

Example-Based Stereo with General BRDFs

Adrien Treuille¹, Aaron Hertzmann², and Steven M. Seitz¹

¹ University of Washington, Seattle, WA, USA
{treuille, seitz}@cs.washington.edu

² University of Toronto, Toronto, ON, Canada
hertzman@dgp.toronto.edu

Abstract. This paper presents an algorithm for voxel-based reconstruction of objects with general reflectance properties from multiple calibrated views. It is assumed that one or more reference objects with known geometry are imaged under the same lighting and camera conditions as the object being reconstructed. The unknown object is reconstructed using a radiance basis inferred from the reference objects. Each view may have arbitrary, unknown distant lighting. If the lighting is calibrated, our model also takes into account shadows that the object casts upon itself. To our knowledge, this is the first stereo method to handle general, unknown, spatially-varying BRDFs under possibly varying, distant lighting, and shadows. We demonstrate our algorithm by recovering geometry and surface normals for objects with both uniform and spatially-varying BRDFs. The normals reveal fine-scale surface detail, allowing much richer renderings than the voxel geometry alone.

1 Introduction

Recovering an object’s geometry from multiple camera orientations has long been a topic of interest in computer vision. The challenge is to reconstruct high-quality models for a broad class of objects under general conditions. Most prior multiview stereo algorithms assume Lambertian reflectance—the radiance at any point is the same in all directions—although progress has begun on relaxing this assumption, e.g. [1,2,3,4]. Also, virtually all approaches to multiview stereo ignore non-local light phenomena, such as interreflections and cast shadows. Because these simplifying assumptions do not hold for large classes of objects, new algorithms are needed.

In this paper, we present an approach for the volumetric reconstruction of objects with *general reflectance*. That is, we do not assume a particular BRDF model. Instead, we assume that one or more reference objects of related material are observed under the same conditions as the object being reconstructed. Our experimental setup handles surfaces with isotropic BRDFs, although a more general setup could in theory be used to capture anisotropic BRDFs as well. Our algorithm also handles varying camera and light positions: we assume only that the cameras and lights are distant and separated from the object by a plane. In addition, we show how to account for cast shadows in voxel coloring. We note

that cast shadows occur when a voxel is not “visible” to a light source. In this way, we can extend voxel coloring’s treatment of visibility to handle shadows as well.

Our approach is based on the *orientation-consistency* cue which states that, under orthographic projection and distant lighting, two surface points with the same surface normal and material exhibit the same radiance. This cue was introduced in the context of photometric stereo [5]. We adapt orientation-consistency to multiview reconstruction within the voxel coloring framework [6]. We chose voxel coloring mainly for simplicity, although orientation-consistency could be used with other stereo algorithms as well. Besides handling general BRDFs, orientation-consistency enables computing per-voxel normals, which is not possible in conventional voxel coloring. These normals reveal fine surface details that the voxel geometry alone does not capture.

2 Related Work

Most previous work in multiview stereo focuses on diffuse objects, e.g. [7,8,9,10,11], although progress has recently been made on treating other types of surfaces. In particular, recent work has addressed the case of completely specular objects. Zheng and Murata [12] reconstruct purely specular objects; Bonfort and Sturm [13] and Savarese and Perona [14] study the case where the surface is a mirror.

Stereo methods have also been proposed for diffuse-plus-specular surfaces. Carceroni and Kutulakos [1] reconstruct moving objects assuming a Phong model with known specular coefficient. They propose an ambitious, complex procedure combining discrete search and nonlinear optimization. Jin *et al.* [2] exploit a BRDF constraint implied by a diffuse-plus-specular model which requires a fixed light source. Yang *et al.* [3] adapt the space carving algorithm for a diffuse-plus-specular model by making heuristic assumptions on the BRDF and illumination; they assume that all observations lie on a line in color space. They also require a fixed light source and color images. All of these algorithms make some diffuse-plus-specular assumption on the reflectance. Our technique improves on these in that we consider a broader class of BRDFs. On the other hand, these techniques do not require a reference object.

The only known stereo method that handles completely general BRDFs is that of Zickler *et al.* [4], which exploits Helmholtz reciprocity, and demonstrates high quality results. Their method applies to perspective projection, and does not need a reference object. However, they require a constrained illumination setup with point light sources and cameras reciprocally placed, whereas we allow arbitrary, distant cameras and illumination. Finally, they do not take into account shadows, though we believe our shadow technique can be used for Helmholtz stereo as well.

