

Symmetric Geodesic Shape Averaging and Shape Interpolation

Brian Avants and James Gee

University of Pennsylvania
Philadelphia, PA, USA 19104-2644
{avants,gee}@grasp.cis.upenn.edu

Abstract. Structural image registration is often achieved through diffeomorphic transformations. The formalism associated with the diffeomorphic framework allows one to define curved distances which are often more appropriate for morphological comparisons of anatomy. However, the correspondence problem as well as the metric distances across the database depend upon the chosen reference anatomy, requiring average transformations to be estimated. The goal of this paper is to develop an algorithm which, given a database of images, estimates an average shape based on the geodesic distances of curved, time-dependent transformations. Specifically, this paper will develop direct, efficient, symmetric methods for generating average anatomical shapes from diffeomorphic registration algorithms. The need for these types of averages is illustrated with synthetic examples and the novel algorithm is compared to the usual approach of averaging linear transformations. Furthermore, the same algorithm will be used for shape interpolation that is independent of the multi-scale framework used.

1 Introduction

An atlas may be used as an instance of anatomy upon which teaching or surgical planning is based [1], a reference frame for understanding the normal variation of anatomy [2], a coordinate system for functional localization studies [3], and as a probabilistic space into which functional or structural features are mapped [4]. Least biased examples are desirable for teaching, as well as for creating coordinate systems that are near the centroid of a population distribution. Performance of algorithms based on manipulating canonical information, such as active shape, should also improve when using an average model.

Computerized atlases based on MRI images may capture either average shape [4], average intensity or both [5] within a single image. Deviations from the mean shape or intensity may be stored separately by statistical models such as principal components [4]. Average intensities are found by first computing transformations from a given anatomical instance to a population dataset. These transformations give intensity correspondence, allowing subsequent averaging. Average shapes are gained by estimating the average of these transformations,

which take a given member of population to the remainder of the data. This average transformation must then be inverted to gain the average shape [5].

One difficulty with this approach is that the process of averaging transformations may destroy the physical and optimal properties of the individual transformations. For example, the average of large deformation elastic displacement fields, each of which satisfy the minimization of a well-defined variational energy, may no longer be an optimizing elastic displacement field. Another example is given by time-parameterized mappings. The flows defining these transformations at each time satisfy the fluid equations, allowing the maps to be interpreted as members of the diffeomorphism group [6]. This invites group theoretical population studies where one bases structural comparisons on the geodesic distances of the group. Thus, it is important to be able to compute atlases which are least-biased within this theoretical framework, as in the small deformation case. This work provides a general algorithm for allowing shape averaging that enables properties of the physical model used in the registration to persist in the average shape transformation. The distances given by the diffeomorphism group are used to illustrate the techniques.

2 Population Shape Averaging

Consider a set of anatomical images defined on bounded domain Ω , each of which contains identical topology initially at positions $\{\mathbf{x}_i\}$. Shape normalization requires a reparameterization of this population dataset, $P = \{\mathbf{x}_i\}$, into a common coordinate space. Each coordinate $\bar{\mathbf{x}}$ then identifies the same anatomical position in each example. Formally, this requires a mapping set $\{g_i: \bar{\mathbf{x}} \rightarrow \mathbf{x}_i\}$ that gives,

$$\begin{aligned} g_1^{-1} \mathbf{x}_1 &= \bar{\mathbf{x}}, \\ &\dots \\ g_n^{-1} \mathbf{x}_n &= \bar{\mathbf{x}}. \end{aligned} \tag{1}$$

Each mapping gives the coordinate transformation between the canonical configuration $\bar{\mathbf{x}}$ and $\mathbf{x}_i$, such that $g_i = \mathbf{Id} + \mathbf{u}$, where $\mathbf{Id}$ is the identity. This gives $g_i(\bar{\mathbf{x}}) = \bar{\mathbf{x}} + \mathbf{u}(\bar{\mathbf{x}}) = \mathbf{x}_i$. If g_i is time-parameterized in interval $[0, 1]$, its value is taken at $g_i(1)$, the final state.

