking about machine intelligence (the term in use in Britain, predating 'artificial intelligence') at least as early as 1941 — in particular about the possibility of computing machines that solved problems by means of searching through the space of possible solutions, guided by rules of thumb (or what would now be called 'heuristic' principles), and about the mechanization of chess.[1] At Bletchley Park, in his spare time, Turing discussed these topics and also machine learning, and he circulated a typescript concerning machine intelligence among some of his colleagues.[2] Now lost, this typescript was undoubtedly the earliest paper in the field of AI. Chess and other board games have been regarded ever since as an important test-bed for ideas in AI (chiefly because, while simple in their essentials, they are a source of difficult, precisely stateable problems), and both search involving heuristics and machine learning are major parts of modern AI.

Turing's Bombe — the electro-mechanical computing device used by GC & CS to break German Enigma traffic — was the first milestone in the history of machine intelligence.[3] Turing completed the logical design of the Bombe in the last months of 1939. The first Bombe, named 'Victory', was installed at Bletchley Park early in 1940 [Mahon, 1945, 292]. The Bombe was a computing machine with a very narrow and specialized purpose — searching through the Enigma machine's settings, at superhuman speed, in order to find the settings at which the particular message under attack had been encrypted.

Modern AI researchers speak of the method of 'generate-and-test'. Potential solutions to a given problem are generated by means of a guided search through the space of possible solutions (e.g. the configurations of a chess board). These potential solutions are then tested by an auxiliary method to find out if any of them actually is a solution. Nowadays in AI, both processes, generate and test, are typically carried out by the same program. The Bombe mechanized the first process. The space of possible solutions consisted of configurations of the Enigma machine. The search could be guided in various ways; one involved what Turing called the 'multiple encipherment condition' associated with a 'crib' [Turing, 1940, 317]. (A crib is a word or phrase that the cryptanalyst believes might be part

[1] Donald Michie in interview with Copeland (October 1995).

[2] Donald Michie in interview with Copeland (February 1998).

[3] We include in this claim the Polish Bomba (see [Rejewski, 1980]), a more primitive form of the Bombe, which also employed guided search (although cribs were not used).

of the German message. For example, it might be conjectured that a certain message contains 'WETTER FUR DIE NACHT' (weather for the night). Cribs were relatively plentiful on many Enigma networks, thanks both to the stereotyped nature of German military messages and to lapses of cipher security. One station sent exactly the same message each day for a period of several months: 'beacons lit as ordered' [Mahon, 1945, 303].) A search guided in this fashion, Turing said, would 'reduce the possible positions to a number which can be tested by hand methods' [1940, 317]. The testing of the potential solutions was carried out by setting up a replica Enigma accordingly, typing in the cipher text, and seeing whether or not German words emerged. Turing described other heuristics for guiding the Bombe's search in his famous 'Treatise on the Enigma'. In sum, the central concept of early AI — heuristic search — originated in the Bombe, and it was at Bletchley Park that Turing glimpsed the possibility of achieving intelligence through mechanical search.

In 1936 the universal Turing machine existed only as an idea. But right from the start Turing was interested in the possibility of actually building such a machine.[4] His wartime acquaintance with electronics (in 1943 GC & CS built the world's first large-scale electronic — but not stored-program — digital computer, Colossus, in order to break the German Lorenz machine messages[5]) was the key link between his earlier theoretical work and his 1945 design for an electronic stored-program digital computer. In that year Turing was recruited by the National Physical Laboratory (NPL) in London. His technical report 'Proposed Electronic Calculator' [Turing, 1945], containing his design for the Automatic Computing Engine (ACE), was the first relatively complete specification of an electronic stored-program digital computer.[6] Turing regarded the ACE as a 'practical version' of the universal

[4]Newman in interview with Christopher Evans ('The Pioneers of Computing: An Oral History of Computing' (London: Science Museum)); 'Dr. A. M. Turing', *The Times*, 16 June 1954, p. 10.

[5]Colossus was designed and built during 1943 by Thomas Flowers and his team at the Post Office Research Station in Dollis Hill, London, in consultation with the Cambridge mathematician Max Newman, head of the section at Bletchley Park known simply as the Newmanry [Copeland, 2006]. (Turing attended Newman's lectures on mathematical logic at Cambridge before the war; these lectures launched Turing on the research that led to his 'On Computable Numbers' (Newman in interview with Christopher Evans, 'The Pioneers of Computing: An Oral History of Computing' (London: Science Museum)).) Colossus first worked at the start of 1944 — two years before the first comparable U.S. machine, the ENIAC, was operational. It was not until 2000 that the British finally declassified the complete account of Colossus' wartime role ('General Report on Tunny, with Emphasis on Statistical Methods', written in 1945 by Good, Michie, and Timms — members of the Newmanry (National Archives / Public Record Office (PRO), document reference HW 25/4, HW 25/5 (2 vols)). A digital facsimile of this report is in The Turing Archive for the History of Computing `http://www.AlanTuring.net/tunny_report`.

[6]In the U.S. von Neumann's [1945] 'First Draft of a Report on the EDVAC' (Electronic Discrete Variable Calculator) appeared some months before Turing's 'Proposed Electronic Calculator' but contained little engineering detail, in particular with regard to electronic hardware. In contrast, Turing provided detailed circuit designs and detailed specifications of hardware units, specimen programs in machine code, and even an estimate of the cost of building the ACE. For a discussion of the relation between Turing's work and von Neumann's, see [Copeland and Proudfoot, 2005].

Turing machine [Turing, 1947, 379]. In 'Proposed Electronic Calculator' he said:

> 'Can the machine play chess?' It could fairly easily be made to play a rather bad game. It would be bad because chess requires intelligence. We stated at the beginning of this section that the machine should be treated as entirely without intelligence. There are indications however that it is possible to make the machine display intelligence at the risk of its making occasional serious mistakes. By following up this aspect the machine could probably be made to play very good chess. [1945, 389]

In working on the ACE, Turing described himself as 'building a brain'[7] and remarked that he was 'more interested in the possibility of producing models of the action of the brain than in the practical applications to computing'.[8] From its very beginnings, AI was pursued in the hope of illuminating how the brain or mind works. Turing neatly summed up the link between Artificial Intelligence and cognitive science:

> The whole thinking process is still rather mysterious to us, but I believe that the attempt to make a thinking machine will help us greatly in finding out how we think ourselves. [1951, 486]

In London in 1947 Turing gave what was, so far as is known, the earliest public lecture to mention computer intelligence. He described the human brain as a 'digital computing machine' [1947, 382] and discussed the prospect of machines acting intelligently, learning, and beating human opponents at chess. He stated that '[w]hat we want is a machine that can learn from experience' and that '[t]he possibility of letting the machine alter its own instructions provides the mechanism for this' [1947, 393]. This possibility — a computer that operates on and modifies its own program as it runs, just as it operates on the data in its memory — is implicit in Turing's concept of a stored-program machine. (Turing also set out in his lecture what he later called the 'Mathematical Objection' to the view that minds are machines. See Section 9.)

In 1948 Turing wrote (but did not publish) a report entitled 'Intelligent Machinery' [Turing, 1948]. It followed a year's sabbatical leave, granted with the purpose of enabling Turing 'to extend his work on the [ACE] still further towards the biological side ... hitherto the machine has been planned for work equivalent to that of the lower parts of the brain, and [Turing] wants to see how much a machine can do for the higher ones; for example, could a machine be made that could learn by experience?'.[9] 'Intelligent Machinery' is a wide-ranging and strikingly original survey of the prospects of AI. It is AI's first manifesto.

[7] [S. Turing, 1959, 75]; Don Bayley in interview with Copeland (December 1997).

[8] Turing in a letter to W. Ross Ashby, in [Copeland, 2004a, 374–5].

[9] Letter from Darwin to Appleton (23 July 1947) (PRO document reference DSIR 10/385; a digital facsimile is in The Turing Archive for the History of Computing `http://www.AlanTuring.net/darwin_appleton_23jul47`).

In this report Turing introduced many of the concepts that were later to become central to the field, in some cases after reinvention by others. He sketched the logic-based approach to problem-solving, now widely used. In a bold development of his war-time thinking about search, he hypothesized that 'intellectual activity consists mainly of various kinds of search' [1948, 431]; eight years later the same hypothesis was put forward independently by Newell and Simon and through their influential work became one of the principal tenets of AI (see, for example, [Newell and Simon, 1976]). Also in 'Intelligent Machinery' Turing described his (paper) experiments on learning, involving the modification of an initially 'unorganised machine' by a process analogous to teaching by reward and punishment [1948, 416–23]. He later called this machine a 'child-machine' [1950a, 461], saying that he had 'succeeded in teaching it a few things, but the teaching method was too unorthodox for the experiment to be considered really successful' [ibid.]. 'Intelligent Machinery' contains probably the first suggestion that computing machines could be built out of simple, neuron-like elements connected together into networks in an initially random manner; Turing proposed that initially random networks be 'trained' (his term) to perform specific tasks (see Section 4).

In 'Intelligent Machinery', in a brief passage concerning what he called 'genetical or evolutionary search' [1948, 431], Turing introduced the concept of a genetic algorithm[10], now important both in AI and in (what since 1987 is called) Artificial Life (A-life).[11] Turing had the idea of (what is now called) evolutionary computing, where programs are developed in a manner reminiscent of evolution by natural selection, the 'criterion being survival value' [1948, 431]. In his final years Turing worked on Artificial Life, the central aim of which is a theoretical understanding of naturally occurring biological life — in particular, of the most conspicuous feature of living matter, its ability to self-organize (i.e. to develop form and structure spontaneously). Turing was the first to use computer simulation to investigate a theory of the development of organization and pattern in living things, such as fircones and hydra [Turing, 1952].[12] He began this investigation in 1951, as soon as the first Ferranti Mark I was installed at the University of Manchester. (This machine was the world's first commercially available electronic stored-program digital computer, built by the Manchester firm of Ferranti in conjunction with the University of Manchester Computing Machine Laboratory — where Turing was Deputy Director, there being no Director.[13]) Turing described this research as 'not altogether unconnected' to his work on neural networks; his study of biological

[10]The term 'genetic algorithm' was introduced circa 1975 [Holland, 1992, x].

[11]The term is due to Christopher Langton [Langton, 1986].

[12]Turing employed nonlinear differential equations to describe the chemical interactions hypothesized by his theory and used the Manchester computer to explore instances of such equations. He was probably the first researcher to engage in the computer-assisted exploration of nonlinear systems.

[13]The first U.S. commercial machine, the Eckert-Mauchly UNIVAC, appeared shortly afterwards. The Ferranti Mark I was the production model of the Manchester Mark I (successor to the Manchester 'Baby' — see Section 2) and was completed in February 1951 [Lavington, 1975. 20].

growth would, he hoped, assist him in developing 'a fairly definite theory' of 'the genetical embryological mechanism' by which 'brain structure ... [is] achieved'.[14]

Turing's 'Intelligent Machinery' distinguished two different approaches to AI, now usually termed 'top-down AI' and 'bottom-up AI'. In top-down AI, cognition is treated as a phenomenon independent of the particular details of the mechanism (e.g. brain or electronic artefact) that is implementing it. Researchers in bottom-up AI take the opposite approach. The difference between the two strategies may be illustrated by considering the task of building a system to discriminate between the letter 'W' and other letters. A top-down approach could involve writing a computer program that checks inputs of letters against a description of a 'W' couched in terms of the angles and relative lengths of intersecting line segments and stored in the computer's memory as a structure of symbols. A bottom-up approach could involve presenting letters one-by-one to a network of artificial 'neurons' configured somewhat like a retina and reinforcing those neurons that respond more vigorously to a 'W' than to other letters. Researchers who construct (or simulate) networks of artificial neurons similar to the neurons in the human brain and who investigate which aspects of cognition can be recreated in these artificial networks work in bottom-up AI.

2 EARLY AI PROGRAMS

The first AI programs ran in Britain in 1951-1952, in Manchester and Cambridge. This was due in part to the fact that the world's first electronic stored-program digital computers ran in Britain (the first was the Manchester 'Baby', in 1948 at the University of Manchester Computing Machine Laboratory,[15] and the second was the EDSAC, in 1949 at the University of Cambridge Mathematical Laboratory), and in part to Turing's influence on the first generation of computer programmers.

Both during and after the war Turing experimented with machine routines for playing chess: in the absence of a computer, the machine's behaviour was simulated by hand, using paper and pencil. In principle, a chess-playing computer can play by searching exhaustively through all the (very many!) available moves, but in practice this is impossible; heuristics are necessary to narrow and to guide the search. Turing experimented with two heuristics which later became common in AI: 'minimax' (assume that your opponent will move in a way that maximizes their gains and make your move in such a way as to minimize the losses that your opponent's expected move will cause you), and 'best first' (rank the moves available to you according to a numerical scoring system and examine first the consequences of the highest-scoring move).[16] In 1948 Turing and David Champernowne, the mathematical economist, constructed the loose system of rules dubbed

[14] Letter from Turing to J. Z. Young (8 February 1951); an extract is in [Copeland, 2004a, 517].

[15] For an account of Turing's input to the developments at Manchester, see [Copeland and Proudfoot, 2005].

[16] Michie in interview with Copeland (October 1995).

'Turochamp'.[17] (Champernowne reported that his wife, a beginner at chess, took on the Turochamp and lost.) Turing began to code the Turochamp for the Ferranti Mark I but never completed the task [Michie, 1966, 189]. He later published a classic early article on chess programming [Turing, 1953].

Dietrich Prinz, who worked for Ferranti, wrote the first chess program to be implemented [Prinz, 1952]. It ran in November 1951 on the Ferranti Mark I [Bowden, 1953, 295]. (Prinz 'learned all about programming the Mark I computer at seminars given by Alan Turing'.[18]) Prinz's program was for solving simple problems of the mate-in-two variety. Unlike the Turochamp, it could not play a complete game and operated by exhaustive search rather than under the guidance of heuristics. The program would examine every possible move until a solution was found; on average, several thousand moves had to be examined in the course of solving a problem and the program was considerably slower than a human player. (Prinz also used the Ferranti Mark I to solve logical problems, and in 1949 and 1951 Ferranti built two small experimental special-purpose computers for theorem-proving and other logical work ([Mays and Prinz, 1950]; [Prinz and Smith, 1953]). In 'Intelligent Machinery' Turing had emphasized that once a computer can prove logical theorems it will be able to search intelligently for solutions to problems.)

Christopher Strachey's Draughts Player was — apart from Turing's 'paper' chess-players — the first AI program to use heuristic search. In May 1951[19] Strachey coded his program for the Pilot Model ACE (running its first program in May 1950, this was the NPL's minimal version of Turing's ACE; see [Copeland, 2005a; 2005b]). His attempt to get the program running on the Pilot ACE was initially defeated by coding errors; when he returned to the NPL with a debugged version, he found that the engineers had made a major hardware change, with the result that the program would not run without substantial revision.[20] Strachey was also dissatisfied with the program's method of evaluating board positions. He transferred his attention to the Ferranti Mark I at Manchester University and wrote an improved version of the program, which he finally got working in mid-1952 (with Turing's encouragement and utilizing the latter's recently completed *Programmers' Handbook* ([Turing, 1950b]; [Campbell-Kelly, 1985, 24])). By the summer of 1952 the program could 'play a complete game of Draughts at a reasonable speed' [Strachey, 1952, 47]. It used simple heuristics and looked ahead 3-4 turns of play; the state of the board was represented on the face of a cathode ray tube — one of the earliest uses of computer graphics.

The earliest functioning AI programs to incorporate learning were written by

[17]Letter from Champernowne, in [Copeland, 2004a, 563–4].

[18]Gradwell, C. 'Early Days', reminiscences in a Newsletter 'For those who worked on the Manchester Mk I computers', April 1994. (Copeland is grateful to Prinz's daughter, Daniela Derbyshire, for sending him a copy of Gradwell's article.) Prinz later wrote a programming manual for the Mark I (Prinz, D. G. 'Introduction to Programming on the Manchester Electronic Digital Computer', no date, Ferranti Ltd; a digital facsimile is in The Turing Archive for the History of Computing `http://www.AlanTuring.net/prinz`).

[19]Letter from Strachey to Michael Woodger (13 May 1951) (in the Woodger Papers).

[20]Letters from Woodger to Copeland (15 July 1999, 15 September 1999).

Anthony Oettinger and ran in 1951 on the Cambridge EDSAC.[21] Like Strachey, Oettinger was considerably influenced by Turing's views on machine learning.[22] (In a letter to Turing on the evening of Turing's May 1951 radio broadcast 'Can Digital Computers Think?' [Turing, 1951], Strachey wrote: 'your remark ... that the programme for making a machine think would probably have great similarities with the process of teaching ... seems to me absolutely fundamental'.[23]) Oettinger's 'response-learning programme' could be taught to respond appropriately to given stimuli by means of expressions of 'approval' or 'disapproval' by the teacher. Oettinger claimed that the 'behaviour pattern of the response-learning ... machine is sufficiently complex to provide a difficult task for an observer required to discover the mechanism by which the behaviour of the ... machine is determined' [1952, 1257]. In 'Intelligent Machinery' Turing had said that we do not call a machine intelligent if we find 'out the rules of its behaviour' ([1948, 431]; see Section 6.2 below).

Using Turing's terminology, Oettinger described another learning program, his 'shopping programme', as a 'child machine'. This program (which we call 'Shopper') simulates the behaviour of 'a small child sent on a shopping tour' [1952, 1247]. In its simulated world of eight shops, Shopper is instructed to buy some item, and searches for it at random. While searching, Shopper memorizes a few of the items stocked in each shop that it visits; if sent out again for the original item, or for some other item whose location it has learned during previous searches, it goes directly to the appropriate shop.

The earliest AI program to run in the US, so far as we have been able to discover, was also a draughts (checkers) program; it was written by Arthur Samuel for the IBM 701 (the company's first mass-produced electronic stored-program digital computer) and first ran at the end of 1952.[24] Samuel took over the essentials of Strachey's program[25] (which Strachey had publicized at a computing conference in Canada in 1952) and over a period of years considerably extended it. In 1955 he added features [Samuel, 1959] that enabled the program to learn from experience and thereby improve its play. (Oettinger's shopping program learned simply by rote and his response-learning program operated 'at a level roughly corresponding to that of conditioned reflexes' [Oettinger, 1952, 1257].) Samuel's program learned enough to outplay its creator and (in 1962) to win a game against a former Connecticut checkers champion (who then, however, immediately won six games straight). Samuel appears to have been the first programmer to use a genetic algorithm. He installed two copies of his checkers program, Alpha and Beta, on the same computer and left them to play game after game with each other. The computer made small changes to Alpha's heuristics, leaving Beta unchanged, and then compared Alpha's and Beta's performances over a few games. If Alpha

[21] Letter from Oettinger to Copeland (19 June 2000).

[22] Oettinger in interview with Copeland (January 2000).

[23] Letter from Strachey to Turing (15 May 1951) (in the Modern Archive Centre, King's College, Cambridge; catalogue reference D5); an extract is in [Copeland, 2003, 521–2].

[24] Letter from Samuel to Copeland (6 December 1988).

[25] Samuel [1959, 104].

played worse than Beta, the changes were discarded; if Alpha played better than Beta, the same changes were made to Beta. Over many such cycles of mutation and selection the program's skill increased.

In 1955-1956 Newell, Shaw, and Simon wrote the Logic Theorist [Newell *et al.*, 1957], a program designed to prove theorems from Whitehead and Russell's *Principia Mathematica* (in one case, the proof devised by the program was several lines shorter than the one given by Whitehead and Russell[26]). The Logic Theorist was the major exhibit at a conference in 1956 at Dartmouth College, organized by John McCarthy and entitled 'The Dartmouth Summer Research Project on Artificial Intelligence'. This was the first use of the term 'Artificial Intelligence'; Turing's original term — 'machine intelligence' — has also persisted, especially in Britain. (It is often falsely claimed that Artificial Intelligence itself was born in the mid-1950s in the U.S. (for example, [Haugeland, 1985, 176].)

Newell, Simon, and Shaw went on to write the General Problem Solver (GPS). The first version of this AI program ran in 1957 and work continued on the project for about a decade [Ernst and Newell, 1969]. GPS could solve an impressive variety of puzzles (for example, the 'missionaries and cannibals' problem), searching for a solution in a trial-and-error fashion under the guidance of heuristics. Other well-known early programs are Eliza [Weizenbaum, 1966] and Parry [Heiser *et al.*, 1966]. These programs simulated the conversational behaviour of, respectively, a human therapist and a human paranoiac. In each case the program relied on canned responses (constructed in advance by the program's designer(s) and stored in the computer's memory) and simple programming tricks (such as editing and returning the human participant's remark).

3 MICRO-WORLD AI

A micro-world is a small, simple, artificial 'world' — consisting, for example, of coloured play blocks of various shapes and sizes arrayed on a flat surface. The world may be real or virtual. Minsky and Papert of the MIT AI Laboratory proposed that, in order to simplify research, AI focus temporarily on programs capable of intelligent behaviour in micro-worlds [Minsky and Papert, 1972]. One early success was SHRDLU at MIT [Winograd, 1972]. In SHRDLU's virtual world, a robot arm operated above a flat surface strewn with play blocks. SHRDLU responded to natural language commands, such as 'Will you please stack up both of the red blocks and either a green cube or a pyramid'. The program planned a sequence of actions and the robot arm arranged the blocks appropriately. SHRDLU could answer questions such as 'Is there anything which is bigger than every pyramid but is not as wide as the thing that supports it?' (SHRDLU replied 'Yes, the blue block'). It could also answer questions about its own actions, for example

[26]Newell, Shaw, and Simon wrote up the proof and sent it to the *Journal of Symbolic Logic*. This was almost certainly the first paper to have a computer listed as a co-author, but unfortunately it was rejected. (Shaw in interview with Pamela McCorduck [McCorduck, 1979, 143].)

'Had you touched any pyramid before you put one on the green block?' (SHRDLU replied 'Yes, the green one').

Shakey was a mobile robot developed at the Stanford Research Institute by Raphael, Nilsson, and their research group during 1968-1972 [Raphael, 1976]. The robot occupied a specially built micro-world consisting of walls, doorways, and a few simply-shaped wooden blocks. Each wall had a carefully painted baseboard to enable the robot to 'see' (Shakey's primary sensor was a black-and-white television camera) where the wall met the floor — a simplification of reality typical of the micro-world approach. Shakey had about a dozen basic abilities, such as TURN, PUSH, and CLIMB-RAMP. These could be variously combined into sequences of actions devised entirely by the robot's planning program, without human intervention. For example, when instructed to move a certain block from one room to another, Shakey obeyed by locating a ramp, pushing the ramp to the platform on which the block was situated, trundling up the ramp, toppling the block onto the floor, descending the ramp, and manoeuvring the block to the required room. Another early product of the micro-world approach was FREDDY, a stationary robot with a TV 'eye' (mounted on a steerable platform) and a pincer 'hand', constructed under Michie's direction at the University of Edinburgh from 1966–1973 [Michie, 1986]. FREDDY could recognize a small repertoire of objects, including a hammer, cup, and ball, with about 95% accuracy; it could also be taught to assemble simple objects, such as a toy car, from a kit of parts.