The individual mappings, g_i , may be found by using non-rigid image registration algorithms [4, 3, 7]. In general, these methods return a displacement field, $\mathbf{u}$, that models the motion of a continuum deforming under external forces. Solutions of this kind minimize a balance of regularization and similarity terms. The displacement field is found in either the Lagrangian reference frame, where the reference configuration $\mathbf{x}$ is fixed, or in the Eulerian frame, where the configuration is a function of time, $\mathbf{x}(t)$. The distinction here is that the Eulerian frame tracks the flow in time, giving the displacement as the time-integration of a velocity field, $\frac{D\mathbf{g}}{dt} = \mathbf{v}(t)$, where the material derivative is used. Given a value

for $\mathbf{u}$, the solution will transform the moving image $J(\mathbf{x}_i)$ to the fixed domain $I(\bar{\mathbf{x}})$ such that $J'(\bar{\mathbf{x}}) = J(\mathbf{Id} + \mathbf{u}_i^{-1})(\mathbf{x}_i) = J \circ \mathbf{g}_i^{-1} \circ \mathbf{x}_i$.

2.1 Shape Average from Mean Deformation: Linear Averaging

Here, we assume a registration algorithm has provided a correspondence field. Given the ability to gain this solution, average shapes are found by choosing an arbitrary instance as the reference configuration, $\mathbf{x}_o = \mathbf{x}_j$, and then computing $\{\mathbf{g}_i\}$ with respect to this configuration. The resulting average displacement field from $\mathbf{x}_o$ to $\{\mathbf{x}_i\}$ is,

$$\bar{\mathbf{u}}(\mathbf{x}_o) = (N)^{-1} \sum_{i=1}^N \mathbf{u}_i. \quad (2)$$

This mean deformation minimizes the energy,

$$\bar{\mathbf{u}}(\mathbf{x}_o) = \underset{\mathbf{u}}{\operatorname{argmin}} \frac{1}{2} \sum_{i=1}^N |\mathbf{u}_i - \mathbf{u}|^2 \quad (3)$$

Note that, as the total displacement fields are used, all scales of information are treated equally. The average configuration is then computed as $\bar{\mathbf{x}} = \bar{\mathbf{g}}^{-1} \mathbf{x}_o$, where an inversion of the mean deformation field is required. Averaging of vector fields does not necessarily preserve the large deformation continuum model, nor does it satisfy the correct optimization model, as illustrated in figure 1.

3 Curved Case: The Diffeomorphism Group and Its Distances

3.1 Definitions and Group Properties

The set of one to one and onto differentiable maps with differentiable inverse gives the diffeomorphism group, $\mathcal{G}$. Elements of the group may be composed with each other and distances between them measured by (in the case of image registration),

$$D(I, J) = \inf_{\mathbf{v} \in \mathcal{G}} \left\{ \int_0^1 \|\mathbf{v}(t)\|_L dt \mid J \circ \mathbf{g}^{-1}(0) = J \wedge J \circ \mathbf{g}^{-1}(1) = I \right\}. \quad (4)$$

The differential Sobolev norm on $\mathbf{v}$ is determined by the associated linear (e.g. Cauchy-Navier) operator L . This is a true distance in that it is positive, symmetric and satisfies the triangle inequality [8]. The incremental integration of the velocity field is what gives this property (consider that the flow along the velocity field is, in the infinitesimal limit, piece-wise linear with equivalent norm forward and backward in time). In contrast, deformation-based norms taken from continuum mechanics are also positive and equal zero at the identity, but may not be symmetric.




Fig. 1. The “C” and “C patch” images (first and second in the top row) are registered with the large deformation diffeomorphic method in the top row, giving transformation $g(t)$. The result of linearly averaging the vector field u at $g(1)$ with Id is shown in the bottom row left and left middle. The averaging given by geodesic distances, at bottom row right, is more natural. Deformed grid images are also shown for the geodesic average. The jacobian of the C to C patch map is strictly positive with minimum value 0.14.