Schank and his group at Yale applied a form of the micro-world approach to language-processing. Their program SAM (1975) could answer questions about simple stories concerning stereotypical situations, such as dining in a restaurant or travelling on the subway ([Schank and Abelson, 1977]; [Schank, 1984]). These included questions about information implicit in a story; for example, when asked 'What did John order?', SAM would reply 'John ordered lasagne', even though the story stated only that John went to a restaurant and ate lasagne. Another program by Schank's group, FRUMP, produced summaries in three languages of wire-service news reports [DeJong, 1977].

Systems such as Shakey and FREDDY contain an internal model or representation of their worlds, i.e. a structure of symbols encoding features of the task domain. In the 1950s Newell and Simon put forward the hypothesis that the processing of structures of symbols by the human brain is the basis of human intelligence and that the processing of structures of symbols by a digital computer is sufficient to produce artificial intelligence. This 'physical symbol system hypothesis' ([Newell and Simon, 1976]; [Newell, 1980]) is the basis of (what is now called) 'symbolic AI'.

Shakey's and FREDDY's internal representations had to be continuously updated as the robot moved or the world changed. The robots' planning programs would juggle with a vast structure of symbols until descriptions were derived of actions that would transform the current situation into the desired situation. This required a large amount of computing time. Although Shakey's micro-world was designed to reduce the complexity of the internal model as much as possible, the

robot's simple sequence of actions described above took days to complete. Likewise, although FREDDY could recognize only a small number of simple objects, identifying a single object took several minutes of processing time. The 'frame problem' of updating, searching, and otherwise manipulating a large structure of symbols in realistic amounts of time is inherent in symbolic AI.

Expert systems form another example of symbolic AI and of the micro-world approach. An expert system is a program dedicated to solving problems and giving advice within a specialized area of knowledge, the programmer's aim being to match the performance of a human specialist. The basic components of an expert system are a knowledge-base and an 'inference engine'. An interviewer ('knowledge engineer') organizes information elicited from human experts in the area in question into a collection of rules, typically of *if ... then* structure. The inference engine enables the expert system to draw conclusions from the knowledge-base. In 1965 Feigenbaum and Lederberg began work on Heuristic Dendral (subsequently DENDRAL), the high-performance program for chemical analysis that was the model for much of the ensuing work on expert systems [Feigenbaum *et al.*, 1971]. DENDRAL was used in industry and universities. Work on MYCIN, an expert system for treating blood infections, began in 1972 [Shortliffe, 1976]. Expert systems have wide commercial uses, including medical diagnosis, chemical analysis, credit authorization, financial management, corporate planning, document routing in financial institutions, oil and mineral prospecting, genetic engineering, automobile design and manufacture, camera lens design, computer installation design, airline scheduling, cargo placement, and automatic customer help for home computer owners. (Dreyfus's arguments that expert systems do not, and will not, match a human expert's performance, save in highly particular task domains ([Dreyfus, 1992]; [Dreyfus and Dreyfus, 1986]) have been influential. See Section 5.)

An expert system does not take into account the aims of the program, the limits of its applicability, or how its recommendations fit into a larger context. If MYCIN were told that a patient who has been stabbed is bleeding and in convulsions, the program would attempt to diagnose a bacterial cause for the patient's symptoms. Expert systems have no 'common sense', i.e. the enormous amount of worldly knowledge required for daily life. (For example, you know that: to reach a place by the quickest route you may have to drive for a time in the opposite direction; you can pull but not (normally) push with a rope; water flows downhill; city dwellers do not usually go outside undressed; causes generally precede their effects; and so on.)

In 1984 (at the Microelectronics and Computer Technology Corporation in Texas) Lenat began the CYC project, the goal to construct a common-sense knowledge-base to serve as the foundation for future generations of expert systems ([Lenat and Guha, 1990]; [Lenat and Feigenbaum, 1991]). Knowledge-enterers ('cyclists') go through newspaper and magazine articles, encyclopaedia entries, advertisements, and so forth, asking themselves what the writer had assumed the reader would know; they then hand-code many millions of assertions, aiming to capture the information that any (human or machine) reader must have in order to

begin to understand an encyclopaedia. In 1991 Lenat and Feigenbaum predicted that CYC would become 'a system with human-level breadth and depth of knowledge' [Lenat and Feigenbaum, 1991, 224]. The CYC project is ongoing today, but many question whether it has lived up to its early promise. So far, the goal of an expert system with common sense remains a distant one.

4 THE DEVELOPMENT OF NEURON-LIKE COMPUTING

Neuron-like computing, also known as *connectionism* and *parallel distributed processing* (PDP), investigates methods of computing inspired by theories about the way neurons function. In 1943 McCulloch and Pitts published a theory according to which each neuron in the brain is a simple digital processor and the brain as a whole is a form of computing machine [McCulloch and Pitts, 1943]. Later McCulloch remarked, 'What we thought we were doing (and I think we succeeded fairly well) was treating the brain as a Turing machine'.[27]

It appears that McCulloch and Pitts did not conceive of building networks of artificial neuron-like elements. Turing (so far as we have been able to discover) was the first to think of this, in his 1948 report 'Intelligent Machinery'. Turing introduced networks of neuron-like Boolean elements connected together in a largely random fashion (we call these networks 'Turing Nets' to mark their distinctive architecture [Copeland and Proudfoot, 1999]). Turing's 'B-type' neural networks foreshadowed modern connectionist nets [Copeland and Proudfoot, 1996; 1999]. B-types consist of randomly connected two-state 'neurons' whose operation is synchronized by means of a central clock. An external agent organizes an initially random B-type network by selectively disabling and enabling connections within it. This arrangement is functionally equivalent to one in which the stored information takes the form of new connections within the network. By the application of 'appropriate interference, mimicking education', a B-type network can, Turing said, be trained to 'do any required job, given sufficient time and provided the number of units is sufficient' [1948, 422].[28] Turing's idea that an initially unorganized neural network can be organized by 'interfering training' did not appear in the work of McCulloch and Pitts, whose discussion of learning is perfunctory [McCulloch and Pitts, 1943].

Turing anticipated the procedure, now in common use, of simulating the neurons and their interconnections within an ordinary digital computer. He also saw the

[27] A remark by McCulloch during the public discussion of a lecture given by von Neumann in 1948 [von Neumann, 1963, 319].

[28] Turing claimed a proof (now lost) of the proposition that an initially unorganized B-type network with sufficient neurons can be organized to become a universal Turing machine with a given storage capacity [1948, 422]. (This proof first opened up the possibility, noted by Turing [ibid., 424], that the human cognitive system is a universal symbol-processor implemented in a neural network.) A difficulty for this claim is that not all Boolean functions can be computed by a B-type. Our solution [Copeland and Proudfoot, 1996, 367–368] is to alter the definition of a B-type in such a way as to include two — rather than one, as in Turing's discussion — of the devices we called an 'introverted pair' [ibid., 365] on each unit-to-unit connection.

need to develop training algorithms for unorganized machines and conceived of the (now standard) practice of programming the training algorithm into a computer simulation of the network. However, Turing's own research on neural networks was carried out before the first stored-program electronic computers became available. It was not until 1954 (the year of Turing's death) that Farley and Clark, working at MIT, succeeded in running the first computer simulations of small networks containing a maximum of 128 neurons. (Farley and Clark were able to train their networks to recognise simple patterns ([Farley and Clark, 1954]; [Clark and Farley, 1955].)

Turing's research into neuron-like computation remained unknown to others working subsequently in the area, even those in Britain. Modern connectionists regard not Turing but the influential Frank Rosenblatt as the founder of their approach. In 1933 the psychologist Thorndike had suggested that human learning consists in the strengthening of some (then unknown) property of neurons [Thorndike, 1933] and in 1949 the psychologist Donald Hebb suggested that it is specifically an increase in the *weights* (strengths) of interconnections between neurons in the brain that accounts for learning [Hebb, 1949]. These ideas found concrete form in Rosenblatt's 'perceptrons' [Rosenblatt, 1957; 1958; 1962] — the networks of artificial neurons that Rosenblatt began investigating in 1957 (at the Cornell Aeronautical Laboratory).

Rosenblatt's perceptrons differed only in matters of detail from types of neural network investigated previously by Farley and Clark and by various researchers in Britain (including [Beurle, 1957], [Taylor, 1956], [Uttley, 1954a-d]). Rosenblatt's contributions were extensive mathematical work in the theory of neural networks (sometimes, however, lacking in rigour) and detailed experimental investigations of the properties of neural networks, using computer simulation. Rosenblatt chose the term 'connectionist' for his approach, to emphasize the importance in learning of the creation and modification of connections between neurons [1958, 387]. (The allied expression 'parallel distributed processing', introduced later, recognizes that in neuron-like computing a large number of relatively simple processors operate in parallel and that networks store information in a distributed fashion, each individual connection participating in the storage of many different items of information.) Rosenblatt distinguished between simple two-layer networks — one layer for input and the other for output — and multi-layer networks, with one or more additional layers between the input and output layers [1962, 287ff, 313ff]. He generalized the type of training procedure used by Farley and Clark [1954], which applied only to two-layer networks, in such a way that it could be applied to multi-layer networks [1962, 292ff]. Rosenblatt used the phrase 'back-propagating error correction' to describe his training method [1962, 292]; the term 'back-propagation' is now in everyday use in connectionism to describe a family of procedures descended from the Clark-Farley-Rosenblatt method. (Key works in the development of modern back-propagation include Widrow and Hoff [1960], Werbos [1974], Rumelhart, Hinton, and Williams [1986].)

Following the publication in 1969 of an influential critique of Rosenblatt's work,

there was a hiatus in research into neuron-like computing. The critique was the work of Marvin Minsky and Seymour Papert of MIT, who demonstrated mathematically that there are a variety of tasks that simple two-layer perceptrons cannot accomplish [Minsky and Papert, 1969]. For example, no two-layer perceptron can correctly indicate at its output neuron (or neurons) whether there is an even or an odd number of neurons firing in its input layer; nor can a two-layer perceptron produce at its output layer the exclusive disjunction of two binary inputs (the so-called 'XOR problem').

Minsky and Papert's mathematical results about two-layer perceptrons, while interesting and technically sophisticated, in fact showed nothing about the abilities of perceptrons in general, since multi-layer perceptrons are able to carry out tasks that no two-layer perceptron can accomplish. Indeed, the 'XOR problem' illustrates this fact: a simple three-layer perceptron can form the exclusive disjunction of two binary inputs (as Minsky and Papert knew). Nevertheless, Minsky and Papert conjectured — without any real evidence — that the multi-layer approach is 'sterile' (their word). Somehow their analysis of the limitations of two-layer perceptrons convinced the AI community (and the bodies that funded it) of the fruitlessness of pursuing work into neural networks, and the majority of researchers turned away from the approach.

This hiatus in research persisted for well over a decade before a renaissance occurred. The causes of renewed interest in neuron-like computing in the 1980s included: a widespread perception that traditional symbol-processing AI was stagnating; the possibility of simulating larger and more complex neural networks, owing to improvements in the speed and memory capacity of digital computers; and the results obtained by a research group led by James McClelland and David Rumelhart, which were widely viewed as a powerful demonstration of the potential of neural networks. In one famous experiment Rumelhart and McClelland trained a network to form the past tenses of English verbs [Rumelhart and McClelland, 1986]. The network consisted of two layers each of 460 neurons. Root forms of verbs — such as 'come', 'look', and 'sleep' — were presented in an encoded form to the input layer. About 400 different verbs were presented one-by-one to the network and the weights of connections between the layers were adjusted after each presentation (by an automatic, local process). This whole procedure was repeated about 200 times. By this stage the network could form the correct past tenses of both the original verbs and new, unfamiliar verbs (e.g. 'cling' and 'weep') — a striking example of learning involving generalization.

Modern work in neuron-like computing is diverse. Research includes the investigation of face recognition: the face (or other object) to be identified may be photographed or drawn from different angles (e.g. [Hummel and Biederman, 1992]) and may have different expressions (e.g. [Poggio and Beymer, 1996]). (WISARD was an early commercially available face-recognition system used for security applications [Aleksander and Burnett, 1987, ch. 9].) Business applications include loan risk assessment, real estate valuation, bankruptcy prediction, and share price prediction. Medical applications include the detection of lung nodules and heart ar-

rhythmias and the prediction of patients' reactions to drugs. Telecommunications applications include control of telephone switching networks and echo cancellation in modems and on satellite links.

5 NOUVELLE AI

In 'Computing Machinery and Intelligence' Turing advocated two distinct approaches to AI: (what we might call) disembodied AI, which aims at building disembodied thinking machines carrying out an 'abstract activity, like the playing of chess', and embodied AI, which aims 'to provide the machine with the best sense organs that money can buy, and then teach it to understand and speak English' by a process that 'could follow the normal teaching of a child' [1950a, 463]. In his 1948 report he anticipated the contemporary quest to build a humanoid robot, remarking:

> One way of setting about our task of building a 'thinking machine' would be to take a man as a whole and to try to replace all the parts of him by machinery. He would include television cameras, microphones, loudspeakers, wheels and 'handling servo-mechanisms' as well as some sort of 'electronic brain'. This would of course be a tremendous undertaking. The object if produced by present techniques would be of immense size, even if the 'brain' part were stationary and controlled the body from a distance. In order that the machine should have a chance of finding things out for itself it should be allowed to roam the countryside ... [1948, 420]

(The early response to Turing's suggestion was strikingly negative. According to Woodger, Turing's assistant, Turing's colleagues at the NPL declared that 'Turing is going to infest the countryside' with 'a robot which will live on twigs and scrap iron!' [Meltzer and Michie, 1969, 1].)

Recently a number of theorists and practitioners in AI have turned away from the disembodied approach, arguing that 'bedridden' programs (Dennett's term for systems interacting with the world only via keyboard and screen or printer) such as SAM and SHRDLU can have no real idea what lasagne or a green block is. Opponents of symbolic AI reject the physical symbol system hypothesis (Section 3), claiming that symbol-processing is insufficient for intelligence. Dreyfus, symbolic AI's best-known critic ([Dreyfus, 1992]; [Dreyfus and Dreyfus, 1986]), claims that it is extremely unlikely that systems that perform 'the manipulation of unambiguously defined context-free elements by precise rules' [Dreyfus and Dreyfusm 1986, 21] will exhibit human-level cognitive expertise. He concludes that it is extremely unlikely that Good Old-Fashioned AI (Haugeland's term for traditional symbol-processing AI) will succeed. According to Dreyfus, GOFAI's proponents wrongly assume that human beings are themselves systems for manipulating context-free elements by means of rules. On the contrary, Dreyfus argues, the human expert

'isn't employing rules at all' [Dreyfus and Dreyfus, 1986, 134] and will nearly always outperform an expert system.

Dreyfus's central argument runs as follows. (For Searle's influential attack on the physical symbol system hypothesis, see Section 7.) Intelligent behaviour requires that an entity determine which facts and rules are relevant to a given problem. A system that could do this only by means of programming, i.e. by appeal to further rules, runs into an intractable problem. All real-world rules contain an implicit *ceteris paribus* condition and 'all the rules for dealing with the exceptions are also *ceteris paribus* rules. So we not only get a regress of rules for applying rules but an exponential explosion of them' [Dreyfus and Dreyfus, 1986, 80]. In Dreyfus's view, which fact or rule is relevant to a given problem is determined by the real-world *situation* of the agent, and the content of a situation is not reducible to a set of facts and rules. Hence an artificial system, to be intelligent, must be situated in the real world. (On 'situated AI' see, for example, [Brooks, 1999], [Hendriks-Jansen, 1996], [Clark, [1997].)

The situated approach that came to be known as 'nouvelle AI' was pioneered by the roboticist Rodney Brooks at the MIT AI Laboratory during the latter half of the 1980s. For Brooks, as for Dreyfus, cognition is a matter of activities, skills, and commonsense 'know-how'. Machine intelligence is grounded in behaviour rather than internal representations. Brooks' alternative to the physical symbol system hypothesis is the 'physical grounding hypothesis' [Brooks, 1990; 1991; 1994; 1995]. This is the thesis that higher-level cognitive performance emerges, not from a general-purpose central processor (according to Brooks, there is 'no explicit place where cognition is done' [1999, x]), but from the interaction of simple behaviours. Workers in AI 'must incrementally build up the capabilities of intelligent systems, having complete systems at each step of the way' [Brooks, 1991, 140] and 'should be driven by puzzles which can naturally arise in a physically grounded context' [1990, 12]. Brooks' 'subsumption' architecture is to implement the physical grounding hypothesis. Brooks' approach was anticipated circa 1960 by the British cyberneticist Grey Walter [Holland, 1997].)

Designers working in nouvelle AI aim to build machines whose behaviour is prompted, not by an internal model of the world, but directly by the dynamic and unpredictable real environment. For example, Brooks' robot Herbert (named after Herbert Simon) has as its environment the offices and work-spaces of the MIT AI Laboratory; Herbert searches on desks and tables for empty soft drink cans, which it picks up and carries away. The robot's seemingly goal-directed behaviour emerges from the interactions of about fifteen simple behaviours. By eliminating complicated internal models, engineers in nouvelle AI hope to avoid the frame problem (Section 3). A nouvelle system refers continuously to its sensors; it 'reads off' the external world whatever information it needs at precisely the time it needs it. (Brooks says that the world is its own best model — always exactly up-to-date and complete in every detail [Brooks, 1995].) Advocates of this approach hope that it will lead to robust and flexible devices that respond in real time, handle multiple constraints and pursue multiple goals, scale gradually to more complex

tasks, and more closely resemble biological systems.

The best-known of Brooks' situated robots is the upper-torso humanoid robot Cog ([Brooks *et al.*, 1998]; [Brooks, 1994]; [Brooks and Stein, 1994]). Cog is to have an 'infancy and childhood' [Dennett, 1997, 358] and is part of a project aimed at 'building the ultimate humanoid' [Brooks and Stein, 1994, 9] — a fully-fledged thinking machine. Like Brooks' other machines (such as his planetary rovers), Cog is said to be 'autonomous' in the sense that its behaviour is not simply the result of programmatic intervention. (It is even suggested that Cog's designers may gain more information concerning the robot's internal states from Cog's own self-reports than from the detailed computational data available to them [Dennett, 1994, 2000].) Cog's behaviour is to be guided in part by social cues provided by human beings. The aims of 'social robotics' are: to build artificial systems capable of learning new behaviours by means of social reinforcement (from human beings or other robots); to build robots (including service and entertainment robots) with which untrained humans can interact 'naturally' and effectively, the machines behaving in socially predictable ways and adapting to humans' needs and interests (e.g. [Breazeal and Fitzpatrick, 2000]; [Scassellati, 2000]); and to test complex models of human social and psychological development in controlled, repeatable experiments (e.g. [Scassellati, 2000; 1998]; [Mataric, 1994]).

Turing's suggestion was to teach a 'child-machine' to understand and speak English. Contemporary roboticists (influenced by work in developmental psychology and neuroscience on the connexion between emotion and intelligence) typically attempt to imitate developmentally earlier communicative abilities such as facial expression. An 'expressive robot' may be a caricature face and head — such as MIT's Kismet, an active robotic head with exaggerated facial features (furry eyebrows, doglike ears, silver eyelids, and a metal mouth with cartoon-ish tubular lips) that are designed to evoke human nurturing responses. (Kismet can also make neck and head movements and produce sounds which the human observer may interpret as expressive vocalizations.) Other face robots are quasi-realistic representations of a human face and head — such as the series of animatronic heads built by Hara and Kobayashi (at the Science University of Tokyo) to recognize and mimic human facial expressions. The face is designed to resemble that of a young female, with quasi-realistic skin, hair, teeth, and make-up.

Kismet's designer claims that the robot 'is able to show expressions analogous to anger, fatigue, fear, disgust, excitement, happiness, interest, sadness, and surprise ... which are easily interpreted by an untrained human observer' [Breazeal and Scassellati, 2000]. Observers quickly anthropomorphize humanoid and pet robots (and face agents in virtual worlds). Roboticists report that human spectators readily take a device's head posture or movement to indicate an interest in its surroundings and its eye movements to indicate a locus of attention and degree of engagement ([Breazeal and Fitzpatrick, 2000]; [Scassellati, 1998]). For example, when one observer 'intentionally put his face very close to [Kismet's] face ... [t]he robot withdrew while displaying full annoyance in both face and voice ... [The observer] immediately pushed backwards, rolling the chair across the floor to

put about an additional three feet between himself and the robot, and promptly apologized to the robot.' [Breazeal and Fitzpatrick, 2000]. An expressive robot may be engineered to take advantage of these reactions, so that the nature and intensity of its exchange with a human suits the robot's perceptual and motor capabilities. Indeed, it is claimed, 'Kismet is designed to receive and send human-like social cues to a caregiver, who can ... shape its experiences as a parent would for a child' [ibid.] — this interaction is to help Kismet 'learn the meaning these acts have for others' [Breazeal and Scassellati, 2000], just as a human infant learns to use (initially purely reflexive) behaviours such as smiling to communicate with a human care-giver ([Breazeal, 1998]; [Breazeal and Velásquez, 1998]). (Hara is more cautious, remarking that 'no real communication' exists between a human and a social robot [Hara, 2004, 9].)

These responses to robots illustrate how easy it is to make believe that a machine is a human being. In his 1948 paper Turing said that playing chess against even a 'paper machine' (a simulation of machine behaviour by a human being using paper and pencil) gives 'a definite feeling that one is pitting one's wits against something alive' [1948, 412]. Examples of make-believe abound in AI. For example, Braitenberg described his (very simple) imaginary robot vehicles as 'inquisitive', 'friendly', and 'optimistic'. They possess 'egotism', 'free will', and 'the a priori concept of 2-dimensional space'. They also 'ponder over their decisions' and '[w]hile dreaming and sleepwalking [transform] the world' [Braitenberg, 1984, 46, 83, 81, 68, 41, 19, 83]. The engineers who built Braitenberg's vehicles out of modified LEGO bricks labelled the devices 'timid', 'indecisive', 'paranoid', 'dogged', 'insecure', 'inhumane', and 'frantic' [Hogg *et al.*, 1991]; the last of these is also said to be 'a philosophical creature' which 'does nothing but think' [ibid., 8]. Yamamoto [1993] attributes joy, desperation, fatigue, sadness, and friendliness to his robot vacuum cleaner.

Cog, it is claimed, will 'want to keep its mother's face in view' and will 'work hard to keep mother from turning away'; it is to 'delight in learning, abhor error, strive for novelty, recognize progress', 'crave human praise and company', and 'exhibit a sense of humour' [Dennett, 1994, 140–141]. Leonardo (a comic fur-covered expressive robot which mimics the configurations of a human observer's facial expression) is said to 'make simple inferences about the emotional state' of human observers [Breazeal *et al.*, 2005]. According to Brooks, such claims can be entirely appropriate; if it is easy and natural to maintain the illusion of communication with a machine, then the machine should be counted as possessing fully-fledged intentional (and emotional) states.[29] However, illusions are cheap. The example of the robot vacuum cleaner demonstrates that make-believe is easily and naturally generated in the most improbable cases. The burden of proof is on the AI theorist to demonstrate, of any candidate artificial intelligence, that optimistic hypotheses about the device's abilities are not merely the product of make-believe.

[29] Rodney Brooks in interview with Copeland and Proudfoot (April 2000).

Part Two
Philosophical Issues

6 THE TURING TEST

The 'imitation game', the basis of what is now known simply as the Turing test, appears (in different forms) in Turing's 1948 report for the NPL, 'Intelligent Machinery', in his 1950 article in *Mind*, 'Computing Machinery and Intelligence', and in his 1952 BBC radio broadcast (alongside other speakers), 'Can Automatic Calculating Machines be Said to Think?'.

6.1 The 1950 and 1952 Presentations of the Test

The 1950 presentation

This is the famous version of the test [Turing, 1950a]. Turing first described an imitation game involving an interrogator and two subjects, one male (A) and one female (B). The interrogator communicates with A and B from a separate room (nowadays this would probably be by means of a keyboard and screen); apart from this the three participants have no contact with each other. The interrogator's task is to find out, by asking questions, which of A and B is the man. A's aim is that the interrogator make the wrong identification. As to B, Turing said 'The object of the game for the third player ... is to help the interrogator. The best strategy for her is probably to give truthful answers.' [1950a, 441].[30]

Turing then asked, 'What will happen when a machine takes the part of A in this game?' [ibid., 441]. The game is now one in which a computer imitates a human being (man or woman).[31] The interrogator's task is to discover which of A or B is the computer; to do so he or she is permitted to ask any question (or put any point), on any topic. The computer is allowed to do everything possible to force a wrong identification. To assess the computer's performance, we ask:

> Will the interrogator decide wrongly as often when the [computer-imitates-human] game is played ... as he does when the game is played between a man and a woman? [ibid., 441]

[30]The role of the man-imitates-woman game is frequently misunderstood. For example, Hodges claims that this game is, as an introduction to the Turing test, irrelevant — a 'red herring' [1992, 415]. However, the man-imitates-woman game is not intended as an introduction to the test at all, but rather as part of the protocol for scoring the test.

[31]Some commentators (e.g. [Sterrett, 2003], [Traiger, 2003]) claim that Turing's 1950 paper presents a test in which the computer is to impersonate a woman (rather than a human being), its degree of success being compared with a male player's degree of success at the same task. However, in 1952 Turing said that '[t]he idea of the test is that the machine has to try and pretend to be a man ... and it will pass only if the pretence is reasonably convincing' [Turing *et al.*, 1952, 495] and in 1951 merely that the point of the test is to determine whether or not a computer can 'imitate a brain' [1951, 485].

If the computer (in the computer-imitates-human game) does no worse than the man (in the man-imitates-woman game), it succeeds in the game.

Having described the computer-imitates-human game, Turing remarked

> [t]he ... question, 'Can machines think?' I believe to be too meaningless to deserve discussion. [ibid., 449]

and

> I shall replace the question ['Can machines think?'] by another ['Are there imaginable digital computers which would do well in the imitation game?'], which is closely related to it and is expressed in relatively unambiguous words. [ibid., 441]

The 1952 presentation

In a BBC radio broadcast recorded in January 1952, Turing said:

> I don't want to give a definition of thinking, but if I had to I should probably be unable to say anything more about it than that it was a sort of buzzing that went on inside my head. But I don't really see that we need to agree on a definition at all. The important thing is to try to draw a line between the properties of a brain, or of a man, that we want to discuss, and those that we don't. To take an extreme case, we are not interested in the fact that the brain has the consistency of cold porridge. We don't want to say 'This machine's quite hard, so it isn't a brain, and so it can't think.' I would like to suggest a particular kind of *test* that one might apply to a machine. You might call it a test to see whether the machine thinks, but it would be better to avoid begging the question, and say that the machines that pass are (let's say) 'Grade A' machines. The idea of the test is that the machine has to try and pretend to be a man, by answering questions put to it, and it will only pass if the pretence is reasonably convincing. A considerable proportion of a jury, who should not be expert about machines, must be taken in by the pretence. They aren't allowed to see the machine itself — that would make it too easy. So the machine is kept in a far away room and the jury are allowed to ask it questions, which are transmitted through to it: it sends back a typewritten answer. ...
>
> [The questions can concern] anything. And the questions don't really have to be questions, any more than questions in a law court are really questions. You know the sort of thing. 'I put it to you that you are only pretending to be a man' would be quite in order. Likewise the machine would be permitted all sorts of tricks so as to appear more man-like, such as waiting a bit before giving the answer, or making spelling mistakes, but it can't make smudges on the paper, any more than one can send smudges by telegraph. We had better suppose that

> each jury has to judge quite a number of times, and that sometimes they really are dealing with a man and not a machine. That will prevent them saying 'It must be a machine' every time without proper consideration.
>
> Well, that's my test. Of course I am not saying at present either that machines really could pass the test, or that they couldn't. My suggestion is just that this is the question we should discuss. It's not the same as 'Do machines think,' but it seems near enough for our present purpose, and raises much the same difficulties. [Turing *et al.*, 1952, 494–495]

The 1950 and 1952 presentations of the test are significantly different. According to the 1950 formulation, the Turing test is a three-party game involving the parallel interrogation by a human of a computer and a human foil. It is implied that the interrogator knows that one of each pair of contestants is a human and one a machine. According to the 1952 formulation, members of a jury question a series of contestants one by one; some of the contestants are machines and some humans. It is implied that the interrogators do not know the ratio of machines to humans. (The arrangements employed in the Loebner Turing Test Contest — the annual competition set up in 1991 to award a grand prize of $100,000 to the first program to pass (a version of) the Turing test — conform in these respects to the 1952 presentation.)

The test as presented in 1950 is harder for the computer to pass. This is because some small feature of the computer's performance which each interrogator might overlook in a non-competitive situation might be the decisive factor in the competitive one. The test as formulated in 1950 is, nevertheless, fairer to the machine. It appears that interrogators are determined not to be fooled by a program; to avoid this outcome, when presented with contestants one by one, interrogators tend to classify contestants as machines. The Loebner Contest provides evidence of this. In the 2000 competition, for example, no machine was mistaken for a human being, but on 10 occasions a judge took a human to be a machine [Moor, 2003b]. In the 2003 competition, on no occasion was a machine judged 'definitely a human' (and only once 'probably a human') but on 4 occasions a human was judged 'definitely a machine'.[32]

Turing's prediction

Turing believed that someday a machine would perform satisfactorily in the imitation game [Turing, 1953, 569]. But when? In the 1952 radio broadcast, Newman remarked to Turing:

> I should like to be there when your match between a man and a machine takes place, and perhaps to try my hand at making up some of the

[32] `http://www.surrey.ac.uk/dwrc/loebner/results.html`.

> questions. But that will be a long time from now, if the machine is to stand any chance with no questions barred?

Turing replied:

> Oh yes, at least 100 years, I should say. [Turing *et al.*, 1952, 495]

Using the protocol for scoring the test proposed in his 1950 paper, Turing's view is as follows: not before 2052 will a computer do as well in the machine-imitates-human game as a male player in the man-imitates-woman game (in each game no questions being barred).

6.2 The Motivation behind the Test

An operational definition?

The orthodox interpretation of the famous 1950 formulation of the test (putting aside disagreements as to the precise rules of the imitation game) is that it yields an operational definition of intelligence: a machine is intelligent (or thinks) if and only if its behaviour in the imitation game matches that of a human being. For example, French claims that '[t]he Turing Test [was] originally proposed as a simple operational definition of intelligence' [2000, 115]. Hodges, Turing's biographer, claims that Turing 'introduced ... an operational definition of "thinking" or "intelligence" ... by means of a ... guessing game' [1992, 415]. Block writes:

> An especially influential behaviorist definition of intelligence was put forward by Turing [1950a]. ... Turing's version of behaviorism formulates the issue of whether machines could think or be intelligent in terms of whether they could pass the following test ... The computer is intelligent *if and only if* the judge cannot tell the difference between the computer and the person. [1990, 248, our italics]

In a footnote Block mentions Turing's suggestion in the 1950 paper that machines may 'carry out something which ought to be described as thinking but which is very different from what a man does' [1950a, 442], and claims that Turing

> jettisoned the claim that being able to pass the Turing test is a necessary condition of intelligence, weakening his claim to: passing the Turing test is a sufficient condition for intelligence. [1990, 249–250]

However, as the long quotation in 6.1 shows, Turing explicitly denied that he was proposing a definition of 'thinking' or 'intelligence'. Even in the 1950 paper, Turing described his proposal that we replace the question 'Can machines think?' with the question 'Are there imaginable digital computers which would do well in the imitation game?' as a 'tentative' suggestion [1950a, 448] and remarked that

> [w]e cannot altogether abandon the original form of the problem [viz. 'Can machines think?'], for opinions will differ as to the appropriateness

> of the substitution and we must at least listen to what has to be said in this connexion. [ibid., 448–449]

Turing never claimed that the ability to pass the Turing test is a necessary condition for intelligence;[33] in consequence, Block's description of Turing as 'jettisoning' such a condition and 'weakening' his claim is entirely misleading.

Intelligence as an emotional concept

In 'Intelligent Machinery' Turing made explicit the philosophical motivation for the Turing test. This report contains the first (restricted, chess-playing) version of the imitation game. In a section entitled 'Intelligence as an emotional concept' Turing said:

> The extent to which we regard something as behaving in an intelligent manner is determined as much by our own state of mind and training as by the properties of the object under consideration. If we are able to explain and predict its behaviour or if there seems to be little underlying plan, we have little temptation to imagine intelligence. With the same object therefore it is possible that one man would consider it as intelligent and another would not; the second man would have found out the rules of its behaviour.
>
> It is possible to do a little experiment on these lines, even at the present state of knowledge. It is not difficult to devise a paper machine which will play a not very bad game of chess. Now get three men as subjects for the experiment A, B, C. A and C are to be rather poor chess players, B is the operator who works the paper machine. (In order that he should be able to work it fairly fast, it is advisable that he be both mathematician and chess player.) Two rooms are used with some arrangement for communicating moves, and a game is played between C and either A or the paper machine. C may find it quite difficult to tell which he is playing.
>
> (This is a rather idealized form of an experiment I have actually done.) [1948, 431].)

Turing here sets out the thesis that whether an entity (real or artificial) is intelligent or thinks is determined, at least in part, by our responses to the entity's behaviour. This thesis is *externalist* — i.e., it is an example of the view that an entity's psychological states are individuated partly in terms of the entity's history or environment, rather than solely in terms of the entity's behaviour (as for the operationalist and behaviourist) or internal states (as for the internalist).

The 1948 imitation game is an experiment to see if the interrogator playing against the paper machine will 'imagine intelligence' in the machine. The

[33] Moor emphasized this point in his classic paper on the Turing test ([1976, 249]; see also [Moor, 1987]).

computer-imitates-human game of Turing's 1950 article is simply a version of the 1948 experiment that is unrestricted with respect to the verbal inputs presented to the contestants. Contrary to the orthodox operational interpretation, the Turing test is externalist. The machine's social environment — the responses of the judges, who 'must be taken in by the [machine's] pretence' (Turing in [Turing *et al.*, 1952, 495]) — plays a crucial part in determining whether the machine thinks.

6.3 Objections to the Test

We consider three influential objections. (The putative counterexample to the Turing test contained in the Chinese room argument is discussed in Section 7. For criticism of other prominent objections to the test, see [Copeland and Proudfoot, forthcoming] and [Copeland, 2003].)

Anthropocentrism

One obvious criticism of the Turing test is that it is anthropocentric, in that it requires a computer to be capable of producing outward behaviour indistinguishable from that of a human being. For example, French (arguing against Turing's supposed operational definition of 'intelligence') objects that

> the Test provides a guarantee not of intelligence but of culturally-oriented *human* intelligence ... [T]he Turing Test [is] a test for *human* intelligence, not intelligence in general. [1990, 12]

Turing mentioned the same objection:

> The game may perhaps be criticized on the ground that the odds are weighted too heavily against the machine. ... May not machines carry out something which ought to be described as thinking but which is very different from what a man does? [1950a, 442]

However, it is clear from this remark that Turing could hardly have intended the imitation game as a test for 'intelligence in general'; it was obvious to him that an intelligent machine might fail the test just because its behaviour was distinctly non-human. The anthropocentrism objection misses the point: Turing intended the imitation game precisely as a means of testing whether or not a given machine emulates the (intellectual behaviour of the) human brain.

In this connection it is important to keep in mind that the imitation game addresses only the general question 'Can *machines* think?', and not every particular question of the form 'Can machine M think?'. Turing's proposal was that we replace the question 'Can machines think?' by the question 'Are there imaginable digital computers which would do well in the imitation game?'; he did not propose that we replace 'Can machine M think?' by 'Does machine M do well in the imitation game?'. That machine M does badly in the imitation game fails to show that it does not think.

If no machine does well in the imitation game, then the question 'Can machines think?' remains open, and must be settled by some other means. However, Turing cheerfully set aside this possibility, assuming that some machine would succeed in the game:

> [The anthropocentrism] objection is a very strong one, but at least we can say that if, nevertheless, a machine can be constructed to play the imitation game satisfactorily, we need not be troubled by this objection. [1950a, 442]

Fiendish expert objections

Fiendish expert objections are all of the form: 'An expert could unmask the computer by asking it . . . '. For example, it is sometimes pointed out (e.g. by Lenat in a paper at the 2000 Loebner Contest) that an interrogator with expert knowledge in psychology could use recent discoveries of characteristic weaknesses in human reasoning to unmask a computer. These discoveries include the facts that in some circumstances human beings fail to notice certain disconfirming instances of conditional statements and that they may assign a higher probability to a conjunction (e.g. 'Linda is a bank teller and a feminist') than to one of the conjuncts ('Linda is a bank teller') [Tversky and Kahnemann, 1983]. An imitation-game interrogator could test for such weaknesses and easily detect any computer not specifically programmed to reproduce them. Likewise, an interrogator who has studied humans' response times to certain sorts of questions (especially in experiments using priming) could employ that knowledge to identify the machine contestant [French, 1990]. Expertise in computer science could also be exploited. For example, the AI theorist is aware that an artificial system may be 'superarticulate' [Michie, 1993]; unless a computer is specially programmed, its superarticulacy may give it away. In the Loebner competitions, interrogators use their knowledge of typical weaknesses in AI programming. For example, judges probe a contestant's common sense by asking questions such as 'Do your parents have two eyes each?' and by requesting definitions of the contestants' words. They use deliberate misspellings ('what do you think of the whether?') and nonsense utterances (e.g. 'carrot dune everlast the open handle') to outfox standard programming strategies.[34]

Turing anticipated this type of objection. Although the Turing test interrogator is permitted to ask any question (or put any point) she likes, not just any interrogator is — officially — permitted. In the 1948 presentation of the test the interrogator is to be 'rather poor' at chess and in the 1952 presentation the jury 'should not be expert about machines'. It is likely that, had Turing been writing after the recent relevant experimental work, he would also have excluded interrogators who are expert about human psychology. (In the first Loebner competition, the interrogators were chosen as 'a diverse group of bright judges who had little or no knowledge of AI or computer science' [Epstein, 1992, 84]. The 2003 rules

[34] Judges' utterances from the 2003 competition `http://www.surrey.ac.uk/dwrc/loebner/docs`.

still state that '[t]he aim will be to choose Judges who are well educated, who are capable typists, and who represent the community at large, rather than experts in psychology, linguistics, artificial intelligence, etc.'.[35])

The Shannon-McCarthy objection

In 1956 Shannon and McCarthy offered this objection:

> The problem of giving a precise definition to the concept of 'thinking' and of deciding whether or not a given machine is capable of thinking has aroused a great deal of heated discussion. One interesting definition has been proposed by A.M. Turing: a machine is termed capable of thinking if it can, under certain prescribed conditions, imitate a human being by answering questions sufficiently well to deceive a human questioner for a reasonable period of time. A definition of this type has the advantages of being operational, or, in the psychologists' term, behavioristic. ... A disadvantage of the Turing definition of thinking is that it is possible, in principle, to design a machine with a complete set of arbitrarily chosen responses to all possible input stimuli ... Such a machine, in a sense, for any given input situation (including past history) merely looks up in a 'dictionary' the appropriate response. With a suitable dictionary such a machine would surely satisfy Turing's definition but does not reflect our usual intuitive concept of thinking. This suggests that a more fundamental definition must involve something relating to the manner in which the machine arrives at its responses — something which corresponds to differentiating between a person who solves a problem by thinking it out and one who has previously memorized the answer. [1956, v-vi]

This objection has occurred to a number of writers but nowadays is usually credited to Block [1981]. A 'blockhead' is a (hypothetical) program able to play the imitation game successfully, for any fixed length of time, by virtue of including a look-up table. This large, but finite, table contains all the exchanges between program and interrogator that could occur during the length of time in question. Such a program faithfully imitates the intellectual behaviour of the brain but, so the objection goes, does not think.

The formal point on which the Shannon-McCarthy objection rests — the encapsulability in a look-up table of all the relevant behaviour — would have been obvious to Turing. In his 1950 paper, Turing pointed out that the behaviour of any 'discrete state machine' (i.e. a machine 'which move[s] by sudden jumps or clicks from one quite definite state to another' [Turing, 1950a, 446]) can be represented by a finite look-up table if the machine has a finite number of states, and that a computer can mimic the machine if supplied with the table:

[35] `http://www.surrey.ac.uk/dwrc/loebner/rules.html`.

> discrete state machines ... can be described by such tables provided they have only a finite number of possible states. ... Given the table corresponding to a discrete state machine ... [and provided the calculation] could be carried out sufficiently quickly the digital computer could mimic the behaviour of [the] discrete state machine. The imitation game could then be played with the machine in question (as B) and the mimicking digital computer (as A) and the interrogator would be unable to distinguish them. Of course the digital computer must have an adequate storage capacity as well as working sufficiently fast. [1950a, 447–448]

Storage capacity and speed are crucial. In the 1952 radio broadcast Max Newman remarked:

> It is all very well to say that a machine could ... be made to do this or that, but, to take only one practical point, what about the time it would take to do it? It would only take an hour or two to make up a routine to make our Manchester machine analyse all possible variations of the game of chess right out, and find the best move that way — *if* you didn't mind its taking thousands of millions of years to run through the routine. Solving a problem on the machine doesn't mean finding a way to do it between now and eternity, but within a reasonable time ...

To this Turing replied:

> To my mind this time factor is the one question which will involve all the real technical difficulty. If one didn't know already that these things can be done by brains within a reasonable time one might think it hopeless to try with a machine. The fact that a brain *can* do it seems to suggest that the difficulties may not really be so bad as they now seem. [Turing *et al.*, 1952, 473–474]

Turing's remarks suggest the following reply to the Shannon-McCarthy objection. First, given practical limitations on storage capacity, the hypothetical 'machine with a complete set of arbitrarily chosen responses to all possible input stimuli' simply cannot be built. Second, even if such a machine could actually be constructed, it would *not* faithfully imitate the intellectual behaviour of the brain, since what the brain can do in minutes would take this machine 'thousands of millions of years'.

Assuming that blockheads do not think, the Shannon-McCarthy objection establishes that, in a possible world in which a blockhead takes only seconds to answer the interrogator's questions, the proposition 'If M does well in the imitation game, then M thinks' is false. (Such a world is, of course, very different from the actual world.) Hence the objection would succeed if, as Shannon and McCarthy believed, Turing intended the test to provide a definition of 'thinking'

— because, in that case, Turing would be claiming that 'If machine M does well in the imitation game, then M thinks' is true in *all* possible worlds. However, since Turing did not offer the Turing test as a definition (Section 6.2), the Shannon-McCarthy objection fails.

There is no textual evidence to indicate that Turing was claiming anything more than that 'If machine M does well in the imitation game, then M thinks' is *actually* true, i.e. true in the actual world. This claim is not susceptible to the Shannon-McCarthy objection.

7 THE CHINESE ROOM ARGUMENT

Searle's Chinese room argument is perhaps the best-known philosophical broadside against the possibility of AI. We distinguish four versions of the Chinese room argument and show that none is successful.

7.1 The Vanilla Version

There are four principal versions of the Chinese room argument, which we call the vanilla version, the outdoor version, the simulator version, and the gymnasium version.[36] The vanilla and outdoor versions [Searle, 1980] are directed against traditional symbol-processing AI; the simulator and gymnasium versions [Searle, 1990] are directed against connectionism. In its basic or vanilla form, the Chinese room argument addresses the case of a human clerk — call him or her Clerk — who 'handworks' a GOFAI program. The program is presented to Clerk in English in the form of a set of rule-books. One not uncommon misunderstanding takes Searle to be asserting that the rule-books contain a look-up table that pairs possible inputs directly with ready-made outputs (see, for example, [Sterelny, 1990, 220ff]). So misinterpreted, Searle's argument is weakened. In fact, Searle's intention is that the rule-books may contain any GOFAI program that is claimed by its creators (or others) to understand Chinese — or indeed to have any 'cognitive states' whatsoever [Searle, 1980, 417]. The symbols processed by the program are not limited to Chinese ideograms (which merely form a vivid example); the symbols may even be generated by a suite of sensors mounted on motor equipment that is itself under the control of the program [ibid., 420].

To the programmers outside, the verbal behaviour of the Room — the system that includes the rule-books, Clerk, Clerk's pencils, rubbers, and worksheets, the input and output provisions (paper sheets and slots) and any clock, random number generator, or other equipment that Clerk may need in order to execute the precise program in question — is indistinguishable from that of a native Chinese speaker. But does the Room understand Chinese?

Here is the vanilla argument:

[36] Our remarks on the vanilla and outdoor versions apply mutatis mutandis to the derivative robot version (found in [Searle, 1980, 420].

> [Clerk] do[es] not understand a word of the Chinese ... [Clerk] ha[s] inputs and outputs that are indistinguishable from those of the native Chinese speaker, and [Clerk] can have any formal program you like, but [Clerk] still understand[s] nothing. For the same reasons, [the] computer understands nothing ...
>
> [W]hatever purely formal principles you put into the computer, they will not be sufficient for understanding, since a human will be able to follow the formal principles without understanding ... [Searle, 1980, 418]

The flaw in the vanilla version is simple: the argument is *not logically valid*. The proposition that the formal symbol manipulation carried out by Clerk does not enable Clerk to understand the Chinese story by no means entails the quite different proposition that the formal symbol manipulation carried out by Clerk does not enable the Room to understand the Chinese story. One might as well claim that the statement 'The organisation of which Clerk is a part has no taxable assets in Japan' follows logically from the statement 'Clerk has no taxable assets in Japan'.

It is important to distinguish this, the *logical reply* [Copeland, 1993a; 1993b; 2002a] to the vanilla argument, from what Searle calls the *systems reply*. The systems reply is the following claim:

> While it is true that the individual person who is locked in the room does not understand the story, the fact is that he is merely part of a whole system and the system does understand the story. [Searle, 1980, 419]

As Searle points out, the systems reply is worthless, since it 'simply begs the question by insisting without argument that the system must understand Chinese' [ibid., 419]. The logical reply, on the other hand, is a point about entailment. The logical reply involves no claim about the truth — or falsity — of the statement that the Room can understand Chinese.

7.2 The Outdoor Version

Of course, any invalid argument can be rendered valid with the addition of further premisses (in the limiting case one simply adds the conclusion to the premisses). The trick is to produce additional premisses that not only secure validity but are sustainable.