3.2 The Diffeomorphic Registration Algorithm

The optimization problem for the registration is to compute the mapping, g , such that,

$$g^* = \underset{v(t)}{\operatorname{argmin}} \left\{ \int_0^1 \|v\|_L^2 + \|Ig^{-1} - J\| dt \right\}, \quad (5)$$

is minimized where the brightness constancy assumption is used as the driving force [9]. Additionally, the maps are fixed to the identity at the boundary of the domain, $d\Omega$.

The Euler-Lagrange (E-L) equations for this problem were recently derived [8]. Rather than using the local E-L equations, we solve the variational problem in the integral form by using the Galerkin finite element (FE) method [10]. This method finds an optimal finite dimensional estimate to the infinite-dimensional solution. Using this method we compute the instantaneously optimal velocity field as,

$$v(\cdot, t) = \underset{v \in \mathcal{G}}{\operatorname{argmin}} \left\{ w_1 \|Ig^{-1} - J\| + w_2 \|v\|_L^2 \right\}, \quad (6)$$

which gives the optimal gradient in the space of diffeomorphic flows. The optimal estimate to the time-integrated map is approximated by using finite differences.

This requires parameterization of $\mathbf{g}$ by arc length, the unique and natural choice for (even infinite-dimensional) curves.

The time dependent integral for $\mathbf{g}$ is estimated with finite differences in time via the trapezoidal rule. Because the optimization process is not locally smooth in time, we measure the arc length at a step-size, h , that is larger than that of the optimization step-size $\|\delta\mathbf{v}\|$, with δ a small scalar. This is illustrated for a one dimensional map in figure 2 (the ratio is typically near 0.1). The incremental velocities are accumulated via gradient descent such that $\mathbf{v}^*(t) \approx \mathbf{Id} + \sum_{i=0}^{i=n} \delta\mathbf{v}_i$. The total map $\mathbf{g}(t)$ is then integrated to $\mathbf{g}(t+h)$ using $\mathbf{v}^*(t)$ when $\|\mathbf{v}^*(t)\|$ reaches the desired constant value, h . If the arc length oversteps h , a local line search on $\delta\mathbf{v}_n$ corrects the size of $\|\mathbf{v}^*\|$. The trapezoidal rule then gives the optimal in time approximation to the length as,

$$\int_{t_a}^{t_b} \|\mathbf{v}(t)\|_L dt \approx \sum_{i=1}^{i=n} \frac{h}{2} (\|\mathbf{v}(t_i + (i-1)h)\|_L + \|\mathbf{v}(t_i + ih)\|_L). \quad (7)$$

This gives a more robust estimate of the geodesic distance and is also beneficial for the shape averaging application (we have found the distances estimated by the greedy method in [7] are too noisy).

To summarize, the following steps are needed in the algorithm for computing geodesics:

1. Solve for the instantaneous regularized velocity field using the FE method.
2. Use the FE solution for gradient descent, accumulating $\delta\mathbf{v}$ until the arc length reaches h .
3. Integrate $\mathbf{v}$ to augment $\mathbf{g}$ in time and to compute a robust estimate of the geodesic distance.

Note the advantage that the velocity field is only needed at two time points, although an optimal in time solution is computed. Regridding is also performed, as in [7].




Fig. 2. The computation of the optimal $\mathbf{g}$ in time is performed using the trapezoidal rule, requiring constant arc length parameterization.

4 The Variational Problem for Diffeomorphic Averaging

We recover the average-shape image from an image dataset by solving an inverse problem involving the geodesic distances between the images and a boundary value constraint on the transformed image appearances. Find $\{g_i\}$ such that,

$$\begin{cases} \forall j \quad I_j \circ g_j^{-1} = I_j g_j^{-1} = \bar{I} \\ E(\{g_j\}) \end{cases} \text{ is minimal.} \quad (8)$$

The existence of geodesic paths on the diffeomorphism manifold as well as the symmetry of the distances given by those paths are important facts for this algorithm. Furthermore, these paths, as mentioned above, are also parameterized with constant arc length.