In his discussion of the systems reply Searle says:

> My response to the systems theory is quite simple: Let the individual ... memoriz[e] the rules in the ledger and the data banks of Chinese symbols, and [do] all the calculations in his head. The individual then incorporates the entire system. ... We can even get rid of the room and suppose he works outdoors. All the same, he understands nothing

> of the Chinese, and a fortiori neither does the system, because there isn't anything in the system that isn't in him. If he doesn't understand, then there is no way the system could understand, because the system is just a part of him. [1980, 419]

We shall represent this, the outdoor version of the argument, as follows:

1. The system is part of Clerk.

2. If Clerk (in general, x) does not understand the Chinese story (in general, does not Φ), then no part of Clerk (x) understands the Chinese story (Φs).

3. The formal symbol manipulation carried out by Clerk does not enable Clerk to understand the Chinese story.

∴ 4. The formal symbol manipulation carried out by Clerk does not enable the system to understand the Chinese story.

The outdoor version is valid. Premiss (1) is perhaps innocent enough. Attention thus centres on (2), which we call the 'Part-Of' principle [Copeland, 1993b, 175]. Searle makes no mention of why he thinks the Part-Of principle is true. Yet the principle is certainly not self-evident. It is conceivable that a homunculus or homuncular system in Clerk's head should be able to understand Chinese without Clerk being able to do so. (Searle has no reservations concerning the application of predicates like 'understand' to sub-personal systems. He writes (against Dennett): 'I find nothing at all odd about saying that my brain understands English. ... I find [the contrary] claim as implausible as insisting "I digest pizza; my stomach and digestive tract don't".' [1980, 451].)

We can construct other counter-examples to the Part-Of principle. Conceivably there is a special-purpose module in Clerk's brain that produces solutions to certain tensor equations, yet Clerk himself may sincerely deny that he can solve tensor equations — he may insist that he does not even know what a tensor equation is. Perhaps it is the functioning of this module that accounts for our ability to catch cricket balls and other moving objects [McLeod and Dienes, 1993]. Clerk himself is unable to produce solutions to the relevant tensor equations even in the form of leg and arm movements, say because the output of the module fails to connect owing to the presence of a lesion. Science fiction can provide other counterexamples. Neuropharmacologists induce Clerk's liver to emulate a brain, the liver remaining in situ and receiving input directly from a computer workstation, to which the liver also delivers its output. Clerk's modified liver performs many acts of cognition that Clerk cannot; for example, Clerk stares uncomprehendingly at the screen of the computer as his liver proves theorems in quantified tense logic.

Of course, one might respond to these and similar examples as follows. Since a part of Clerk is proving a theorem of quantified tense logic (solving a set of tensor equations, etc.) *then so is Clerk* — there he is doing it, albeit to his own surprise. However, this response is not available to Searle. If Clerk's sincere denial that he is

able to solve tensor equations (or what have you) counts for nothing, then likewise his denial that he understands the Chinese story. Yet it is a cornerstone of Searle's overall case that Clerk's sincere report 'I don't speak [sic] a word of Chinese' [1980, 418] suffices for the truth of premiss (3). One might call this Searle's Incorrigibility Thesis. It, like the Part-Of principle, is left totally unsupported by Searle. (Searle sometimes says, as if to give independent support to (3), 'there is no way [Clerk] could come to understand Chinese in the [situation] as described, since there is no way that [Clerk] can learn the meanings of any of the symbols' [1990, 20]. This rhetoric simply begs the question, since the matter at issue is whether 'just having the symbols by themselves ... [is] sufficient for semantics' and that this cannot be sufficient is, allegedly, 'the point that the Chinese room demonstrated' [ibid., 21].)

If the Part-Of principle is taken to be a modal claim equivalent to 'It is not possible that (some part of Clerk understands Chinese and Clerk does not understand Chinese)' then, assuming the Incorrigibility Thesis, possible scenarios such as the foregoing do more than bear on the plausibility of the principle: they settle its truth-value. If, on the other hand, the Part-Of principle is said to be a purely contingent claim (i.e. a claim that happens to be true in the actual world but is not true in possible alternatives to the actual world), then Searle's difficulty is to produce reasons for thinking the principle true.

7.3 The Simulator Version

Searle claims:

> The Churchlands are correct in saying that the original Chinese room argument was designed with traditional AI in mind but wrong in thinking that connectionism is immune to the argument. It applies to any computational system ... whether [the computations] are done in serial or parallel; that is why the Chinese room argument refutes strong AI in any form. [Searle, 1990, 22]

Directed against the thesis that a suitable connectionist network might understand Chinese, the simulator version of the Chinese room argument is this:

> Computationally, serial and parallel systems are equivalent: any computation that can be done in parallel can be done in serial. If the man in the Chinese room is computationally equivalent to both, then if he does not understand Chinese solely by virtue of doing the computations, neither do they. [ibid., 22]

The first sentence of this argument, while not completely clear in meaning, appears too strong. Networks of universal Turing machines operating in an asynchronous manner [Copeland and Sylvan, 1999, 54] are in a sense irreducibly parallel in nature, since in general the actions of the individual processors cannot be interleaved in order to form a sequence of actions that is performable by a single

universal Turing machine. (The claim, often heard, that the activity of any finite collection of universal Turing machines can be simulated by a single universal Turing machine is simply false.[37]) We shall, therefore, represent the argument in such a way that the questionable generalisation in the first premiss does not figure. Let N be a connectionist network that is said to understand Chinese:

1. Given an appropriate program, Clerk (together with his or her pencils, paper, erasers and rule-books) is 'computationally equivalent' to N.

2. Clerk does not understand Chinese (solely by virtue of doing computations).

$\therefore$ 3. N does not understand Chinese (solely by virtue of doing computations).

The situation is that we have a serial system — call it S — simulating the parallel system N. We are being asked to endorse the inference that, since S and N are 'computationally equivalent', and S does not understand Chinese, then neither does N. But, for the reason given in 7.1, (2) does not entail that S (Clerk plus rule-books, etc.) does not understand Chinese; and so the argument stalls.

This is not the only problem. Searle has issued frequent warnings on the perils of confusing a simulation with the thing being simulated. Here are some characteristic passages.

> [N]o one supposes that a computer simulation of a storm will leave us all wet ... Why on earth would anyone in his right mind suppose a computer simulation of mental processes actually had mental processes? [1989, 37–38]

> Barring miracles, you could not run your car by doing a computer simulation of the oxidation of gasoline, and you could not digest pizza by running the program that simulates such digestion. It seems obvious that a simulation of cognition will similarly not produce the effects of the neurobiology of cognition. [1990, 23]

[37] The standard textbook proof that any finite assembly of Turing machines can be simulated by a single universal Turing machine involves the idea of the universal machine interleaving the processing steps performed by the individual machines in the assembly. The proof is sound in the case where the machines in the assembly are operating in synchrony. Under certain conditions, however, a simple network of two non-halting Turing machines, m_1 and m_2, writing binary digits to a common, initially blank, single-ended tape, T, cannot be simulated by universal Turing machine [Copeland and Sylvan, 1999]. m_1 and m_2 work unidirectionally along T, never writing on a square that has already been written on, and writing only on squares all of whose predecessors have already been written on. (If m_1 and m_2 attempt to write simultaneously to the same square, a refereeing mechanism gives priority to m_1.) If m_1 and m_2 operate in synchrony, the evolving contents of T can be calculated by the universal machine. This is true also if m_1 and m_2 operate asynchronously and the timing function associated with each machine, D_1 and D_2 respectively, is Turing-machine-computable. (D_1 is defined as follows: $D_1(n) = k$ ($n, k \geq 1$) if and only if k moments of operating time separate the n^{th} atomic operation performed by m_1 from the $n+1^{th}$; similarly for m_2 and D_2.) Where D_1 and D_2 are both Turing-machine-computable, the universal machine can calculate the necessary values of these functions in the course of calculating each digit of the sequence being inscribed on T. If at least one of D_1 and D_2 is not Turing-machine-computable, and m_1 and m_2 are not in synchrony, then m_1 and m_2 may be in the process of inscribing an uncomputable number on T.

Searle's examples show forcefully that the following form of inference is logically invalid:

(α) x is a simulation of y; y has property Φ $\therefore$ x has property Φ.

We term the fallacy pointed out by Searle the *simulation fallacy* [Copeland, 2002a]. Notice that even where x and y are both computational processes, the inference-form has exemplifications with true premisses and false conclusion. For example: x is a simulation of y; y is a parallel process; therefore x is a parallel process.

Contraposing (α) produces:

(β) x is a simulation of y; it is not the case that x has property Φ $\therefore$ it is not the case that y has property Φ.

(To contrapose an argument one swaps the conclusion with any one of the premisses and negates each of the swapped statements: P, $\sim$R $\therefore$ $\sim$Q is a contrapositive of P, Q $\therefore$ R.) It is a theorem of elementary logic that contraposing an invalid inference-form always produces an invalid inference-form.)

A crucial step in the simulator version of the Chinese room argument can be represented thus:

(γ) S is a simulation of N; it is not the case that S understands Chinese $\therefore$ it is not the case that N understands Chinese.

This reasoning exemplifies the invalid form (β). In short, Searle's attempt to extend the Chinese room argument to connectionism appears to involve the simulation fallacy.

There is another difficulty. Premiss (1) cannot be relied upon. Clerk, who we may assume has access to unlimited amounts of paper and never breaks down, is acting as a universal Turing machine. Siegelmann and Sontag have described a particular type of connectionist network that cannot be simulated by a universal Turing machine ([Siegelmann and Sontag, 1994]; [Siegelmann, 2003]). (In essence Siegelmann and Sontag take a perfectly run-of-the-mill connectionist network and allow the value of a single one of its synaptic weights to be a *real* number rather than, as is more usually the case, a rational number.) We shall call any neural network that cannot be simulated by a universal Turing machine an *o-network* (this terminology is explained in section 8.1). If N is an *o*-network, then premiss (1) is false. Some connectionists believe that the brain — abstracted out from sources of inessential boundedness, such as mortality — is an *o*-network. Searle's claim that the Chinese room argument 'refutes strong AI in *any* form' (our emphasis) is very bold.

We say more about the idea of mechanisms that cannot be simulated by a universal Turing machine in section 8.

7.4 The Chinese Gym

According to Searle, connectionism is

> subject even on its own terms to a variant of the objection presented by the original Chinese room argument. Imagine that instead of a Chinese room, I have a Chinese gym: a hall containing many monolingual English-speaking men. These men would carry out the same operations as the nodes and synapses in a connectionist architecture ... and the outcome would be the same as having one man manipulate symbols according to a rule book. No one in the gym speaks a word of Chinese ... Yet with appropriate adjustments, the system could give the correct answers to Chinese questions. [1990, 22]

We might imagine that the people in the gym simulate the network by passing each other plastic tokens, green tokens representing input along an excitatory connection and red tokens along an inhibitory connection. The number of tokens passed from one player to another in a single transaction represents the weight of the connection. Each player has a list detailing to whom they must pass their tokens and how many should be handed over. During the training phase of the simulation, the players make changes to their lists in accordance with the shouted instructions of the trainer.

Once again the logical reply is sufficient to rebut the argument. One can agree with Searle that no amount of handing around tokens and fiddling with lists of recipients will enable the individual players to understand Chinese. But there is no entailment from this to the claim that the simulation as a whole does not come to understand Chinese. The fallacy involved in moving from part to whole is even more glaring here than in the original version of the room argument.

The gymnasium version in fact consists of two inferences:

1. No individual player understands Chinese (by virtue of doing the computations).

$\therefore$ 2. The simulation as a whole — call it G — does not understand Chinese (by virtue of doing the computations).

$\therefore$ 3. The network being simulated, N, does not understand Chinese (by virtue of doing the computations).

Since (1) does not entail (2), the upper inference is unsound. Might the lower inference alone nevertheless suffice for Searle's purposes?

With the upper inference gone, how is (2) to be supported? (2) is certainly not *intuitively* compelling. It is only if we place reality firmly to one side and enter a kind of fairyland that we can entertain as true the proposition that a group of humans with their pockets full of plastic tokens is simulating a connectionist network that contains maybe as many as one thousand million million connections. We have no intuitions about fairyland and least of all the intuition that a 'brain' of token-passing slaves cannot understand Chinese. (Many of Searle's attempted objections to cognitive science involve a studied refusal to take physical and biological realities seriously. For example: 'if we are trying to take seriously the idea

that the brain is a digital computer, we get the uncomfortable result that we could make a system that does just what the brain does out of pretty much anything ... cats and mice and cheese or levers or water pipes or pigeons or anything else ...' [Searle, 1992, 207]. The theory that the brain is a digital computer does not (need we say) imply that a system that does what the brain does can *really* be made out of pigeons. Only an absurd theory could imply this — Searle is surely right about that much. Likewise connectionism does not imply that *really* a team of gymnasts can simulate the brain.)

Letting the issue of the truth of (2) hang for now, and turning to the question of validity, the inference from (2) to (3) is this: G is a simulation of N; it is not the case that G understands Chinese ∴ it is not the case that N understands Chinese. As in the case of (γ), above, this reasoning exemplifies inference-form (β), contrapositive of the rejected (α). So the inference from (2) to (3) commits the simulation fallacy.

Let us see if there is any useful way of fending off this charge. After all, not every way of cashing out the variables in a fallacious inference-form need produce an invalid argument — for example, cashing out P and Q with one and the same sentence makes a valid argument even out of the invalid form P ∨ Q ∴ P ('∨' represents inclusive disjunction).

When one computing device, x, is said to simulate another, y, it is often the case that only the input-output relationship is duplicated: in other respects, the processes carried out by x and y may differ considerably. (For example, x may replicate y's input-output behaviour by means of an algorithm that is completely different from the algorithm actually employed by y.) G, on the other hand, seems to mimic N much more closely than is the case in a common-or-garden simulation where only the input-output relationship matters. In fact, let's assume that, relative to some favoured way of distinguishing between process and implementation, G and N are carrying out one and the same computational process. This may be expressed by saying that G is *isomorphic* to N.

With the proviso that G is isomorphic to N, the inference from (2) to (3) is trivially valid: if a given computational process is insufficient for understanding then it is insufficient for understanding! The rub is that there is no more leverage for establishing that the process is insufficient in the case of one isomorph than in the case of the other. It is undeniable that Searle is home and dry *if* he can show that N, or any isomorphic copy, fails to understand Chinese by virtue of 'doing the computations'. But *that* is where we came in.

8 HYPERCOMPUTATION

Hypercomputation is the computation of functions or numbers that cannot be computed in the sense of Turing [1936], i.e. cannot be computed with paper and pencil in a finite number of steps by a human clerk working mechanically.

As has often been remarked, in 1936 (when Turing thought up the universal Turing machine) a computer was not a machine at all, but a human being, a

mathematical assistant [Copeland, 1997a; 2000]. The human computer calculated by rote, in accordance with some effective method supplied by an overseer, prior to the calculation. The overseer would also supply all requisite data, in the form of symbols on paper. Many thousands of human computers were employed in business, government, and research establishments. Turing used the term 'computer' and its cognates in this sense in his 1936 paper, saying, for example, 'Computing is normally done by writing certain symbols on paper' [1936, 75] and 'The behaviour of the computer at any moment is determined by the symbols which he is observing, and his "state of mind" ' [ibid.]. The Turing machine (or, as Turing called it, 'computing machine') is an idealization of the human computer:

> We may compare a man in the process of computing a real number to a machine which is only capable of a finite number of conditions ... The machine is supplied with a 'tape' ... [1936, 59]

There is also a technical concept of computability: a number (function) is Turing-machine-computable if and only if there is a Turing machine that can write down each successive digit (value) of the number (function), starting from a blank tape and producing each digit (value) in a finite number of steps.

The Church-Turing thesis connects the sense of 'computer' just explained with this technical sense of computability: all numbers and functions computable in the former sense — 'by human clerical labour, working to fixed rules, and without understanding' [Turing, 1945, 386] — are computable by the universal Turing machine. This thesis and its converse together state that the predicates 'is Turing-machine-computable' and 'is computable by a human rote-worker unassisted by machinery' are equivalent in extension (assuming that the rote-worker is free of limitations on time, paper, pencils, etc).

A hypercomputer is any information-processing machine, notional or real, that is able to achieve more than the traditional human clerk working by rote [Copeland and Proudfoot, 1999]. Hypercomputers compute functions or numbers, or more generally solve problems or carry out tasks, that lie beyond the reach of the universal Turing machine of 1936. (Some working in the field speak of 'breaking the Turing barrier', 'computing beyond the Turing limit', 'escaping the Turing tar pit', and the like.) The additional computational power of a hypercomputer may arise because the machine possesses, among its repertoire of fundamental operations, one or more operations that no human being unaided by machinery can perform. Or the additional power may arise because certain of the restrictions customarily imposed on the human computer are absent in the case of the hypercomputer — for example, the restrictions that data take the form of symbols on paper, that all data be supplied in advance of the computation, and that the rules followed by the computer remain fixed for the duration of the computation. In one interesting family of hypercomputers, what is relaxed is the restriction that the human computer produce the result, or each digit of the result, in some finite number of steps.

One set of functions (or numbers) is of special interest: the functions (or num-

bers) that are in principle computable in the real world. The exact membership of this set is an open question. Could a hypercomputer actually exist? When we introduced the term 'hypercomputation' [Copeland and Proudfoot, 1999, 77] we were naming an emerging field with a substantial history, and speculation that there may be physical processes — and so, potentially, machine-operations — whose behaviour cannot be simulated by Turing machine stretches back over at least four decades. (See, for example, [Da Costa and Doria, 1991]; [Doyle, 1982]; [Geroch and Hartle, 1986]; [Komar, 1964]; [Kreisel, 1967; 1974]; [Penrose, 1989; 1994]; [Pour-El and Richards, 1979; 1981]; [Scarpellini, 1963]; [Stannett, 1990; 2003]; [Vergis *et al.*, 1986]. An historical survey is given in Copeland [2002b] (see also Copeland (ed.) [2002-3]).)

Is the mind — or the brain — some form of hypercomputer? Is hypercomputation the route to machine intelligence, including full mathematical creativity? The view that the mind is some form of hypercomputer — and the relevance of this view to AI — is yet to be fully explored [Copeland 1997b; 2000; 2004b]. As Mario Bunge remarked, the traditional computational approach 'involves a frightful impoverishment of psychology, by depriving it of nonrecursive functions' [Bunge and Ardila, 1987, 109].[38]

8.1 Oracle Machines

Turing's 'oracle machines' or *o*-machines may be regarded as an abstract form of hypercomputer (the earliest to appear in the literature). Turing introduced the concept of an *o*-machine in his PhD thesis [Princeton, 1938]. Subsequently published as Turing [1939], this is a classic of recursive function theory.

In detail, the fundamental processes of the standard Turing machine are: (1) move left one square; (2) move right one square; (3) identify the symbol in the scanner; (4) write a symbol on the square of tape in the scanner (first deleting the symbol already inscribed there, if the square is not blank); (5) change internal state. These fundamental processes are made available by unspecified subdevices of the machine — 'black boxes'. (The question of what mechanisms might appropriately occupy these black boxes is not relevant to the machine's logical specification.)

An *o*-machine is a Turing machine augmented with a fundamental process that produces the values of some non Turing-machine-computable function — for example, the Turing-machine halting function, $H(x,y)$. (The halting function is defined

[38] Hilary Putnam is one of the few writers on the philosophy of mind to question the proposition that Turing machines provide a maximally general formulation of the notion of machine:

> [M]aterialists are committed to the view that a human being is — at least metaphorically — a machine. It is understandable that the notion of a Turing machine might be seen as just a way of making this materialist idea precise. Understandable, but hardly well thought out. The problem is the following: a 'machine' in the sense of a physical system obeying the laws of Newtonian physics need not be a Turing machine. [Putnam, 1992, 4]

thus: $H(x,y) = 1$ if and only if the x^{th} Turing machine eventually halts if set in motion with the integer y inscribed on its tape, say in binary code (think of y as being the machine's input); and $H(x, y)= 0$ otherwise.) The new fundamental process is carried out by a black box called an 'oracle'. When a function g is computable by an *o*-machine whose oracle produces the values of function f, then g is sometimes said to be computable *relative* to f.

How is the *o*-machine organized? As in the case of the ordinary Turing machine, the behaviour of the *o*-machine is governed by a table of instructions (program). The table provides an exhaustive specification of which fundamental processes the machine is to perform when in such-and-such internal state with such-and-such symbol in the scanner. Let p be the new fundamental process. The machine inscribes the sequence of symbols that is to be input into p on any convenient block of squares of its tape, using some symbol to mark the beginning and end of the sequence. p is called into action by means of a special internal state χ. When an instruction in the program puts the machine into state χ, the marked sequence is delivered to the subdevice that effects p, which then returns the corresponding value of the function, 0 or 1 (Turing considered two-valued functions). (On Turing's way of handling matters, the value is not written on the tape; a pair of states is employed in order to record values of the function. Thus a call to p ends with a subdevice placing the machine in one or other of these two states, according to whether the value of the function is 1 or 0.)

Turing's description of *o*-machines is entirely abstract; he introduced them in order to exhibit an example of a certain type of mathematical problem (the section in which he introduces *o*-machines is entitled 'A type of problem which is not number-theoretic'). In his 1936 paper he exhibited a problem that cannot be solved by effective means, the problem of determining, given any Turing machine, whether or not it prints infinitely many binary digits. All problems equivalent to this one he termed 'number-theoretic' (noting that he was using the term 'number-theoretic' in a 'rather restricted sense') [1939, 152–6]. The *o*-machine concept enabled him to describe a new type of problem, not solvable by a uniform process even with the help of a number-theoretic oracle.

The class of machines whose oracles solve number-theoretic problems is, Turing showed, subject to the same diagonal argument that he used in the 1936 paper: the problem of determining, given any such machine, whether or not it prints infinitely many binary digits is not one that can be solved by any machine in the class and, therefore, is not number-theoretic [1939, 157]. By means of his diagonal argument of 1936 Turing defeated the 'Hilbert program' — the attempt to replace all the intuitive judgments of mathematics by a finite number of rules. By means of this later diagonal argument concerning *o*-machines, Turing also defeated a liberalised form of the Hilbert program, whereby mathematical intuition is replaced not by a finite number of rules but by a specific oracle. (See further [Copeland, 2004c].)

Is an *o*-machine really a *machine*? Those who hold that Turing machines provide a maximally general 'notion of mechanism' (e.g. Newell, quoted below in Section 8.3) will say not. In earlier work we have argued that an *o*-machine is a machine in

the sense of 'machine' crucial to the historical debate between mechanists and anti-mechanists about physiological and psychological mechanism (a debate involving such figures as Descartes, Hobbes, and de la Mettrie) [Copeland, 2000]. The core of the claim, as put forward by the historical mechanists, that such-and-such naturally occurring item — a living body, say — is a machine is this: the item's operation can be accounted for in monistic, materialist terms and in a manner analogous to that in which the operation of an artefact, such as a clockwork figure or church organ, is explained in terms of the nature and arrangement of its components.[39] The *o*-machine, like the universal Turing machine, is a machine in the sense that its behaviour is the product of the nature and arrangement of its material parts. Moreover, the claim that the mind is a machine, in the sense of 'machine' used by the historical mechanists, is evidently consistent with the hypothesis that the mind is a form of *o*-machine.

For Turing an oracle was simply 'some unspecified means of solving number-theoretic problems' and he did 'not go any further into the nature' of an oracle [1939, 156]. One way to conceive of an oracle — perhaps the simplest — is as a device accessing an infinite internal tape upon which there have been inscribed, in order, all the infinitely many arguments and values of whatever function it is that the oracle generates. This device can produce any of the function's values after only a finite search along the tape. The literature now contains various other conceptualisations of an oracle. There is considerable diversity among them, some being more physical in flavour than others. To mention only a few examples (there is a fuller survey in [Copeland, 2002b]), an oracle is in principle realisable by: certain classical electrodynamical systems; the physical process of equilibration; an automaton travelling through relativistic spacetime; a quantum mechanical computer; an inter-neural connection; and a temporally evolving sequence of Turing machines, representing for example a learning mind.

The concept of oracular computation can profitably be employed in theorizing about hardware, about brains, and about minds. Turing, more than anyone else, is to be thanked for uniting historical mechanism with modern mathematics. He enriched mechanism with an abstract theory of (information-processing) machines, presenting us with an indefinitely ascending hierarchy of possible machines, of which the Turing machines form the lowest level. His work posed a new question: if the mind is a machine, where in the hierarchy does it lie?

8.2 The Church-Turing Fallacy

To commit what we call the *Church-Turing fallacy* [Copeland, 1998; 2000] is to believe that the Church-Turing thesis, or some formal or semi-formal result established by Turing or Church, secures the following proposition: *if the mind-brain is a*

[39] Bechtel and Richardson speak aptly of the mechanists' twin strategies of *decomposition* and *localisation* [1993, 23]. The former seeks to decompose the activity of the system whose functioning is to be explained into a number of subordinate activities; the latter attributes these subordinate activities to specific components of the system.

machine, then the Turing-machine computable functions provide sufficient mathematical resources for a full account of human cognition. (The Church-Turing thesis is also known as 'Church's thesis' and 'Turing's thesis'.)

Searle believes it follows from the Church-Turing thesis that the activity of the brain can be simulated by a universal Turing machine:

> Can the operations of the brain be simulated on a digital computer [i.e. Turing machine]? ... [G]iven Church's thesis that anything that can be given a precise enough characterization as a set of steps can be simulated on a digital computer, it follows trivially that the question has an affirmative answer. ([1992, 200–201]; see also [Searle, 1997, 87].)

Searle's statement of the Church-Turing thesis is mistaken. As mentioned above, the Church-Turing thesis is a proposition concerning the extent of what can be achieved by a human mathematician who is unaided by any machinery save paper and pencil and who is working in accordance with an 'effective' method (i.e. a method set out in the form of a finite number of exact instructions that call for no insight or ingenuity on the part of the person carrying them out).

The Church-Turing thesis states that whatever can be calculated by a mathematician so working, even a mathematician idealised to the extent of being free of all constraints on time, patience, concentration, and so forth, can also be calculated by a Turing machine. This thesis carries no implication concerning the extent of what can be calculated by a *machine* (say one that operates in accordance with a finite program of instructions), for among the machine's repertoire of primitive operations there may be those that no human being unaided by machinery can perform. Nor does the thesis imply that every process admitting of a precise characterisation 'as a set of steps' can be simulated by a Turing machine, for the steps need not be ones that a human mathematician working in accordance with some effective method can carry out. Trivially, the processing of an *o*-machine is always characterisable as a set of steps, namely the set of steps specified by the machine's program.

Employing Searle's interpretation (in the above quotation) of Church's thesis yields the absurdity that an *o*-machine can be simulated by a Turing machine. Searle's attempt to recruit Church's thesis in support of his view about the brain is fallacious.

Searle's mistake concerning the extent of what is asserted by the Church-Turing thesis is, in fact, a common one, frequently encountered in recent writing on the philosophy of mind. For example, the Churchlands, like Searle, purport to deduce from the Church-Turing thesis that the mind-brain can in principle be simulated by a Turing machine [Churchland and Churchland, 1983, 6]. They further assert that Turing's

> results entail something remarkable, namely that a standard digital computer, given only the right program, a large enough memory and sufficient time, can compute *any* rule-governed input-output function.

> That is, it can display any systematic pattern of responses to the environment whatsoever. [1990, 26]

Turing had no results entailing this. Rather, he had a result that entails the opposite. The various functions that, in 1936, he proved to be not Turing-machine-computable, are mathematical characterisations of patterns of responses to the environment — perfectly systematic patterns — that cannot be displayed by a standard digital computer (even one unfettered by resource constraints).

Here is a third example of the fallacy at work, this time from Johnson-Laird:

> If you assume that [consciousness] is scientifically explicable ... [and] [g]ranted that the [Church-Turing] thesis is correct, then the final dichotomy rests on ... functionalism. If you believe [functionalism] to be false ... then ... you hold that consciousness could be modelled in a computer program in the same way that, say, the weather can be modelled ... If you accept functionalism, however, then you should believe that consciousness is a computational process. [Johnson-Laird, 1987, 252]

No less common in the literature are statements which, while not explicitly involving the Church-Turing fallacy, make exaggerated claims on behalf of Turing machines. For example, the entry on Turing in Guttenplan's 'A Companion to the Philosophy of Mind' contains the following assertions: 'we can depend on there being a Turing machine that captures the functional relations of the brain' because so long as 'these relations between input and output are functionally well-behaved enough to be describable by ... mathematical relationships . . . we know that some specific version of a Turing machine will be able to mimic them' [Guttenplan, 1994, 595]. Also typical are the following statements, the first by Fodor and the second by Dreyfus:

> If a mental process can be functionally defined as an operation on symbols, there is a Turing machine capable of carrying out the computation. ([Fodor, 1981, 130]; see also [Fodor, 1983, 38–39].)

> *any* process which can be formalised so that it can be represented as a series of instructions for the manipulation of discrete elements can, at least in principle, be reproduced by [a universal Turing machine]. [Dreyfus, 1992, 72]

Closely related to the Church-Turing fallacy is the *equivalence* fallacy [Copeland, 2000]. Paramount among the evidence for the Church-Turing thesis properly so-called is the fact that all attempts to give an exact analysis of the intuitive notion of a human clerical method have turned out to be *equivalent*, in the sense that each analysis has been proved to pick out the same class of functions, namely those that are computable by Turing machine. (For example, there have been analyses in terms of lambda-definability, recursiveness, Post canonical systems, and Markov algorithms.) Because of the diversity of these various analyses, their equivalence is

generally considered very strong evidence for the Church-Turing thesis (although for a sceptical point of view see [Kreisel, 1965, 144]). However, the equivalence of these diverse analyses is sometimes taken to be evidence also for stronger theses such as: *all functions that can be generated by machines (working on finite input in accordance with a finite program of instructions) are computable by Turing machine*. This is nothing more than a confusion — the equivalence fallacy. The analyses under discussion are of the notion of a human clerical method, not of the notion of a machine-generable function; the equivalence of the analyses bears only on the issue of the extent of the former notion and indicates nothing concerning the extent of the latter.

8.3 Artificial Intelligence and the Equivalence Fallacy

Newell, discussing the possibility of artificial intelligence, argues that (what he calls) a 'physical symbol system' can be organized to exhibit general intelligence [Newell, 1980]. A 'physical symbol system' is a universal Turing machine, or any equivalent system, situated in the physical — as opposed to the conceptual — world. (The tape of the machine is accordingly finite; Newell specifies that the storage capacity of the tape (or equivalent) be unlimited in the practical sense of finite yet not small enough to 'force concern'.)

> A [physical symbol] system always contains the potential for being any other system if so instructed. Thus, a [physical symbol] system can become a generally intelligent system. [Newell, 1980, 170]

Is the premiss of this pro-AI argument true? A physical symbol system, being a *universal* Turing machine situated in the real world, can, if suitably instructed, simulate (or, metaphorically, become) any other physical symbol system (modulo some fine print concerning storage capacity). If this is what the premiss means, then it is true. If, on the other hand, the premiss is taken literally, then it is false, since, as previously remarked, systems can be specified which no Turing machine — and so no physical symbol system — can simulate. The rub is that, if the premiss is interpreted in the former manner, so that it is true, the conclusion fails to follow. Only to one who believes, as Newell does, that 'the notion of machine or determinate physical mechanism' is 'formalised' by the notion of a Turing machine [ibid.] will the argument appear deductively valid.

Newell's defence of his view that the universal Turing machine exhausts the possibilities of mechanism involves an example of the equivalence fallacy:

> [An] important chapter in the theory of computing ... has shown that all attempts to ... formulate ... general notions of mechanism ... lead to classes of machines that are equivalent in that they encompass in toto exactly the same set of input-output functions. In effect, there is a single large frog pond of functions no matter what species of frogs (types of machines) is used. ... A large zoo of different formulations of maximal classes of machines is known by now —

> Turing machines, recursive functions, Post canonical systems, Markov algorithms ... [Newell, 1980, 150]

Newell's a priori argument for the claim that a physical symbol system can become generally intelligent founders in confusion.

8.4 The Church–Turing Fallacy and Formal Symbol Manipulation

Even if some version of the Chinese Room argument were sound (Section 7), the argument could not possibly establish — as Searle claims it does — his key thesis that whatever is 'purely formal' or 'syntactical' is neither constitutive of nor sufficient for mind. Searle has repeatedly emphasised that it is the fact that computers have 'no more than just formal symbols' which entails that programs 'cannot constitute the mind' and that this entailment is demonstrated by the Chinese room argument ([1989, 33]; [1992, 200]):

> The whole point of the original [i.e. indoor] example was to argue that ... symbol manipulation by itself couldn't be sufficient for understanding Chinese. [1980, 419]

> [F]ormal syntax ... does not by itself guarantee the presence of mental contents. I showed this a decade ago in the Chinese room argument. [1992, 200]

However, *o*-machines point up the fact that the notion of a programmed machine whose activity consists of the manipulation of formal symbols is more general than the notion of a universal Turing machine. If there is an implication from 'x's operation is defined purely formally or syntactically' to 'x's operation is neither constitutive of nor sufficient for mind', it is not one that could possibly be established by the Chinese room argument.

9 PENROSE'S 'GÖDEL OBJECTION' TO AI

Roger Penrose argues, on the basis of formal results due to Gödel and Turing, that the brain neither is, nor can be perfectly simulated by, a Turing machine [1989; 1990; 1994]. This form of objection to the mechanical view of mind was anticipated by Emil Post as early as 1921.[40] John Lucas made the objection well-known in a famous article [1961]. Turing was aware of the objection, which he dubbed the 'Mathematical Objection' [1950a, 450]. In 'Intelligent Machinery' Turing formulated the objection as follows:

> Recently the theorem of Gödel and related results ... have shown that if one tries to use machines for such purposes as determining the truth or falsity of mathematical theorems and one is not willing to tolerate

[40] Post, E. L. 'Absolutely Unsolvable Problems and Relatively Undecidable Propositions: Account of an Anticipation', p. 417. In [Davis, 1965].

> an occasional wrong result, then any given machine will in some cases be unable to give an answer at all. On the other hand the human intelligence seems to be able to find methods of ever-increasing power for dealing with such problems 'transcending' the methods available to machines. [1948, 410–11]

Penrose notes that his formulation of the Gödel argument can be extended to apply to oracle machines:

> the arguments of Part I of this book can be applied equally well against an oracle-machine model of mathematical understanding as they were against the Turing-machine model, almost without change. ([1994, 380]; see also [Penrose, 1996, sects 3.10, 13.2].)

There is, as Penrose is clearly aware, a whiff of reductio ad absurdum about this. Let the *first-order* o-machines be those whose oracle returns the values of the Turing-machine halting function $H(x,y)$. The second-order o-machines are those with an oracle that can say whether or not any given first-order o-machine eventually halts if set in motion with such-and-such a number inscribed on its tape; and so on for third-order, and in general α-order. Penrose's argument was originally marketed as demonstrating that human mathematicians do not use a knowably sound Turing-machine algorithm in order to ascertain mathematical truth (e.g. in [1994, ch. 2]). But the argument appears to be so powerful that it can equally well be employed to show that human mathematicians do not use any knowably sound procedure capable of being executed by an α-order oracle machine (given any number α at all for which there is a notation) in order to ascertain mathematical truth. Penrose's argument moves relentlessly up through the orders, stopping nowhere. This discovery evidently disconcerts Penrose:

> The final conclusion of all this is rather alarming. For it suggests that we must seek a non-computable physical theory that reaches beyond every [recursive] level of oracle machines (and perhaps beyond). No doubt there are readers who believe that the last vestige of credibility of my argument has disappeared at this stage! I certainly should not blame any reader for feeling this way. [1994, 381]

Penrose does hint at a way out:

> [I]t need not be the case that human mathematical understanding is in principle as powerful as *any* oracle machine at all. ... [T]he conclusion [that human mathematicians are not using a knowably sound algorithm in order to ascertain mathematical truth] does *not* necessarily imply that human insight is powerful enough, in principle, to solve each instance of the halting problem. Thus, we need not necessarily conclude that the physical laws that we seek reach, in principle, beyond every computable level of oracle machine (or even reach the first order). We need only seek something that is not equivalent to *any* specific oracle

> machine (including also the zero*th*-order machines, which are Turing machines). Physical laws could perhaps lead to something that is just *different*. [1994, 381]

On that enigmatic note Penrose leaves it. Just when we seem to be approaching a crucial part of the exposition, he suddenly falls silent. What, indeed, does Penrose mean in the last two sentences of this passage?

It is natural to think of the functions, or problems, that are solvable by a first-order oracle machine as being *harder* than those solvable by Turing machine, and those solvable by a second-order oracle machine as being harder still, and so forth. (To say that a Turing machine or *o*-machine 'solves a problem' is to say that, when the machine is given the problem (suitably encoded) on its tape, it halts with the answer, 1 (Yes) or 0 (No), under its scanner.) It is customary in recursion theory to say that a class of problems of equal hardness are of the same *degree*. Problems that are solvable by Turing machine are said to be of degree 0. We shall write 1 for the degree of problems that are solvable by a first-order oracle machine (but not by Turing machine). It is known that there are degrees *between* 0 and 1 ([Friedberg, 1957]; [Sacks, 1964]; [Simpson, 1977] is a survey of the area). That is to say, there are classes of problems that are too hard to be solved by Turing machine and yet are less hard than some of the problems that a first-order oracle machine can solve. Here 'less hard' has the precise sense that while a first-order machine can solve any of the problems in such an intermediate class, an *o*-machine that is equipped only with an oracle for delivering the solutions ('Yes' or 'No') to problems in that class is unable to solve every problem that the first-order machine can solve. This notion of degrees lying between 0 and 1 seems to make sense of much of what Penrose says about what it is that he seeks (although certainly not of the specific claim that 'it need not be the case that human mathematical understanding is in principle as powerful as *any* oracle machine at all'). For some degree between 0 and 1, the 'physics of mind' is exactly that hard. This is certainly a coherent position; and for all that anyone presently knows, such may in fact be the case.

However, the *reductio* threat to Penrose's position now becomes extreme. Let i (for 'intermediate') be an *o*-machine of the sort just described (any arbitrarily selected one), and let I be the set of all machines with the same oracle as i. Do mathematicians use, in ascertaining mathematical truth, a knowably sound procedure that can be executed by a machine in I? Apparently not, if Penrose's Gödelian argument is sound. For just as the argument can be pressed to yield the conclusion that mathematicians do not use a knowably sound procedure that can be executed by an α-order *o*-machine, it can equally well be pressed to show the same regarding machines in I. The two cases appear to be parallel in all relevant respects; in particular, it can be shown that there is no machine in I that can calculate all values of the halting function for machines in I (i.e. the function whose definition is just like that of $H(x,y)$ given above, except that the phrase 'the x^{th} Turing machine' is replaced by 'the x^{th} machine in I'). So what knowably sound procedure *can* mathematicians be supposed to use (recall that i was arbitrarily selected)? It is by no means clear how an upholder of the Gödelian

argument might respond to this difficulty.

Penrose has failed to make it clear what scientific conception of the mind can remain for one who endorses the Gödel argument. Lucas was happy to conclude from the Gödel argument that 'no scientific enquiry can ever exhaust the ... human mind' [1961, 127]. Penrose, it seems, wants there to be a fully scientific conception of the mind. But in response to the fact that the crucial step of the Gödel argument 'can be applied in very general circumstances indeed', he says only that the mind is 'something very mysterious' [Penrose, 1996, sect. 13.2].

9.1 Turing's Answer to the Mathematical Objection

> The short answer to this argument is that although it is established that there are limitations to the powers of any particular machine, it has only been stated, without any sort of proof, that no such limitations apply to the human intellect. [Turing, 1950a, 451]

This remark might appear to cut to the heart of the matter. However, Turing expresses dissatisfaction with it, saying that the Mathematical Objection cannot 'be dismissed quite so lightly' [ibid.]. He goes on to broach a further line of attack on the argument, pointing out that humans 'often give wrong answers to questions' [ibid.]. It is this line of attack that he pursues in his lecture 'Intelligent Machinery, A Heretical Theory' [Turing, c.1951].

In his 1948 statement of the Mathematical Objection Turing notes that the objection rests on a proviso (quoted at the beginning of Section 9) that the machine is not allowed to make mistakes and that 'the condition that the machine must not make mistakes ... is not a requirement for intelligence' [1948, 411]. Turing envisages machines no more limited than the human intellect:

> My contention is that machines can be constructed which will simulate the behaviour of the human mind very closely. They will make mistakes at times, and at times they may make new and very interesting statements ... [c.1951, 473]

The use of heuristic search carries with it the risk of the computer producing a proportion of incorrect answers. This fact would have been very familiar to Turing from his experience with the Bombe. Probably Turing was thinking of heuristic search when he said:

> There are indications however that it is possible to make the machine display intelligence at the risk of its making occasional serious mistakes. By following up this aspect the machine could probably be made to play very good chess. ([1948, 389]; see Section 1 above.)

In 'Intelligent Machinery, A Heretical Theory', Turing suggests further that the 'danger of the mathematician making mistakes is an unavoidable corollary of his power of sometimes hitting upon an entirely new method' [c.1951, 472]. In this

lecture Turing passes immediately from his remarks on the Mathematical Objection to a discussion of machine learning. This juxtaposition perhaps indicates that Turing's view was this: it is the possibility of a machine's learning *new* methods and techniques that ultimately defeats the Mathematical Objection.

How may new methods be learned? In the simplest possible case, akin to classroom learning, the machine's tutor (a human mathematician, say) will present the machine with a better method whenever the machine's existing methods fail to produce the correct answer to a problem. This new input in effect alters the machine's program, transforming the machine into a different Turing machine. More speculatively, the machine may itself be able to search around for better methods, with the 'unavoidable corollary' that it may sometimes make mistakes. (Turing emphasised that the machine's search might involve the use of a random element.) As in the preceding case, the machine's program is altered in consequence of the learning process, as the machine updates or overwrites its previous methods of proof.

Thus the learning machine successively mutates from one Turing machine into another, becoming capable of dealing with increasingly many mathematical problems as new methods of proof, of 'ever-increasing power', are acquired. The trajectory of the learning machine through the space of Turing machines might indeed be uncomputable, in the sense that the function on the non-negative integers whose value at i is the i^{th} Turing machine on the trajectory need not be computable by the universal Turing machine.

Turing: 'In short, then, there might be men cleverer than any given machine, but then again there might be other machines cleverer again, and so on.' [1950a, 451]

BIBLIOGRAPHY

[Aleksander and Burnett, 1987] I. Aleksander and P. Burnett. *Thinking Machines: The Search for Artificial Intelligence*. Oxford: Oxford University Press, 1987.

[Bechtel and Richardson, 1993] W. Bechtel and R. C. Richardson. *Discovering Complexity: Decomposition and Localization as Strategies in Scientific Research*. Princeton: Princeton University Press, 1993.

[Beurle, 1957] R. L. Beurle. Properties of a Mass of Cells Capable of Regenerating Pulses. *Philosophical Transactions of the Royal Society of London*, Series B, 240: 55–94, 1957.

[Block, 1981] N. Block. Psychologism and Behaviorism. *Philosophical Review*, 90: 5–43, 1981.

[Block, 1990] N. Block. The Computer Model of the Mind. In D. N. Osherson, H. Lasnik (eds.), *An Invitation to Cognitive Science*, Vol. 3. Cambridge, Mass.: MIT Press, 1990.

[Bowden, 1953] B. V. Bowden (ed.). *Faster Than Thought*. London: Sir Isaac Pitman & Sons, 1953.

[Braitenberg, 1984] V. Braitenberg. *Vehicles: Experiments in Synthetic Psychology*. Cambridge, Mass.: MIT Press, 1984.

[Breazeal, 1998] C. Breazeal. Regulating Human-Robot Interaction using 'Emotions', 'Drives', and Facial Expressions. *Proceedings of the 1998 Autonomous Agents Workshop: Agents in Interaction — Acquiring Competence Through Imitation*, 14–21, 1998.

[Breazeal *et al.*, 2005] C. Breazeal, D. Buchsbaum, J. Gray, D. Gatenby and B. Blumberg. Learning From and About Others: Towards Using Imitation to Bootstrap the Social Understanding of Others by Robots. *Artificial Life*, 11: 31–62, 2005. `http://robotic.media.mit.edu/Papers/Breazeal-etal-AL04.pdf`.

[Breazeal and Fitzpatrick, 2000] C. Breazeal and P. Fitzpatrick. That Certain Look: Social Amplification of Animate Vision. *Proceedings of the AAAI Fall Symposium, Socially Intelligent Agents — The Human in the Loop*, 2000. `http://people.csail.mit.edu/paulfitz/pub/AAAIFS00.pdf`.

[Breazeal and Scassellati, 2000] C. Breazeal and B. Scassellati. Infant-like Social Interactions Between a Robot and a Human Caretaker. *Adaptive Behavior*, 8: 49–74, 2000. `http://www.ai.mit.edu/projects/ntt/projects/NTT9904-01/documents/kismet.pdf`.

[Breazeal and Velásquez, 1998] C. Breazeal and J. Velásquez. Toward Teaching a Robot "Infant" using Emotive Communication Acts. In B. Edmonds, K. Dautenhahn (eds.), *Socially Situated Intelligence*. University of Zurich Technical Report, 1998.

[Brooks, 1990] R. A. Brooks. Elephants Don't Play Chess. *Robotics and Autonomous Systems*, 6: 3–15, 1990.

[Brooks, 1991] R. A. Brooks. Intelligence without Representation. *Artificial Intelligence*, 47: 139–159, 1991.

[Brooks, 1994] R. A. Brooks. Coherent Behaviour from Many Adaptive Processes. In [Cliff *et al.*, 1994].

[Brooks, 1995] R. A. Brooks. Intelligence Without Reason. In L. Steels, R. A. Brooks (eds.), *The Artificial Life Route to Artificial Intelligence*. Hillsdale, NJ: Erlbaum, 1995.

[Brooks, 1999] R. A. Brooks. *Cambrian Intelligence: the Early History of the New AI*. Cambridge, Mass.: MIT Press, 1999.

[Brooks *et al.*, 1998] R. A. Brooks, C. Breazeal, R. Irie, C. C. Kemp, M. Marjanović, B. Scassellati and M. M. Williamson. Alternative Essences of Intelligence. *Proceedings of the Fifteenth National Conference on Artificial Intelligence (AAAI-98)*. Madison, WI: AAAI Press, 1998.