The simplest non-trivial case of this inverse problem is given by a single pair of images. First, consider the naive algorithm for minimization.

1. Register images I and J while measuring the distance, $D(I, J)(t)$, between them.
2. Repeat the registration stopping at $t_{\frac{1}{2}}$, where $D(I, J)(t_{\frac{1}{2}}) = \frac{1}{2}D(I, J)(t = 1)$.

This numerical minimization can be achieved successfully and was used to generate the results of figure 1. However, considering that, in practice, the optimization process continues in a coarse to fine fashion, one observes that coarse scale corrections will occur first in time. This is an undesirable bias that makes the averages appear visually incorrect when features exist at multiple scales, such as in anatomical images, as in figure 3. This caveat also makes the naive approach highly asymmetric.

The solution used is to solve the variational problem given in 8 explicitly. The variational problem for averaging pairs of images is then,

$$g_1^*, g_2^* = \underset{v_1(t)}{\operatorname{argmin}} \underset{v_2(t)}{\operatorname{argmin}} \left\{ \int_0^1 \|v_1\|_L^2 + \|v_2\|_L^2 + \|I g_1^{-1} - \bar{I}\| \, dt \right\} + \|J g_2^{-1} - \bar{I}\| \, dt \}.$$

Rearranging terms using the equality constraints given in the original problem gives,

$$g_1^*, g_2^* = \underset{v_1(t)}{\operatorname{argmin}} \underset{v_2(t)}{\operatorname{argmin}} \left\{ \int_0^1 \|v_1\|_L^2 + \|v_2\|_L^2 + \|I g_1^{-1} - J g_2^{-1}\| \, dt \right\}. \quad (9)$$

Solving this problem, via alternating minimization with respect to g_i and all h_i constant, provides average deformations that are optimized *symmetrically* using information at all scales. The geodesic averaging constraint $E(g_1) = E(g_2)$ is upheld by construction and the configurations $I g_1^{-1}$ and $J g_2^{-1}$ are both in average position. Note also that the transformation from I to J is given by $g_1^{-1} \circ g_2$. We will denote the output of this algorithm as $\mathcal{A}(\cdot, \cdot)$ where the input is a pair of images. Intuitively, the algorithm lets the images I and J “meet” at the mean configuration. A similar idea was introduced recently in [11], in which intensity averaging was incorporated.




Fig. 3. The top row shows, in the center, the geodesic shape average of the image on the left and the image on the far right. For comparison, the bottom row shows, in the center, the naive geodesic shape average, which is biased towards coarse-scale corrections. Note that the outline of the naive average is similar to the image on the right, but the ventricles are similar to the image on the left. The geodesic average improves upon this bias.

4.1 Recursive Algorithm for Diffeomorphic Averaging and Diffeomorphic Shape Interpolation

The constraints given above allow one to recursively compute either the minimal energy configuration over a dataset or, alternatively, the full-scale time-dependent shape interpolation between a pair of images. The latter application is a practical way to resolve the fact that measuring physical energies within a multiple-scale optimization process introduces a significant bias in the resulting computed intermediate mappings. That is, deformations near the zero time-point will accommodate coarse scale differences, while transformation adjustments occurring later in time will be at a fine scale.

The full-scale minimization of (8) is achieved as a recursive least squares estimation problem. That is, given A, B, C, D , averages may be computed as,

$$\mathcal{A}(A, B, C, D) = \mathcal{A}(\mathcal{A}(A, B), \mathcal{A}(C, D)), \quad (10)$$

where $\mathcal{A}$ is the averaging function. The symmetric diffeomorphic approach used here avoids exclusive dependence on the (possibly individual specific) topology of a specific template anatomy I_i , as in deformation field averaging. Furthermore, the group theoretical framework and geodesics used by our algorithm and those described in [6] will ensure the $\mathbf{g}_i$ are in $\mathcal{G}$ after composition. An average transport o.d.e. may also be used to minimize the energy, as in [12].