[Brooks and Stein, 1994] R. A. Brooks and L. A. Stein. Building Brains for Bodies. *Autonomous Robots*, 1: 7–25, 1994.

[Bunge and Ardila, 1987] M. Bunge and R. Ardila. *Philosophy of Psychology*. New York: Springer-Verlag, 1987.

[Campbell-Kelly, 1985] M. Campbell-Kelly. Christopher Strachey, 1916-1975: A Biographical Note. *Annals of the History of Computing*, 7: 19–42, 1985.

[Church, 1937] A. Church. Review of Turing (1936), *Journal of Symbolic Logic*, 2: 42–43, 1937.

[Churchland and Churchland, 1983] P. M. Churchland and P. S. Churchland. Stalking the Wild Epistemic Engine, *Nous*, 17: 5–18, 1983.

[Churchland and Churchland, 1990] P. M. Churchland and P. S. Churchland. Could a Machine Think? *Scientific American*, 262 (1): 26–31, 1990.

[Clark, 1997] A. Clark. *Being There*. Cambridge, Mass.: MIT Press, 1997.

[Clark and Farley, 1955] W. A. Clark and B. G. Farley. Generalization of Pattern Recognition in a Self-Organizing System. *Proceedings of the Western Joint Computer Conference*, 1955.

[Cliff *et al.*, 1994] D. Cliff, P. Husbands, J.-A. Meyer and S. Wilson (eds.). *From Animals to Animats: Proceedings of the Third International Conference on Similarities of Adaptive Behavior (SAB-94)*. Cambridge, Mass.: MIT Press, 1994.

[Copeland, 1993a] B. J. Copeland. *Artificial Intelligence: A Philosophical Introduction*. Oxford: Blackwell, 1993.

[Copeland, 1993b] B. J. Copeland. The Curious Case of the Chinese Gym. *Synthese*, 95: 173–186, 1993.

[Copeland, 1997a] B. J. Copeland. The Church-Turing Thesis. In E. Zalta (ed.), *Stanford Encyclopedia of Philosophy*, 1997. `http://plato.stanford.edu`.

[Copeland, 1997b] B. J. Copeland. The Broad Conception of Computation, *American Behavioral Scientist*, 40: 690–716, 1997.

[Copeland, 1998] B. J. Copeland. Turing's O-machines, Penrose, Searle, and the Brain, *Analysis*, 58: 128–138, 1998.

[Copeland, 2000] B. J. Copeland. Narrow Versus Wide Mechanism, *Journal of Philosophy*, 96: 5–32, 2000.

[Copeland, 2002a] B. J. Copeland. The Chinese Room from a Logical Point of View. In J. Preston, M. Bishop (eds.), *Views into the Chinese Room: New Essays on Searle and Artificial Intelligence*. Oxford: Oxford University Press, 2002.

[Copeland, 2002b] B. J. Copeland. Hypercomputation. In [Copeland, 2002-3].

[Copeland, 2002-3] B. J. Copeland (ed.). *Hypercomputation*. Special issue of *Minds and Machines*, Vols 12(4), 13(1), 2002-3.

[Copeland, 2003] B. J. Copeland. The Turing Test. In [Moor, 2003a].

[Copeland, 2004a] B. J. Copeland, ed. *The Essential Turing*. Oxford: Oxford University Press, 2004.

[Copeland, 2004b] B. J. Copeland. Hypercomputation: Philosophical Issues. *Theoretical Computer Science*, 317: 251–267, 2004.

[Copeland, 2004c] B. J. Copeland. Introduction to "Systems of Logic Based on Ordinals". In [Copeland, 2004a].

[Copeland, 2005a] B. J. Copeland (ed.). *Alan Turing's Automatic Computing Engine: The Master Codebreaker's Struggle to Build the Modern Computer*. Oxford: Oxford University Press, 2005.

[Copeland, 2005b] B. J. Copeland. The Origins and Development of the ACE Project. In [Copeland, 2005].

[Copeland *et al.*, 2006] B. J. Copeland *et al.*. *Colossus*. Oxford: Oxford University Press, 2006.

[Copeland and Proudfoot, 1996] B. J. Copeland and D. Proudfoot. On Alan Turing's Anticipation of Connectionism, *Synthese*, 108: 361–377, 1996.

[Copeland and Proudfoot, 1999] B. J. Copeland and D. Proudfoot. Alan Turing's Forgotten Ideas in Computer Science, *Scientific American*, 280 (4): 99–103, 1999.

[Copeland and Proudfoot, 2005] B. J. Copeland and D. Proudfoot. Turing and the Computer. In [Copeland, 2005a].

[Copeland and Proudfoot, forthcoming] B. J. Copeland and D. Proudfoot. The Turing Test. In R. Epstein, G. I., Peters (eds.), *The Turing Test Sourcebook*. Dordrecht: Kluwer, forthcoming.

[Copeland and Sylvan, 1999] B. J. Copeland and R. Sylvan. Beyond the Universal Turing Machine. *Australasian Journal of Philosophy*, 77: 46–66, 1999.

[da Costa and Doria, 1991] N. C. A. da Costa and F. A. Doria. Classical Physics and Penrose's Thesis. *Foundations of Physics Letters*, 4: 363–373, 1991.

[Davis, 1965] M. Davis, ed. *The Undecidable: Basic Papers On Undecidable Propositions, Unsolvable Problems And Computable Functions*. New York: Raven, 1965.

[Dennett, 1994] D. C. Dennett. The Practical Requirements for Making a Conscious Robot. *Philosophical Transactions of the Royal Society of London*, Series A, 349: 133–146, 1994.

[Dennett, 1997] D. C. Dennett. When HAL Kills, Who's to Blame? Computer Ethics. In D. G. Stork (ed.), *Hal's Legacy: 2001's Computer as Dream and Reality*. Cambridge, Mass.: MIT Press, 1997.

[Dennett, 2000] D. C. Dennett. The Case for Rorts. In R. Brandom (ed.), *Rorty and his Critics*. Oxford: Blackwell, 2000.

[DeJong, 1977] G. DeJong. Skimming Newspaper Stories by Computer. *Proceedings of the Fifth International Joint Conference on Artificial Intelligence*. Cambridge, Mass.: Massachusetts Institute of Technology, 1977.

[Dreyfus, 1992] H. L. Dreyfus. *What Computers* Still *Can't Do*. Cambridge, Mass.: MIT Press, 1992.

[Dreyfus and Dreyfus, 1986] H. L. Dreyfus and S. E. Dreyfus. *Mind Over Machine*. Oxford: Blackwell, 1986.

[Doyle, 1982] J. Doyle. What is Church's Thesis? Laboratory for Computer Science, MIT, 1982. Reprinted in [Copeland, 2002-3].

[Epstein, 1992] R. Epstein. The Quest for the Thinking Computer. *AI Magazine*, 13(2): 81–95, 1992.

[Ernst and Newell, 1969] G. W. Ernst and A. Newell. *GPS: A Case Study in Generality and Problem Solving*. New York: Academic Press, 1969.

[Farley and Clark, 1954] B. G. Farley and W. A. Clark. Simulation of Self-Organizing Systems by Digital Computer. *Institute of Radio Engineers Transactions on Information Theory*, 4: 76–84, 1954.

[Feigenbaum *et al.*, 1971] E. A. Feigenbaum, B. G. Buchanan and J. Lederberg. On Generality and Problem Solving: A Case Study Using the DENDRAL Program. In B. Meltzer, D. Michie (eds.), *Machine Intelligence 6*. Edinburgh: Edinburgh University Press, 1971.

[Feigenbaum and Feldman, 1963] E. A. Feigenbaum and J. Feldman (eds.). *Computers and Thought*. New York: McGraw-Hill, 1963.

[Fodor, 1981] J. A. Fodor. The Mind-Body Problem. *Scientific American*, 244 (1): 124–132, 1981.

[Fodor, 1983] J. A. Fodor. *The Modularity of Mind*. Cambridge, Mass.: MIT Press, 1983.

[French, 1990] R. French. Subcognition and the Limits of the Turing Test. *Mind* 99: 53–65, 1990. Reprinted in [Millican and Clark, 1996]. Page references are to [Millican and Clark, 1996].

[French, 2000] R. French. The Turing Test: the First 50 Years. *Trends in Cognitive Science*, 4: 115–122, 2000.

[Friedberg, 1957] R. M. Friedberg. Two Recursively Enumerable Sets of Incomparable Degrees of Unsolvability (Solution of Post's Problem, 1944). *Proceedings of the National Academy of Sciences*, 43: 236–238, 1957.

[Geroch and Hartle, 1986] R. Geroch and J. B. Hartle. Computability and Physical Theories. *Foundations of Physics*, 16: 533–550, 1986.

[Guttenplan, 1994] S. Guttenplan. *A Companion to the Philosophy of Mind*. Oxford: Blackwell, 1994.

[Hara, 2004] F. Hara. Artificial Emotion of Face Robot Through Learning in Communicative Interactions with Humans. *Proceedings of the 2004 IEEE International Workshop on Robot and Human Interactive Communication*, 7–15, 2004.

[Haugeland, 1985] J. Haugeland. *Artificial Intelligence: The Very Idea*. Cambridge, Mass.: MIT Press, 1985.

[Hebb, 1949] D. O. Hebb. *The Organization of Behavior: A Neuropsychological Theory*. New York: John Wiley, 1949.

[Heiser *et al.*, 1980] J. F. Heiser, K. M. Colby, W. S. Faught and R. C. Parkison. Can Psychiatrists Distinguish a Computer Simulation of Paranoia from the Real Thing? *Journal of Psychiatric Research*, 15: 149–162, 1980.

[Hendriks-Jansen, 1996] H. Hendriks-Jansen. *Catching Ourselves in the Act: Situated Activity, Interactive Emergence, Evolution, and Human Thought*. Cambridge, Mass.: MIT Press, 1996.

[Hodges, 1992] A. Hodges. *Alan Turing: The Enigma*. London: Vintage, 1992.

[Hogg *et al.*, 1991] D. W. Hogg, F. Martin and M. Resnick. Braitenberg Creatures. *Epistemology and Learning Memo #13*, MIT Media Laboratory, 1991.

[Holland, 1992] J. H. Holland. *Adaptation in Natural and Artificial Systems*. Cambridge, Mass.: MIT Press, 1992.

[Holland, 1997] O. Holland. Grey Walter: The Pioneer of Real Artificial Life. In C. G. Langton and K. Shimohara (eds.), *Artificial Life V*. Cambridge, Mass.: MIT Press, 1997.

[Hummel and Biederman, 1992] J. E. Hummel and I. Biederman. Dynamic Binding in a Neural Network for Shape Recognition. *Psychological Review*, 99: 480–517, 1992.

[Johnson-Laird, 1987] P. Johnson-Laird. How Could Consciousness Arise from the Computations of the Brain? In C. Blakemore and S. Greenfield (eds.), *Mindwaves*. Oxford: Blackwell, 1987.

[Komar, 1964] A. Komar. Undecidability of Macroscopically Distinguishable States in Quantum Field Theory. *Physical Review*, second series, 133B: 542–544, 1964.

[Kreisel, 1965] G. Kreisel. Mathematical Logic. In T. L. Saaty (ed.), *Lectures on Modern Mathematics*, Vol. 3. New York: Wiley, 1965.

[Kreisel, 1967] G. Kreisel. Mathematical Logic: What has it done for the philosophy of mathematics? In R. Schoenman (ed.), *Bertrand Russell: Philosopher of the Century*. London: George Allen and Unwin, 1967.

[Kreisel, 1974] G. Kreisel. A Notion of Mechanistic Theory. *Synthese*, 29: 11–26, 1974.

[Langton, 1986] C. G. Langton. Studying Artificial Life with Cellular Automata. *Physica D*, 22: 120–149, 1986.

[Lavington, 1975] S. H. Lavington. *A History of Manchester Computers*. Manchester: NCC Publications, 1975.

[Lenat and Guha, 1990] D. B. Lenat and R. V. Guha. *Building Large Knowledge-Based Systems: Representation and Inference in the Cyc Project*. Reading, Mass.: Addison-Wesley, 1990.

[Lenat and Feigenbaum, 1991] D. B. Lenat and E. A. Feigenbaum. On the Thresholds of Knowledge. *Artificial Intelligence*, 47: 185–250, 1991.

[Lucas, 1961] J. R. Lucas. Minds, Machines and Gödel. *Philosophy*, 36: 112–127, 1961.

[Mahon, 1945] P. Mahon. History of Hut 8. 1945, in [Copeland, 2004a].

[Mataric, 1994] M. J. Mataric. Learning to Behave Socially. In [Cliff *et al.*, 1994].

[Mays and Prinz, 1950] W. Mays and D. G. Prinz. A Relay Machine for the Demonstration of Symbolic Logic. *Nature*, 165: 197–198, 1950.

[McCorduck, 1979] P. McCorduck. *Machines Who Think*. New York: W. H. Freeman, 1979.

[McCulloch and Pitts, 1943] W. S. McCulloch and W. Pitts. A Logical Calculus of the Ideas Immanent in Nervous Activity. *Bulletin of Mathematical Biophysics*, 5: 115–133, 1943.

[McLeod and Dienes, 1993] P. McLeod and Z. Dienes. Running to Catch the Ball. *Nature*, 362: 23, 1993.

[Meltzer and Michie, 1969] B. Meltzer and D. Michie (eds.). *Machine Intelligence 5*. Edinburgh: Edinburgh University Press, 1969.

[Michie, 1966] D. Michie. Game-Playing and Game-Learning Automata. In L. Fox (ed.), *Advances in Programming and Non-numerical Computation*. New York: Pergamon, 1966.

[Michie, 1986] D. Michie. *On Machine Intelligence*. 2nd edition. Chichester: Ellis Horwood, 1986.

[Michie, 1993] D. Michie. Turing's Test and Conscious Thought. *Artificial Intelligence*, 60: 1–22, 1993. Reprinted in [Millican and Clark, 1996]. Page references are to [Millican and Clark, 1996].

[Millican and Clark, 1996] P. Millican and A. Clark (eds.). *Machines and Thought: The Legacy of Alan Turing*. Oxford: Oxford University Press, 1996.

[Minsky and Papert, 1969] M. L. Minsky and S. Papert. *Perceptrons: An Introduction to Computational Geometry*. Cambridge, Mass.: MIT Press, 1969.

[Minsky and Papert, 1972] M. L. Minsky and S. Papert. Progress Report on Artificial Intelligence. Memo 252, MIT AI Laboratory, 1972.

[Moor, 1976] J. H. Moor. An Analysis of the Turing Test. *Philosophical Studies*, 30: 249–257, 1976.

[Moor, 1987] J. H. Moor. Turing Test. In [Shapiro, 1987].

[Moor, 2003a] J. H. Moor (ed.). *The Turing Test: The Elusive Standard of Artificial Intelligence*. Dordrecht: Kluwer, 2003.

[Moor, 2003b] J. H. Moor. The Status and Future of the Turing Test. In [Moor, 2003a].

[Newell, 1980] A. Newell. Physical Symbol Systems. *Cognitive Science*, 4: 135–183, 1980.

[Newell *et al.*, 1957] A. Newell, J. C. Shaw and H. A. Simon. Empirical Explorations with the Logic Theory Machine: a Case Study in Heuristics. *Proceedings of the Western Joint Computer Conference*, Vol. 15, 1957.

[Newell and Simon, 1976] A. Newell and H. A. Simon. Computer Science as Empirical Inquiry: Symbols and Search. *Communications of the Association for Computing Machinery*, 19: 113–126, 1976.

[Oettinger, 1952] A. G. Oettinger. Programming a Digital Computer to Learn. *Philosophical Magazine*, 43: 1243–1263, 1952.

[Penrose, 1989] R. Penrose. *The Emperor's New Mind: Concerning Computers, Minds, and the Laws of Physics*. Oxford: Oxford University Press.

[Penrose, 1990] R. Penrose. Précis of *The Emperor's New Mind: Concerning Computers, Minds, and the Laws of Physics*. *Behavioral and Brain Sciences*, 13: 643–655, 692-705, 1990.

[Penrose, 1994] R. Penrose. *Shadows of the Mind: A Search for the Missing Science of Consciousness*. Oxford: Oxford University Press, 1994.

[Penrose, 1996] R. Penrose. Beyond the Doubting of a Shadow, *Psyche*, 2, 1996. `http://psyche.cs.monash.edu.au`.

[Poggio and Beymer, 1996] T. Poggio and D. Beymer. Learning to See. *IEEE Spectrum*, May 1996: 60–69, 1996.

[Pour-El and Richards, 1979] M. B. Pour-El and I. Richards. A Computable Ordinary Differential Equation Which Possesses No Computable Solution. *Annals of Mathematical Logic*, 17: 61–90, 1979.

[Pour-El and Richards, 1981] M. B. Pour-El and I. Richards. The Wave Equation With Computable Initial Data Such That Its Unique Solution Is Not Computable. *Advances in Mathematics*, 39: 215–239, 1981.

[Prinz, 1952] D. G. Prinz. Robot Chess. *Research*, 5: 261–266, 1952.

[Prinz and Smith, 1953] D. G. Prinz and J. B. Smith. Machines for the Solution of Logical Problems. In [Bowden, 1953].

[Putnam, 1992] H. Putnam. *Renewing Philosophy*. Cambridge, Mass.: Harvard University Press.

[Raphael, 1976] B. Raphael. *The Thinking Computer: Mind Inside Matter*. San Francisco: W. H. Freeman, 1976.

[Rejewski, 1980] M. Rejewski. Jak Matematycy polscy rozszyfrowali Enigme [How the Polish mathematicians broke Enigma]. *Annals of the Polish Mathematical Society, Series II: Mathematical News*, 23: 1–28, 1980. (English translation in Appendix D of Kozaczuk, W. *Enigma: How the German Machine Cipher Was Broken, and How It Was Read by the Allies in World War Two*. London: Arms and Armour Press, 1984.)

[Rosenblatt, 1957] F. Rosenblatt. The Perceptron, a Perceiving and Recognizing Automaton. Report No. 85-460-1, Cornell Aeronautical Laboratory, 1957.

[Rosenblatt, 1958] F. Rosenblatt. The Perceptron: a Probabilistic Model for Information Storage and Organisation in the Brain. *Psychological Review*, 65: 386–408, 1958.

[Rosenblatt, 1962] F. Rosenblatt. *Principles of Neurodynamics*. Washington, D.C.: Spartan, 1962.

[Rumelhart *et al.*, 1986] D. E. Rumelhart, G. E. Hinton and R. J. Williams. Learning Representations by Back-propagating Errors. *Nature*, 323(9): 533–536, 1986.

[Rumelhart and McClelland, 1986] D. E. Rumelhart and J. L. McClelland. On Learning the Past Tenses of English Verbs. In J. L. McClelland, D. E. Rumelhart, and the PDP Research Group *Parallel Distributed Processing: Explorations in the Microstructure of Cognition*, Vol. 2: *Psychological and Biological Models*. Cambridge, Mass.: MIT Press, 1986.

[Sacks, 1964] G. E. Sacks. The Recursively Enumerable Degrees are Dense. *Annals of Mathematics*, 2nd series, 80: 300–312, 1964.

[Samuel, 1959] A. L. Samuel. Some Studies in Machine Learning Using the Game of Checkers. *IBM Journal of Research and Development*, 3: 211–229, 1959. Reprinted in [Feigenbaum and Feldman, 1963].

[Scarpellini, 1963] B. Scarpellini. Zwei unentscheitbare Probleme der Analysis. *Zeitschrift für mathematische Logik und Grundlagen der Mathematik*, 9: 265–289, 1963. Published in translation as 'Two Undecidable Problems of Analysis' (with a new commentary by Scarpellini) in [Copeland, 2002-3].

[Scassellati, 1998] B. Scassellati. Imitation and Mechanisms of Shared Attention: A Developmental Structure for Building Social Skills. Presented at Autonomous Agents 1998 workshop *Agents in Interaction: Acquiring Competence through Imitation*, Minneapolis/St. Paul, May 1998. `http://www.ai.mit.edu/projects/cog/Publications/scaz-imitation.pdf`.

[Scassellati, 2000] B. Scassellati. Investigating Models of Social Development Using a Humanoid Robot. In B. Webb, T. Consi (eds.), *Biorobotics*. Cambridge, Mass.: MIT Press, 2000.

[Schank, 1984] R. C. Schank. *The Cognitive Computer*. Reading, Mass.: Addison-Wesley, 1984.

[Schank and Abelson, 1977] R. C. Schank and R. Abelson. *Scripts, Plans, Goals and Understanding*. Hillsdale, N.J.: Lawrence Erlbaum, 1977.

[Searle, 1980] J. Searle. Minds, Brains, and Programs, *Behavioral and Brain Sciences*, 3: 417–424, 450–456, 1980.

[Searle, 1989] J. Searle. *Minds, Brains and Science*. London: Penguin, 1989.

[Searle, 1990] J. Searle. Is the Brain's Mind a Computer Program? *Scientific American*, 262(1): 20–25, 1990.

[Searle, 1992] J. Searle. *The Rediscovery of the Mind*. Cambridge, Mass.: MIT Press, 1992.

[Searle, 1997] J. Searle. *The Mystery of Consciousness*. New York: New York Review of Books, 1997.

[Shannon and McCarthy, 1956] C. E. Shannon and J. McCarthy (eds.). *Automata Studies*. Princeton: Princeton University Press, 1956.

[Shapiro, 1987] S. C. Shapiro (ed.). *Encyclopedia of Artificial Intelligence*, Vol. 2. Wiley: New York, 1987.

[Shortliffe, 1976] E. H. Shortliffe. *Computer-Based Medical Consultations: MYCIN*. New York: Elsevier, 1976.

[Siegelmann, 2003] H. T. Siegelmann. Neural and Super-Turing Computing. In [Copeland, 2002-3].

[Siegelmann and Sontag, 1994] H. T. Siegelmann and E. D. Sontag. Analog Computation via Neural Networks. *Theoretical Computer Science*, 131: 331–360, 1994.

[Simpson, 1977] S. G. Simpson. Degrees of Unsolvability: a Survey of Results. In J. Barwise (ed.), *Handbook of Mathematical Logic*. Amsterdam: North-Holland, 1977.

[Stannett, 1990] M. Stannett. X-machines and the Halting Problem: Building a Super-Turing Machine. *Formal Aspects of Computing*, 2: 331–341, 1990.

[Stannett, 2003] M. Stannett. Computation and Hypercomputation. In [Copeland, 2002-3].

[Sterelny, 1990] K. Sterelny. *The Representational Theory of Mind*. Oxford: Blackwell, 1990.

[Sterrett, 2003] S. G. Sterrett. Turing's Two Tests for Intelligence. In [Moor, 2003a].

[Strachey, 1952] C. S. Strachey. Logical or Non-Mathematical Programmes. *Proceedings of the Association for Computing Machinery*. Toronto, 1952.

[Taylor, 1956] W. K. Taylor. Electrical Simulation of Some Nervous System Functional Activities. In C. Cherry (ed.), *Information Theory*. London: Butterworths, 1956.

[Thorndike, 1933] E. L. Thorndike. The Law of Effect. *Psychological Review*, 40: 434–439, 1933.

[Traiger, 2003] S. Traiger. Making the Right Identification in the Turing Test. In [Moor, 2003a].