Pseudocode is given in algorithm 1. One is guaranteed that the final mean configuration is derived from the composition of diffeomorphic transformations,

insuring that topology of the mean anatomy is preserved and that the mean transformation exists in the shape space provided by the continuum model. This may be proven inductively by using the fact that if $g, h \in \mathcal{G}$ then $f = g \circ h \in \mathcal{G}$, that is, the composition of a geodesic diffeomorphic (g.d.) transformation with another g.d. transformation is within $\mathcal{G}$ and g.d. Note also that the set of transformations $\{g_i\}$ may be recovered after the end of the algorithm. For simplicity, we assume here that the size of the database is dyadic, though this is not a necessary condition. An illustration of a non-dyadic case is shown in figure 4.




Fig. 4. The progress of the recursive averaging algorithm is shown above. Note the adjustment made for the uneven number of averages performed.

A similar algorithm may be used for geodesic shape interpolation, an example of which is shown in figure 5.

Algorithm 1 : Diffeomorphic Shape Averaging

Divide the database of images into unique pairs, $\mathcal{P} = \{\mathcal{P}_i = (I_i, I_j) | i \neq j\}$. Set $i = 1$.
 Denote the cardinality of $\mathcal{P}$ as $\#\mathcal{P}$.
while $i < \#\mathcal{P}$ **do**
 $\mathcal{P} \cup (\mathcal{A}(\mathcal{P}_i), \mathcal{A}(\mathcal{P}_{i+1}))$.
 $i = i + 2$.
end while
 $\bar{I} = \mathcal{A}(\mathcal{P}_{\#\mathcal{P}})$.




Fig. 5. Three time-points of scale consistent shape interpolation are generated from a pair of anatomical instances. The pair-wise symmetric registration method was used to generate the results here and to insure that the interpolated anatomy is consistent for all scales of information. The original images are at top left and bottom left respectively. This anatomical pair is in correspondence at time 0.5, which is the average of the anatomy at time 0 with the anatomy at time 1. Note that here we have used the topology of the closest anatomy at all points. It would also be natural to use an intensity average where the relative weights are determined by the time values.

5 Experiments

As was shown in figure 1, figure 3 and figure 5, the symmetric geodesic averaging algorithm is essential in some cases. We now investigate the effect of diffeomorphic averaging on normal anatomical images. A dataset of 6 normal female human cortices were manually segmented from volumetric MRI. The cortices were initially aligned by similarity transformation to an arbitrarily chosen reference. The current experiments are in two dimensions, although the implementation is equally functional in three dimensions.

5.1 Comparison of Geodesic and Linear Averaging

Linear Average. The diffeomorphic fluid algorithm was used to register an arbitrarily chosen reference topology to the dataset. The deformation fields provided by the non-rigid registration were then averaged in the Lagrangian reference frame, meaning vectors from the original to the final configuration were used. The average transformation was inverted to find the average shape atlas. The registration was repeated and the root mean square distance from the average shape atlas to the dataset was computed. A second reference image was also chosen and the study was repeated. The overlap ratio between the pair of average shape cortices is also computed. The results are summarized in figure 6(c,d).

Geodesic Average. The same registration algorithm was used with the geodesic averaging procedure of algorithm 1. The computational cost of this study is a logarithmic factor larger than the linear averaging, but without the (not costly, but potentially error prone) step of estimating the average transformation and its inverse. Because the algorithm moves from dataset instance to dataset instance, the order of the dataset was randomized and the study repeated. Results for these studies are summarized in figure 6(a,b). The similarity in appearance between the curved and linear averages suggest that the transformations computed are not highly curved, in contrast to those shown in figure 1. Finally, Table 1 shows a summary of metrics computed from the atlases to the dataset and also between the two instances of the atlases.

Algorithm	Overlap between 1 and 2	Intensity Squared Difference (SSD)
Geodesic	0.989	0.478
Linear	0.990	0.503

Table 1. Summary of Algorithm Dependency on Dataset. The overlap and sum of squared differences in intensity (SSD) are measured between the atlases generated by the same method but with a different dataset ordering and/or reference anatomy.