[Turing, 1936] A. M. Turing. On Computable Numbers, with an Application to the Entscheidungsproblem. In *The Essential Turing* [Copeland, 2004a].

[Turing, 1939] A. M. Turing. Systems of Logic Based on Ordinals. In *The Essential Turing*.

[Turing, 1940] A. M. Turing. The steckered Enigma. Bombe and Spider. In *The Essential Turing*.

[Turing, 1945] A. M. Turing. Proposed Electronic Calculator. In *Alan Turing's Automatic Computing Engine* [Copeland, 2005a].

[Turing, 1947] A. M. Turing. Lecture on the Automatic Computing Engine. In *The Essential Turing*.

[Turing, 1948] A. M. Turing. Intelligent Machinery. In *The Essential Turing*.

[Turing, 1950a] A. M. Turing. Computing Machinery and Intelligence. In *The Essential Turing*.

[Turing, 1950b] A. M. Turing. Programmers' Handbook for Manchester Electronic Computer. University of Manchester Computing Machine Laboratory. A digital facsimile is in The Turing Archive for the History of Computing. `http://www.AlanTuring.net/programmers_handbook`.

[Turing, 1951] A. M. Turing. Can Digital Computers Think? In *The Essential Turing*.

[Turing, c. 1951] A. M. Turing. Intelligent Machinery, A Heretical Theory. In *The Essential Turing*.

[Turing, 1952] A. M. Turing. The Chemical Basis of Morphogenesis. In *The Essential Turing*.

[Turing, 1953] A. M. Turing. Chess. In *The Essential Turing*.

[Turing *et al.*, 1952] A. M. Turing, R. Braithwaite, G. Jefferson and M. Newman. Can Automatic Calculating Machines Be Said To Think? In *The Essential Turing*.

[Turing, 1959] S. Turing. *Alan M. Turing*. Cambridge: W. Heffer, 1959.

[Tversky and Kahneman, 1983] A. Tversky and D. Kahneman. Extensional Versus Intuitive Reasoning: The Conjunction Fallacy in Probability Judgement. *Psychological Review*, 90: 293–315, 1983.

[Uttley, 1954a] A. M. Uttley. Conditional Probability Machines and Conditioned Reflexes. Memorandum No. 1045, Radar Research Establishment, 1954.

[Uttley, 1954b] A. M. Uttley. The Classification of Signals in the Nervous System. Memorandum No. 1047, Radar Research Establishment, 1954.

[Uttley, 1954c] A. M. Uttley. The Probability of Neural Connections. Memorandum No. 1048, Radar Research Establishment, 1954.

[Uttley, 1954d] A. M. Uttley. The Stability of a Uniform Population of Neurons. Memorandum No. 1049, Radar Research Establishment, 1954.

[Vergis *et al.*, 1986] A. Vergis, K. Steiglitz and B. Dickinson. The Complexity of Analog Computation. *Mathematics and Computers in Simulation*, 28: 91–113, 1986.

[von Neumann, 1963] J. von Neumann. *Collected Works*, Vol. 5. A. H. Taub (ed.). Oxford: Pergamon Press, 1963.

[Weizenbaum, 1966] J. Weizenbaum. ELIZA — a Computer Program for the Study of Natural Language Communication Between Man and Machine. *Communications of the Association for Computing Machinery*, 9: 36–45, 1966.

[Werbos, 1974] P. J. Werbos. Beyond Regression: New Tools for Prediction and Analysis in the Behavioral Sciences. PhD dissertation, Harvard University, 1974. Reprinted in P. J. Werbos, *The Roots of Backpropagation: from Ordered Derivatives to Neural Networks and Political Forecasting*. New York: John Wiley, 1994.

[Whitehead and Russell, 1910-1913] A. N. Whitehead and B. Russell. *Principia Mathematica*, Vols I-III. Cambridge: Cambridge University Press, 1910-1913.

[Widrow and Hoff, 1960] B. Widrow and M. E. Hoff. Adaptive Switching Circuits. *IRE WESCON Convention Record*, Vol. 4. New York: Institute of Radio Engineers, 1960.

[Winograd, 1972] T. A. Winograd. *Understanding Natural Language*. New York: Academic Press, 1972.

[Yamamoto, 1993] M. Yamamoto. "SOZZY": A Hormone-Driven Autonomous Vacuum Cleaner. *Proceedings of the International Society for Optical Engineering*, 2058: 211–222, 1993.

INDEX

abducens, 300
abstract logical rules, 21
abstract object, 220–222, 221, 224, 235
abstract thought, 20
abstraction, 22, 302, 303
    mechanism, 8
achromatopsia, 161
action, 295, 297, 333
    potential, 87, 307, 308, 313
    ageneration, 318
    action-oriented, 334
Adams, F., 351
adaptation, 318, 370, 371, 373, 382–384, 387–391, 394–398, 402, 405
adaptationism, 371
adaptive
    explanation, 386
    thinking, 383, 407
Adolphs, R., 251, 280
affect, 295, 298, 310
affect programs, 262
aggregativity, 58, 61
aggression, 370
agnosia, 166, 171
Agre, P., 413
Ahlquist, J., 399
akinesia, 161
Albright, T., 99
Alcoff, L., 415
alcoholism, 348
algorithm, 86
allocentric coordinate system, 164, 165
alometric scaling, 402
Alzheimer's
    disease, 342
    type dementia, 344
American Pragmatism, 414
amygdala, 280, 298, 299
analog computation, 326
analogical
    matches, 17
    reasoning, 422
analyticity, 223, 224, 235
anatomy, 302, 305
anchoring-and-adjustment, 287
Andreasen, N., 341, 347, 352, 358
anger, 281
animal models, 303
anomia, 238
anorexia nervosa, 357
anterior insula, 281
anthropocentrism, 452
anti-dualist, 340
anti-nativism, 224, 225, 232–234
Antonovics, J., 389
Antony, L., 237, 239, 244
antonyms, 15
apes, 398
Aplysia (sea slug), 12
apparent depths, 178
apparent motion, 175, 176, 180, 181, 191
appraisal, 255, 257
archaeopteryx, 383, 384
architectural question, 268
Ardrey, R., 378
Aristotle, 255, 263
arithmetic word problems, 22
Armon-Jones, C., 255
Arnell, K. M., 170
Arnold, , 255
Artificial Intelligence, 40, 42, 85, 314
    classical, 23
    Good Old-Fashioned AI, 443

Gödel Objection to AI, 429, 471
history of AI, 429
history of AI, 429
situated AI, 429, 444
strong, 85
symbolic AI, 438
Artificial Life (A-life), 433
artificial systems, 5
associationism, 37
associationist, 36
associative visual object agnosia, 166
attention, 260, 295, 297, 298
attentional blink, 170, 171
attractor states, 12, 13
auditory, 300
cortex, 298
Augustine, 347
australopithecuss, 392, 399, 403
Australopithicene, 393
autonomy, 58, 98
availability, 422
axonal bundles, 307

Babbage, C., 38
back-propogation, 441
backward projections, 327
Bain, A., 36
Ballenger, J., 354
Barber, A., 221
Barkow, J., 422
Baron-Cohen, S., 270, 374
Bartsch, K., 275
Barwise, J., 413–415
basic features, 162, 163, 165, 167, 169, 180–182
bats, 301
Baum, D. A., 395
Beatty, J., 383, 386
Bechtel, W., xiii, 81, 99, 375, 416, 467
behaviorism, 225, 226
behaviourist, 31, 38, 40, 42, 66, 250
being appeared-to, 159, 181, 182, 185, 187, 190
Bell, C., 36
Bell, Q, 349
Bennett, S. C., 163, 169
Bentham, J., 66
Benton, T., 380
Bernoulli, J., 385
best first, 434
Beurle, R. L., 441
binding, 213, 214
binocular disparity, 306
bioinforamtics, 313
biologically plausible, 321
Birch, S. A. J., 285
Bisiach, E., 189
Bletchley Park, 430
blindsight, 185–187, 190, 191, 204, 211–213
Block, N. J., 100, 172, 174, 183, 184, 195, 197, 199, 200, 202–205, 207, 214, 215, 269, 450
blockhead, 454
Bloom, P., 285, 372, 390, 391, 393, 406
Bloomfield, P., 222, 223
board games, 429
body-centered coordinate scheme, 189
Boghossian, P., 235
Bombe, 430
borderline personality disorder, 357
Boven, L. van, 284, 286
Boyle, R., 31, 72
Bradley, B. P., 252
Bradshaw, A. D., 389
brain, 295, 296, 298–306, 308–310
development, 296
disease, 353, 363
diseases, 342
disorder, 352
malfunction, 342
plasticity, 301
platicity, 299
size, 399
stem, 300, 304
Braitenberg, V., 446

Brandon, R., 388, 389
breaking the Turing barrier, 464
Breazeal, C., 445
Broca, P., 37
Brodman area 6, 306
Brodmann, K., 99
Brooks, D. R., 395, 397
Brooks, R., 413, 415, 444, 446
bug detector, 317
Buller, D. J., 375
Bunge, M., 465
Burian, R. M., 388
Buss, D. M., 370, 372, 385, 387
Byron, A. A., 39

cabel equations, 318
Cacioppo, J., 254
calcium channel
    gill withdrawal reflex, 12
Calder, A. J., 281, 283
call systems, 298
canonical models, 318
capacity-limited, 168
Carlson, S. M., 276
Cartwright, N., 426
case-based reasoning, 22
categorization, 22, 298
category structure, 22
cats, 303
caudal midbrain, 298
causal, 47
    theories, 325
causation
    principle of causal inheritance, 98
    mental, 83, 99
causes, 343
cell death, 296
central nervous system, 300
cerebellum, 304
cerebral cortex, 304
cerebrum, 296
Chagnon, N. A., 370, 374
Chalmers, D. J., 183, 184, 199, 203
Chamberlain, A. T., 399
Champernowne, D., 434
change blindness, 170, 171
channels, 160, 167, 168
checkers, 436
Cheesman, J., 185
Chess, 430
child-machine, 433
child-scientist theory, 269
chimpanzees, 399
Chinese gym, 462
Chinese room argument, 429
Chisholm, R., 173
Chomsky, N., 42, 222, 230, 234, 238, 243, 377
Chris Eliasmith, xv
Church, A., 429
Church's thesis, 468
Church–Turing
    fallacy, 468
    thesis, 464
Churchland, P. M., 459, 468
Churchland, P. S., 459, 468
Chwalisz, K., 249, 250
clade, 395, 399
cladistics, 395
Clancey, W., 413, 414
Clark, A., xiv, 413
Clark, W. A., 441
classical definition, 343
classical model, 4
    five extra properties, 11
clinical, 341
Clore, G., 261
Coddington, J. A., 395
Cog, 445
cognition, 295, 296, 299, 316, 331, 333, 404, 405
cognitive, 298, 305
    architecture, 331
    behaviour, 332, 333
    competence, 360
    development, 18
    dynamics, 20
    function, 316, 331, 332, 334

neuropsychologists, 329
neuroscience, 43, 295, 297, 298, 310
methodology, 299
phenomenon, 332
cognitive processing, 14, 305, 306
cognitive psychology, 31, 40, 305, 314
science, vii, 3, 22, 229, 239, 314–316, 328, 329
coarse-grained, 25
evidence for Linguistics, 219
scientists, 331
system, 4, 5, 12–13, 24, 329, 333
and discrete representations, 13
complex, 15
defining, 5
example of (thermostat), 5
Cohen, A., 176, 185
Colby, K. M., 437
color, 298
Colossus, 431
common descent, 395
common sense
psychological theory, 269
reasoning, 429
communication, 376, 391, 393, 405, 406
comparative method, 394, 395–398, 407
competition for attention, 160, 168, 169, 189, 190
competitors, 168
complex representations, 333
complex systems, 315
compositional
semantics, 220
structure, 19
complex, 19
symbol system, 18
computability, 464
computation, 86, 296, 302, 306, 316, 326
computational
biology, 313
hypothesis, 25
modeling, 296
computationalism, 3, 10
computer
graphics, 435
simulation, 53, 313
computers, 309
computing machine, 464
comuptationalism, 82
concepts, 15
conceptual
change, 275
role, 4, 325
theories, 325
question, 268
concrete
objects, 220
physical objects, 220
conditional, 332
reasoning, 415
probability
threshhold, 7
conditioning, 31
conjunction of features, 162
connection weights, 328, 329
connectionism, 314, 324, 440, 441, 456
connectionist model, 12, 13, 20, 318
weights, 12
conscious, 259
representational content, 344, 345, 348, 349
consciously, 260
consequences, 343
content, 205, 325
focused treatment, 360–363
information, 8
of representation, 316
representational, 4
semantic, 4, 8
context
of discovery, 61
of justification, 61
context-sensitive processing, 333

control
system, 328
theoretic, 327
theory, 327, 328
convention, 1, 230
of language, 228
convergent evolution, 395, 396, 404
conversion disorders, 344
coordinate system, see also reference frame, 164, 166, 188, 189
Copeland, B. J., xv
cortex, 303, 304, 306, 309
cosine tuning, 322
Cosmides, L., 369, 371–378, 380, 381, 384, 385, 387, 388, 405, 406, 422
coupled oscillatyors, 313
Cowey, A., 186, 187
cranial capacity, 402, 403
Craver, C., xiv, 101
crib, 430, 431
Crick, F., 53
Critchley, H. D., 251
Csibra, G., 271
cultural scaffolding, 357
culturally scaffolded, 354
Cummins, D., 416, 417
Currie, G., 272, 279, 280
curse of knowledge, 285
cybernetic movement, 38
CYC, 439

D-N model, 48
Dalmiya, V., 415
Damasio, A. R., 248, 250, 258, 260, 263
Dartmouth College, 437
Dartmouth Summer Research Project on Artificial Intelligence, 437
Darwin's gift, 370
Darwin, C., 371, 372, 378, 394, 395
Darwinian algorithms, 375
data, 302, 303, 307, 308
Davidson, D., 226, 271
Davies, P. S., 374
Davis, G., 177, 178, 189
Dawkins, R., 372, 377
De Sousa, R., 264
Deacon, T. W., 393, 402, 406
decision, 263
decoding, 317, 319, 321, 324, 330
decomposition, 62, 68, 70, 71, 73, 82, 86
deductive-nomological (D-N) model, 46
degree of problems, 473
delayed response alternation, 306
DENDRAL, 439
dendrites, 307, 321
dendritic currents, 329
Dennett, D. C., 267, 271, 383–387, 406, 443
depersonalization, 349
depression, 339, 346
Descartes, R., 31–35, 38, 43, 72, 163, 377, 466
design, 386, 387
Desimone, R., 168
determination, 85
development, 295
Developmental
neurobiology, 296
psychology, 23
question, 268
Devitt, M., 221, 224
Dewey, J., 414
diagonal argument, 466
Dietrich, E., xiii
differential equations, 313, 327, 329
digital computer, 38
discontinuities of intensity, 164, 165
discrete symbol, 14
disgust, 281
distributed
cognition, 414, 424
representation, 318, 324
Doi, T., 262
domain-general meechanism, 270

Donaghue, M. J., 396
Donders, F. C., 42, 64
dopamine, 281
dorsal horn, 304, 305
dot product, 15
double dissociation, 63, 161
Downes, S. M., 378
Drefus, H. L., 439
Dretske, R., 257
Driver, J., 177, 178, 186–190
dualist, 340
Dummett, M., 226–228, 231
Duncan, J., 168
Dunning, D., 286
Dupré, J., 426
dynamic system, 12, 14, 18
dynamicism, 318
dynamicist, 328, 329
dynamics, 316, 327, 329, 330
    of neural systems, 328, 329
dystonia, 306

early vision, 159, 160, 171
edge detection, 163–166
edge-detection, 177
EDSAC, 434
EDVAC, 39, 432
effective method, 468
egocentric attribution, 272, 283
Ekman, P., 262, 282
electrode, 307, 308
electrodes, 307, 308
eliminativist, 251
Eliza, 437
embedded, 333
embodied, 257
    cognition, 333
emergence, 60, 61, 94
emergentism, 232
emergentist, 235
emotion, 247, 298, 306, 334
    recognition, 280
empathy, 272
empiricism, 224, 225, 232, 233
encephalisation, 402
encoded information, 327
encoders, 329
encoding, 317, 318, 322, 324, 330
    process, 325
endogenous selection, 167
endowment effect, 286
engineering, 388
    design, 382, 384, 385
ENIAC, 39, 431
Enigma, 430
Enns, J. T., 179
environment, 296, 297, 413, 415, 417, 419
epistemology, viii, ix
equivalence fallacy, 469
ERPs, 315
ethics, viii, ix
Evarts, E. V., 307
event related potentials (ERPs), 241, 331
evolution, 369, 370, 372, 374, 391, 404
    by natural selection, 388
evolutionlary computing, 429
evolved, 262
exchange, 376
exogenous selection, 167, 180, 181
experience, 309
experimental neuroscience, 313
experiments
    bottom up, 91
    top down, 91
expert system, 429, 439
explanation, xi, 299
explanatory realisation, 87
explicit, 298
externalist, 451
extinction, 188, 190

face
    -based emotion recognition (FaBER), 280
    -recognition, 442

robot, 445
-selective cells, 317
facial expression, 445
feedback, 252
faculties, 345
Fadiga, L., 277
false-belief task, 267, 274
Farah, M., 166
Farley, B. G., 441
fear, 280, 299, 310
feature, 172, 173, 181
hierarchy, 165–167
maps, 160, 171
feedback, 305
feedforward computation, 327
feeling, 247, 254, 260
Feigenbaum, E. A., 439
Feinberg, J., 346
Feldman, J., 439
Felsenstein, J., 396
feminist epistemology, 414
Ferranti, M. I., 433
Ferrier, D., 36
field potentials, 307, 308
finite state device, 39, 42
firing rates, 318, 319
Fischoff, B., 285
fitness, 374, 382, 388
Fleagle, J. F., 399
Flourens, J.-P.-M., 35, 37
Flowers, T., 431
fMRI, 43, 64, 298, 305, 315
Fodor, J., 100, 195, 230, 236, 258, 270, 275, 374, 469
folk psychology, 267
form, 387
primitives, 163, 166, 169
Forster, K. I., 241
foveae, 300
foveated, 301
Fowler, C., 241
frame problem, 444
FREDDY, 438
Frege, G., 220, 221, 223, 238
French, R., 450
Freud, S., 340, 341
frogs, 301, 303, 307
FRUMP, 438
functioanlism, 86
function, 92, 387
functional
impairment, 345
magnetic resonance imaging (fMRI), 331
neurimaging, 43
role, 4, 8
functionalism, 83, 269, 469
Turing machine, 86
fusiform face area, 167

Galanter, E., 41
Galileo, 31, 72
Galileo, G., 417
Galison, P., 413, 421
Gall, F. J., 35–37
Gallese, V., 272, 277, 278
Galvani, L., 307
Garrett, R., 346
Gaussian, 322
General Problem Solver (GPS), 437
generalisations, 268, 302
generate-and-test, 430
generic neural subsystem, 329, 330
genes, 296
genetic
algorithm, 433
variance, 392
George, A., 221, 224
Georgopoulos, A., 321
Graham, G., xv
Ghiselin, M., 385
Gibbons, 398
Gibson, J., 414
Gigerenzer, G., 373, 378
gill withdrawal reflex, 12
Gillett, C., 81
Glossip, 397
goal state, 5

goals, 415–417, 420, 423
Gödel objection to AI, 427, 469
GOFAI, 443
Goldman, A., xiv, 272, 277, 280, 282, 283, 287, 414, 425
Good Old-Fashioned AI, 441
Goodman, N., 166
Gopnik, A., 269, 275
Gordon, R., 272
gorilla, 399
Gould, S. J., 378, 385, 402
Government Code and Cypher School (GC & CS), 430
Graham, G., xv
grammar, 223, 224, 236–238, 243
Grantham, T., 372
Greenspan, P., 255
Griffiths, P. E., 261, 262, 378, 383
Grine, 399
Guttenplan, S., 469

Hacking, I., 413, 421
halting
  function, 465
  problem, 472
Hara, F., 445
Haraway, D., 413, 414, 417–419, 421
Hardcstle, V. G., xv
Harding, S., 413, 418, 419
harm, 345
Harris, P., 272
Harrison, B. G., 354
Hartley, D., 36
Hartsock, N., 418, 419
Harvey, P. H., 395
Harvey, W., 33
Haugeland, J., 443
Heal, J., 272
heavy metal tolerance, 389, 390
Hebb, D., 441
Hebbian learning, 309
Heidegger, M., 414
hemi-neglect, 185–191
hemianopia, 188
Hennig, W., 395
Herbert, 444
heritability, 390, 392, 393
heritable, 389
heuristic search, 431
Higginbotham, J., 221, 235, 243
higher-order thought theory, 260
Hilbert program, 466
hippocampus, 295
historical
  context, 413, 415, 421, 422
  function, 388
Hobbes, T., 66, 466
Hobson J. A., 341
Hodges, A., 447, 450
Hodgkin, A. L., 314
Hohmann, G. W., 249, 250
Holling, P. S., 383
Holyoak, K., 422
hominid, 392
  evolution, 398
  phylogeny, 401
homo, 392, 393, 399
  erectus, 391, 403
  habilis, 391, 403
  sapiens, 391, 399, 402–404
homologous, 300
homunculus, 170, 180
Horwitz, A., 340, 357
Hubel, D., 317
Hull, C., 66
Hull, D., 418
human
  computer, 464
  Evolution, 398
  health, 341
  language, 377, 390–393
  reasoning, 373, 392
  violence, 370
humanoid, 445
humans, 399
Hume, D., 225, 263
Hurtshouse, R., 264

Hutchins, E., 413–415, 417, 418, 421, 422
Huxley, A., 314
hypercomputation, 429, 475
hypercomputer, 464
hypotheses, 307

IBM 701, 436
identity, 85
  type, 85
ideomotor apraxia, 351
illness, 343, 345
imagery, 279, 323
images, 323, 324
imaging, 305
  studies, 298
imitation game, 447
impairment, 345, 347
implementation, 82, 324
  and reduction, 10–11
  constraints, 317
implicit, 298
impulse control disorder, 344, 357
incest, 378, 379
  taboos, 378–380
incongruous colors, 185
indeterminacy, 243
  of grammars, 243
influence stream, 358
information, 319, 320, 322, 323, 326
  causal link, 6
  in the brain, 3
  low-level, 9
  mathematical theory of, 6
  processing, 41, 43, 63, 64, 72, 313
  state, 5
  theory, 40, 316
inhibitory control, 275
innate
  information, 225
  mechanisms, 225
inputs
  categorizing and descriminating, 16
intelligence, 399, 402, 403
intentional, 255
inter-spike intervals, 319
interfield theory, 59
internal dynamics, 327
internalist, 451
interpretation, 2
  infinite regress of explanation, 2
intertheoretic reduction, 59
intuition, 466
ion
  channels, 318
  conductance, 314
irregular verbs, 239
Irschick, 397
isomorphic copy, 463
itch, 306
Iten, C., xiv
Ivry, R. B., 176, 185

Jackman, 397
Jackson, J. H., 36
Jacquard, J.-M., 38
Jagger, A., 419
Jain, S. K., 389
James, W., 248–252
James–Lang theory, 248, 249, 250–253, 259
Jaspers, K., 363
Johnson, M., 413
Johnson-Laird, P., 252, 469
judgements, 255, 257
Julesz, B., 162
justification, 299

Kahneman, D., 422, 423
Kamin, L., 379
Kanizsa, 189
  figures, 177
  square, 176–178, 182
  triangles, 175
Kant, I., 263, 377

Katz, J., 219–221, 223, 224, 229, 235–238, 243, 244
Keller, E. F., 413, 415, 418, 419
Kempley, S., 241
Keysar, B., 284
Keysers, C., 278
Kim, J., 82, 94, 98
Kinsbourne, M., 422
kinship, 374, 379
Kirsh, D., 413, 415, 420–422, 426
Kismet, 445
Kitcher, P., 378–380, 414, 416, 418
Kluge, 399
Knorr-Cetina, K., 413
knowledge, 6
  how, 425
  of mathematics, 221
  that, 425
Kobayashi, H., 445
Kohn, C., 185
Kosslyn, S., 279
Kosslyn, S. M., 280
Kriegel, U., xiv
Kuhn, T., 417, 422
Kunda, Z., 416
Kuniecki, M., 251

La Mettrie, J. O. de, 35
labyrinthectomy, 301, 302
Lakoff, G., 413
lamprey, 313
Lang, C., 248, 250–252
Langton, C., 433
language, 295, 298, 323, 331, 388, 390–392, 402–405
  acquisition, 232, 234, 405
  as a social practice, 231
  language-game, 228–230
  as abstract objects, 221, 235
  as mental objects, 222
  as social practice, 226
  learnability, 238
  learning, 225
  normative dimension, 226
Larson, A., 395, 397, 398
Lashley, K., 37
latent Semantic Indexing, 14
lateral intraparietal cortex (LIP), 322
Latour, B., 413, 421
Lauder, G. V., 386, 387, 390, 404
Laurence, S., 234, 238, 239
Lave, J., 413–415
Lawrence, A. D., 281
laws, 31, 44, 46, 48, 51, 53, 66, 72
Lazarus, R. S., 255
learn represntations, 324
learnability, 238
learning, 296, 309, 375, 376, 429
  theory, 224
Lederberg, J., 439
LeDoux, J. E., 256
leech, 313
left hemisphere, 19
Lenat, D. B., 439, 453
Leonardo, 446
Leroi, A. M., 404
lesion, 305
Leslie, A., 270, 275, 276
levels, 56, 92, 101
  of analysis, 57
  of description, 329
  of organisation, 55, 57, 61, 299
  orders, 92
  vs. orders, 94
Levenson, R. W., 252, 254
  item Lewis, C. S., 81
Lewis, D., 221, 169
Lewontin, R. C., 379, 385
lexical processing, 15
Lieberman, P., 391, 399
linear
  decoding, 320, 324, 326
  transformation, 327
  , time-invariant (LTI), 327
lingistics, 42, 223, 230
  abilities, 393
  a priori restriction of the evidence base, 219, 224

as a science, 230, 231
Chomskian, 23
evidence from the cognitive sciences, 219, 232
facts, 237
knowledge, 226
meaning as use, 227
modern, 23
neuropsychological evidence, 219
norms, 228
phenomena, 230
psychological evidence, 227
restrictions on the evidence-base for linguistics, 229
structuralist, 222
the discipline of, 244
LIP, 327
local
edges, 164
features, 163
field potential, 304
localisation, 62, 69–71, 73, 82
Locke, J., 36
Loebner Turing Test Contest, 449
Loebner, H., 453
Loewenstein, G., 286
Logic Theorist, 437
logical reply, 457
long-term potentiation, 298, 299
Longino, H., 413, 417–419
look-up table, 454
Lorenz machine, 431
Lorenz, K., 378
Losos, J. G., 396–398
Lowenstein, G., 285
Lucas, J., 471
Luck, S. J., 170
Lumsden, C., 378, 380
Lycan, B., 101

Macaulay, C. P. R., 39
Machester Baby, 434
MacNair, 390
macroevolution, 394
Maddison, W. P., 396
Magendie, F., 36
major depression, 357, 359
major depressive episode, 349
make believe, 446
mammals, 300, 301, 304
Margolis, E., 234
Marquand, H. G., 39
Marr, D., 94, 163, 179, 425
Mason, K., xiv
mate preference, 370
material realisation, 87
mathematical formalism, 313
Mathematical Objection, 432
mathematical reasoning, 298
mathematics, 220, 221, 235
Mattingley, J. B., 189
McCarthy, J., 437, 454
McCauley, R., xiv
McClelland, J., 442
McClintock, B., 415
McCollum, 399
McCulloch, W. S., 440
McHenry, 399
McLennan, D. A., 395, 397
Mead, G., 414
meaning, 316, 325
mechanism, 31, 33, 34, 38, 40, 43–45, 49, 51–53, 61, 63, 65, 67, 71–73, 82, 85, 91, 471 73
organisatin of, 91, 94
mechanistic, 319
Meltzoff, A., 269, 272
Melville, H., 235
memories, 297
memory, 295, 297, 298
mental
categories, 314
distress, 346
grammar, 238
illness, 339, 340
prototype, 343
imagery, 298
object, 220, 222

promiscuity, 351
representation, 2, 314, 316
mediators, 6
simulation, 53, 267
mentalising, 267
mentalistic approach, 223
Merikle, P. M., 185
Merleau-Ponty, M., 414
metaphysics, viii, ix, xii
Mettrie, J. O. de la, 466
Michie, D., 438
micro-world, 429, 437
microelectrode, 315
microevolution, 394
Micronesian navigation, 416, 424
Miles, 396
Mill, J. S., 36, 66
Miller, G., 40
Millikan, R., 97
Mills, C. W., 414
Milner, P., 68
mind-brain
implementation, 9
supervenience, 9
mind-mind relation, 4
mind-reading 267
in infants, 274
minimax, 434
Minsky, M., 437, 442
mirror neurons, 251, 277, 297
modelling, 313
models, 53, 60, 65, 73
modular, 262
modularity, 270, 375, 376
of mind, 374
modus ponens, 21–22
Mogg, K., 252
Moniz, 399
morphology, 220, 239, 241
Morris, D., 378
Morton, J., 241
Moses, L. J., 276
motivation, 65–67, 71, 72
motivational
focus, 360
oscillatory theory (MOT), 329
motor
areas, 298
control, 315, 331
cortex, 321
deficits, 351
imagery, 280
planning, 297
system, 304, 321
systems, 297
movement disorders, 351
movements, 306
Mueller-Lyer illusion, 173
multi-layer, 441
multielectrodes, 315
multiple realisability, 83,85
Mundale, J., 99
Münte, T. F., 241, 242
muscles, 297
MYCIN, 439

Nagel, T., 174, 183
naive causal theory, 317
natural, 394, 404
kinds, 261
sciences, 224
selection, 369–373, 388, 390, 391, 393, 404
naturalism, viii
naturalistic, 316
nature of language, 221
naval navigation, 424
Neander, K., 97
necessary truths, 223, 235
necessity
nomic, 82
negative feedback, 58
Neisser, U., 160
neocortex, 298
nervous system, 303–305
Neumann, J. von, 39, 431
neural
action potential, 317

architecture, 2
architectures, 328
computation, 326
computing, 314
control theoretic, 328
correlates, 82
mechanisms, 313
nets, 309
nonlinearity, 329
plausibility, 333
population, 323
representation, 328, 329
spike, 320
spikes, 329
system, 313, 316, 324
neurally
plausible, 333
realistic, 333
neuroanatomy, 315
neurobiological systems, 333
neuroimaging, 63, 71, 250, 262
neuron, 295, 298, 301, 303–305, 307–309
migration, 296
tuning curve, 319
-like computing, 429
neurophysiology, 309
neuropsychological evidence, 220, 224, 229
neuropsychologists, 43, 63
neuroscience, 219, 229, 295, 296, 299, 301–303, 306, 307, 309, 310, 317
neuroscientific theories, 299, 302
neuroscientists, 296, 305, 308, 309
neurotransmitters, 295, 300
neurotrophins, 295
Nevell, H., 99
Newell, A., 40, 437
Newman, M., 449, 455
Newmanry, 431
Nichols S., 284, 372
Nietsche, F., 66
Nilsson, 438
Nisbett, R., 422
nociceptive information, 304
noise, 320
noisy, 323
nonlinear c
omputation, 327
encoding, 320, 324, 326
systems, 433
theory, 313
Norman, D., 413, 420
number-theoretic, 466
numbers, 223
Nussbaum, M. C., 255
nystagmus, 301

o-machine, 465
o-network, 461
Oatley, K., 252
object
-centered coordinate system, 189
permanence, 18, 20
Piagetian task, 18
processing, 303
recognition, 323
observation, 307, 308
obsessive-compulsive disorder, 357
ocular motor, 301
Oettinger, A., 436
Olds J., 68
ontological considerations, 221
ontology of languages, 238
Oppenheim, P., 56
optimal linear decoder, 322
optimality, 382, 383
oracle, 466
machine, 465
oracular computation, 467
orangutan, 398
orders
vs. levels, 94
osteoarthritis, 352

Pagel, 395
pain, 260, 306
pains, 260

palaeoecology, 392
panglossianism, 385
panic
  attack, 354
  disorder, 339, 344, 349
Papert, S., 437, 442
parallel
  distributed processing, 440
  evolution, 395
  processing, 160, 162, 167, 168, 171, 176–179, 182
  system, 460
Parkinson's disease, 342
Parry, 437
Parsons, L., 280
Part-Of principle, 458
parts/wholes, 95
past tense, 442
Pastore, A., 39
patch clamps, 315
Peirce, C., 39
Penrose, R., 429, 471
peptides, 300
perception, 211–213, 295–297, 331, 333, 441
perceptual
  /motor loop, 334
  process, 305
  systems, 296, 334
perisylvian areas, 298
Perner, J., 267, 274, 275
Perry, J., 413–415
personality
  disorders, 344, 349
  disturbance, 342
  transformation, 362
PET, 43, 64, 315
Petersen, E. E., 170
phase spaces, 309
phenomenal
  character, 182, 196, 197, 202, 205, 206
  consciousness, 159, 172, 175, 181–183
  content, 184
  porperty, 183
  property, 159, 172–177, 179, 181–185, 187, 190, 191
phenylketonuria, 353
Philippot, P., 254
philosophical despair, 346, 347
philosophy, viii
phonation, 298
phonetic codes, 298
Phonology, 222
photopic vision, 186, 187
phylogenetic
  systematics, 395
  tree, 395
phylogeny, 394–396, 398
physical grounding hypothesis, 444
physical objects, 220, 221
physical symbol system, 438, 470
physicalism, 82, 83, 85
physiologies, 302
physiology, 303
Piaget, J.
  object permanence tasks, 18
  stage theory, 23
Pickering, A., 413, 421
pictorial depth cues, 163, 166, 175, 178, 179, 184, 191
pigmy chimpanzees, 399
Pilbeam, D., 399
Pinker, S., 371, 372, 390, 391, 393, 406
Pitcher, G., 255
Pitt, Brad, 97
Pitts, W., 440
Place, U. T., 81
planetary rover, 445
plasticity, 295, 296, 305
  neural, 100
Platonism, 223, 235–237
pleistocene, 373, 388, 392, 403, 404
pop out, 161–163, 168, 176, 177, 179–182, 189
population

decoders, 329
of neurons, 322
representation, 318, 324
-temporal decoder, 325, 326
positive feedback, 58
positron emission tomography (PET), 331
Posner, M. I., 168, 170
Post Office Research Station, 431
Post, E., 471
post-attentive
attention, 190
processing, 167–169, 171, 172, 182
post-traumatic stress disorder, 357
postsynaptic currents (PSCs), 321, 329
Poverty of Stimulus, 234
poverty of stimulus, 234
practice, 413, 415–417, 419, 421–423, 426
preattentive
features, 163, 164
processing, 159, 162, 163, 167–169, 171, 175, 177–180, 182, 184, 188, 189, 191
preferred
direction vector, 323
functions, 323
stimulus, 322
pregnancy sickness, 384, 385
Premack, D., 267
Pribam, K., 41
primary sensory, 298
primates, 300, 398
priming
experiment, 241
experiments, 240
reaction times, 241
Prinz, D., 435
Prinz, J., xiv
Prinz, J. J., 248, 251, 252, 257, 258, 260
problem
in living, 346, 363
of living, 341, 342, 347
solving, 22
procedural, 298
processing, 303
Profet, M., 384
properties
relational, 96
property
aggregative, 89
structural, 89
propositional attitude, 20
prosopagnosia, 166
prototype, 350
semantic structure, 343
Proudfoot, D., xv
psycholinguistics, 236
psychological
and neurological priming effects, 242
development, 233
evidence, 230, 235, 237
explanation, 1, 11, 333
faculties, 346
psychologism, 236
psychology, vii, 221, 223, 224, 229, 233, 235, 240
psychopathology, 339
Purkinje shift, 187
Putnam, H., 56, 84, 465

qualia, 173, 183, 187
quantum mechanical computer, 467
Quine, W., 424
Quine, W. V. O., 219, 224–227, 229, 232–235, 243, 244, 271, 272

radical empiricism, 224, 225, 233, 234
Rafal, R. D., 185
raining, 440
Ramachandran, V. S., 177
random element, 475
Raphael, B., 438
rate code, 319, 320,321
rational, 263
decision-making, 298
rationality theory, 271

rats, 299
Ravenscroft, I., 279, 280
Raymond, J. E., 170
reaction time, 42, 64
realisation
  aggregate, 89
  aggregative, 92
  assymetry, 93
  canonical statement, 93
  constitutivity thesis, 96
  contextual, 96
  emergent, 102
  explanatory, 89
  historical, 97
  material, 88, 92
  mechanistic, 90–92
  multiple, 98
  structural, 89, 92
  sufficiency, 93
  synchronic, 96
reasoning, 263, 331, 388
recall, 298
recollection, 298
recurrence, 327
recursiveness, 469
reduction, 9, 54, 60, 61
  and implementation, 10–11
  example (the word-processor), 10
reductionism, 73, 82, 83
reductionist, 251
regular verbs, 239
Reisenzein, R., 254
relationship between ontology and methodology, 239
reliance, 403
Rensink, R. A., 179
representation, 50, 52, 73, 195, 257, 297, 298, 309, 316, 330
  context free, 22
  continuous, 13–15
  discrete, 13–17
    categorize inputs, 16
    definition, 17
    symbols, 13
  dynamic, 12
  feature-list, 13
  fundamental problem of, 1
  higher-level, 8
  interpretation, 1
  noncompositional, 20
  of a function, 323
  pictures, 1
  representational error, 6
  transient, 12
  view of the mind, 18
representational
  content, 316, 325
  decoder, 326, 329
  vehicles, 325
representativeness, 422
Resnick, M., 415
retina, 300
retinotopic coordinate system, 164, 188
reverse, 388
  engineering, 382–385, 394, 407
reward, 65, 67, 70–72, 309
Richardson, R., xv, 375, 383, 388, 416, 467
right hemisphere, 19
Rizzolatti, G., 277, 278
Robinson, T. E., 248
robot, 437
robotics, 429
rodents, 300
role-argument structure, 17
role/occupant distinction, 85
Rorty, A. O., 261
Rose, H. R., 377, 379, 404
Roseman, , 255
Rosenblatt, F., 441
Rosenthal, D., 260
Ross, L., 422
Rothbart, M. K., 168, 170
Rugg, M. D., 241
rule, 333
  and representations, 390
  -based systems, 24

-use, 24
rules, 333
Rumelhart, D., 442
Ruse, M., 378
Russell, B., 346, 437
Ryle, G., 261

Sahlins, M., 380
salience, 422
Salmon, W., 44
SAM, 438
Samuel, A., 436
Schachter, S., 253, 254
Schank, R. C., 438
schemas, 13
Scherer, K., 255
Scherer, K. R., 257
schizophrenia, 362
Schopenhauer, A., 346
scotoma, 186, 190
scotopic vision, 186, 187
search, 429
Searl, J., 429
segmentation, 160, 162, 178, 184
selection, 386, 389, 391, 394, 404
selective attention, 159, 167–171, 176, 180, 181, 189, 190
self-organisation, 58
semantic, 298
networks, 13
promiscuity, 351
theories, 325
value, 316
views of theory, 53
semantics, 241, 316, 325
semicircular canal, 300
sense-data, 165, 173
sensorimotor
boundary, 11
interaction, 20
sensory
modalities, 297
substitution, 302
systems, 296
serial
processing, 167, 168, 179
search, 162
system, 460
severe delusions, 362
sexual selection, 389
Shakey, 438
Shannon, C., 40, 454
shape from shading, 163, 179, 184
Shapiro, L., 81, 100, 170
Shaw, J. C., 437
Shoemaker, S., 269
Shopper, 436
SHRDLU, 437
Sibley, C. G, 399
Siegelmann,, H. T., 461
similarity, 15
Simon, H., 40, 56, 437
simulated, 471
simulation, 73, 468
fallacy, 461
since-cell studies, 305
Singer, C., 253, 254
single cell recording, 307, 309
single neurons, 317
single unit activity, 308
single-cell learning, 331
single-cell recordings, 304, 305, 307
situated action, 22
situation sensitive, 354
mood disorders, 344
situational
identity, 356
sensitivity, 356
Skelton, 399
Skinner, B. F., 38
sleep, 295
Slovic, P., 422
Smart, J. J. C., 81
Smith, C., 255
Smith, D., 418, 419
Soames, S., 221, 224, 237
Sober, E., 97
social, 376

behavior, 378
contracts, 377
exchange, 373, 375–377
relations, 373
robotics, 445
socially constructed, 262
sociobiology, 377–381
sodium channel, 87
software/hardware distinction, 86
Solomon, M., xv, 414, 418
Solomon, R. C., 255, 256
somatic marker, 258
somatosensory, 302
Sontag, E. D., 461
spandrel, 397
special science
autonomy of, 98
species, 301–303
speech processing, 306
Spencer, 378
spike, 320
histograms, 319
train, 317
spiking neurons, 333
Spinoza, B., 248, 255, 263
spotlight analogy, 169, 170
Sprengelmeyer, R., 281
Sripada, C. S., 280, 282, 283
stabilizing selection, 394, 396
Stainton, R., xiv, 239
standard
model, 4
social science model, 375, 381
standpoint epistemology, 419
Stanhope, C., 38
Stanners, R. J., 241
state variable, 327, 328
state-space, 327
statistial inference, 323
steam engine governor, 24
Stenberg, S., 42
Stepan, N., 417
Sterelny, K., 378
Stevens, G. L., xv
Stich, S., 243, 284
stimuli, 304
inputs, 297
Stoerig, P., 186, 187
stored program concept, 430
Strack, F., 252
Strahan, E. J., 261
Strait, 399
stream of influence, 350
stress-related adjustment disorders, 344
stressors, 344, 355
Stroop interference, 185
structure-mapping theory, 17
structure-sensitive processing, 333
sub-cultural deviance, 346
sub-personal information processing, 351
subjective contour, 177–179, 182, 190, 191
subjective contours, 175–177
subliminal perception, 172, 208, 210, 214
substance abuse, 342
subtraction
method, 306
studies, 306
Suchman, L., 413–415
superarticulacy, 453
supercomputer, 3
supervenience, 9, 82
example (hardware and software), 9
Sur, M., 100
symbol manipulation, 458, 471
symbolic
models, 318
representation, 326
Symons, D., 374
symptoms, 343
synapses, 298
synaptic
circuitry, 296
connectivity, 296
weights, 330, 331

syntactic
structure, 316
generalisation, 333
syntax, 220, 241
systemacity, 333
systematics, 396
systems reply, 457
Szasz, T, 341

Taylor, M. A., 352
Taylor, R., 346
Taylor, W. K., 441
telelogical explanation, 33
temporal
decoders, 321, 329
decoding, 319, 320
lobe epilepsy, 342
lobes, 298
representation, 318, 319, 321, 324
texture segmentation, 162, 163, 176, 180
Thagard, P., 263, 414, 418, 422, 425, 426
thalamus, 304
theoretical, 303
theories, 302
theory, 302, 307
of mind, 267
theory, 268
-laden observations, 307
-ladeness, 308
therapeutic feedback, 361
Thorndike, E. L., 441
three layer network, 324
time-varying signals, 318
timing code, 319
timing codes, 320
tokens, 222
Tooby, J., 369, 371–378, 380, 381, 384, 385, 387, 388, 405, 406, 442
tools, 413, 415, 419–421, 423, 424
TOTE unit, 41
training, 462
trait polarity, 390
transducers, 296
transformational decoder, 326, 327, 329
transformations, 324
treatment, 343, 360
Treisman, A., 161, 162, 165
tuning curve, 323
Turing, A., 39, 86, 429
Archive for the History of Computing, 431
machine, 39, 42, 84, 86
Nets, 440
test, 447
thesis, 468
Turner, R.G., 389
Tversky, A., 287, 422, 423
two-factor theory, 325
two-layer, 441
types, 222

uncertainty, 323
unconscious, 260, 340
memory, 356
UNIVAC, 39, 433
universal Turing machine, 429
unorganised machine, 433
Uttal, W. R., 375
Uttley, A. M., 441

V1, 297
valanced, 251, 262
Vallar, G., 189
van den Berghe, P., 378
variance, 392
variation, 391
vector representation, 322, 323
vehicles, 325
of represntation, 316
ventromedial prefontal cortex, 334
vernier offset, 163
vestibular, 301
nuclei, 301, 304
system, 300
-colic response, 300, 301

-ocular reflex, 301, 302
-ocular response, 300, 301
visual
agnosias, 166
appearance, 159, 172–176, 178–180, 182, 183, 185, 187
awareness, 159, 169, 170, 172, 175, 180–188, 190
experience, 172, 183
features, 160, 166, 171
fixation, 300
illusion, 172, 173, 175–177, 190
imagery, 279
masking, 184, 185
priming, 184, 185
processing, 303, 315
search, 161, 162
system, 19
vitalism, 94
Vogel, E. K., 170
Vogel, S., 385, 386
voluntary, 306
Vonnegut, K., 235
Vuilleumier, P., 186–190
Vygotsky, L., 414

Walter, G., 444
waltz, 227, 228
Ward, 398
Wason card selection task, 332
Wason selection task, 21
Watson, J., 53
waveform, 308
Wearing, C., xiv
Wegener, A., 423
Weiskrantz, L., 186
Weizenbaum, J., 437
well-formedness, 243
Wellman, H., 269, 275
Wernicke, C., 37
what it is like, 172, 174, 183, 184
Whitehead, A. N., 437
Wicker, B., 278
Wiesel, T., 317
Williams, G. C., 386
Wilson, E. O., 377, 378, 380, 381
Wilson, R., xiv, 81
Wimmer, H., 267, 274
Wimsatt, W. C., 56
Winograd, T. A., 437
WISARD, 442
Wittgenstein, L., 219, 226–232, 236, 238, 244, 250
Wolfe, J. M., 161–163, 169, 176, 179
Wood, B. A., 399
Woodger, M., 443
Woodruff, G., 267
Woolfe, L., 349
Woolfe, V., 349, 350, 360
working memory, 306, 314, 332
world-mind relation, 4
Wright, C., xiii

XOR problem, 442

Yamamoto, M., 446
Yanamamö, 370, 374

Zajonc, R. B., 256
Zeno, 237