6 Conclusion

We have described an algorithm for geodesic shape averaging and interpolation, argued its correctness and illustrated its results. Furthermore, a finite element




Fig. 6. The geodesic and linear approach to shape averaging are compared. The template anatomy for the first geodesic average example is the same as that of the first linear example. The second examples also share the same topology. One may note that, qualitatively, the geodesic results appear to be more plausible shapes, as the anatomy is deformed more smoothly. (a) Geodesic shape average one. (b) Geodesic shape average two. (c) Linear shape average one. (d) Linear shape average two.

method algorithm for estimating geodesic distances between images was given. This algorithm is symmetric with respect to an image pair without the need for explicit penalization of the difference between the forward and inverse transformations. It also has the property that it divides the deformation needed for correct registration evenly between each image, as is guaranteed by robust measures of the geodesic distance. Future work will investigate improvements to the current geodesic estimation scheme as well as its properties as an image registration algorithm on its own. We also intend to investigate solving problem (8) in parallel. This method will combine an initial estimate to $\bar{I}$ with the symmetric image registration algorithm given here. The solution will then be a piecewise linear estimate to the average curved transformation which can be compared with linear averaging and the recursive estimation in this work.

References

- [1] R. Kikinis, M. E. Shenton, D. V. Iosifescu, R. W. McCarley, P. Saiviroonporn, H. H. Hokama, A. Robatino, D. Metcalf, C. G. Wible, C. M. Portas, R. M. Donino, and F. A. Jolesz, "A digital brain atlas for surgical planning, model-driven segmentation, and teaching," *IEEE Trans. Visualization and Comp. Graph.*, vol. 2, pp. 232–241, 1996.
- [2] J. Talairach and P. Tournoux, *Coplanar stereotaxic axis of the human brain*, Thieme, New York, 1988.
- [3] J. Ashburner and K. Friston, "Fully three-dimensional nonlinear spatial normalization: A new approach," in *2nd Int. Conf. Functional Mapping of the Hum. Brain*, A. W. Toga, R. S. J. Frackowiak, and J. C. Mazziotta, Eds., 1996, vol. 3, p. S111.
- [4] L. Le Briquer and J. C. Gee, "Design of a statistical model of brain shape," in *Information Processing in Medical Imaging*, J. S. Duncan and G. Gindi, Eds., pp. 477–482. Springer-Verlag, Heidelberg, 1997.
- [5] A. Guimond, J. Meunier, and J.-P. Thirion, "Average brain models: A convergence study," *Computer Vision and Image Understanding*, vol. 77, no. 2, pp. 192–210, 2000.
- [6] M. I. Miller and L. Younes, "Group actions, homeomorphisms and matching: a general framework," *Int. J. Computer Vision*, vol. 41, pp. 61–84, 2001.
- [7] G. E. Christensen, R. D. Rabbitt, and M. I. Miller, "A deformable neuroanatomy textbook based on viscous fluid mechanics," in *27th Annual Conference on Information and Systems*, J. L. Prince and T. Runolfsson, Eds., Baltimore, MD, 1993, pp. 211–216.
- [8] M. Miller, A. Trounev, and L. Younes, "On the metrics and euler-lagrange equations of computational anatomy," *Annu. Rev. Biomed. Eng.*, vol. 4, pp. 375–405, 2002.
- [9] P. Dupuis, U. Grenander, and M. I. Miller, "Variational problems on flows of diffeomorphisms for image matching," *Quarterly of Applied Mathematics*, vol. 56, no. 3, pp. 587–600, 1998.
- [10] O. C. Zienkiewicz, *The Finite Element Method in Engineering Science*, McGraw-Hill, New York, 1971.
- [11] B. Davis, P. Davies, and S. Joshi, "Large deformation minimum mean squared error template estimation for computational anatomy," *IEEE International Symposium on Biomedical Imaging*, 2004.
- [12] B. Avants and J.C. Gee, "Shape averaging with diffeomorphic flows for atlas creation," *IEEE International Symposium on Biomedical Imaging*, 2004.