Our algorithm also handles completely general BRDFs, by using a cue called *orientation-consistency*. This cue, proposed by Hertzmann and Seitz for photometric stereo, has been shown to work for BRDFs as complex as velvet and brushed fur [5]. Other techniques in photometric stereo also handle non-Lambertian

BRDFs, e.g. [15,16,17]. However, all of these photometric stereo approaches are constrained by solving only for a normal map, which can lead to geometric distortions and makes depth discontinuities troublesome. Instead of just a normal map, we solve for a full object model with normals.

Our work integrates orientation-consistency into the the voxel coloring algorithm [6], which was originally designed for Lambertian surfaces. The algorithm handles visibility, but requires constraints on camera placement. Space carving [7] relaxes the constraints on camera placement. We adapt voxel coloring to apply to general BRDFs, and show how voxel coloring’s treatment of visibility can be extended to handle shadows. We also adapt the camera placement constraints to orthographic projection. Generalizing our work to space carving is straightforward; we chose voxel coloring because it is simpler than space carving, and provides a testbed for evaluating the novel features presented in this paper.

While voxel coloring uses voxels to represent the geometry, level sets are becoming an increasingly popular geometric representation in computer vision. In particular, Faugeras and Keriven [11] variationally solve for a level set describing the geometry. We believe our technique could be applied to level-set stereo by integrating orientation-consistency into the objective function; other diffuse multiview stereo algorithms could also benefit from orientation-consistency.

Most previous multiview stereo methods do not explicitly handle cast shadows. One exception is the work of Savarese *et al.* [18] who demonstrate a volumetric carving technique that uses only shadows. Their work differs from ours in that they assume shadows can be detected *a priori*, and in that they do not make use of reflectance information.

3 Reconstructing Objects with a Single Material

We now show how we adapt voxel coloring to reconstruct objects with general BRDFs. The central component of voxel coloring is the *photo-consistency* test, which determines if a voxel is consistent with the input photographs. Conventional voxel coloring uses a test that is suitable only for Lambertian surfaces. We replace this test with *orientation-consistency*, which applies to general BRDFs.

We begin with a brief summary of voxel coloring. The target object must be photographed to satisfy the *ordinal visibility constraint*. This condition ensures that there exists a traversal order of the voxels so that occluding voxels are always visited before those that they occlude. For example, in our experiments the camera is placed above the target object, and we traverse the voxels layer-by-layer from top to bottom. A consistency test is applied to each voxel ν in turn. If ν is deemed *consistent* with all input images, then it is included in the volumetric model. Otherwise, ν is *inconsistent* and is discarded. The consistency test only considers views in which the voxel is not occluded (see Section 3.3).

3.1 Diffuse Photo-consistency

We call the consistency metric of the original voxel coloring algorithm *diffuse photo-consistency*. The test consists of projecting the voxel ν into each input image to produce a vector $\mathbf{V}_\nu$ of intensities

$$\mathbf{V}_\nu = [I_{1,\nu}, I_{2,\nu}, \dots, I_{n,\nu}]^T \quad (1)$$

which we call a target *observation vector*. For color images, we have separate observation vectors $\mathbf{R}_\nu$, $\mathbf{G}_\nu$, and $\mathbf{B}_\nu$ for the red, green, and blue channels, respectively. These can be concatenated into the *color observation vector*:

$$\mathbf{V}_\nu = [\mathbf{R}_\nu^T, \mathbf{G}_\nu^T, \mathbf{B}_\nu^T]^T. \quad (2)$$

The voxel is consistent if all intensities are nearly the same, measured by testing whether the sum of the variance in intensity over all color channels falls below a specified threshold. The algorithm ensures that only views in which the voxel is visible are considered during the variance test.

3.2 Orientation-Consistency

Diffuse photo-consistency is physically meaningful for Lambertian surfaces, but does not apply to more general BRDFs, since we cannot assume that the radiance from a voxel will be the same in all directions. Instead, we propose an example-based consistency metric called *orientation-consistency*, which is adapted from [5].

To begin, assume that the target object consists of only one material and that a reference object of the same material and with known geometry has been observed in all the same viewpoints and illuminations as the target object. Moreover, assume the reference object exhibits the full set of visible normals for each camera position. Now consider a voxel ν on the surface of the target object. There must exist a point $\mathbf{p}$ on the reference object that has the same normal, and, thus, the same observation vector. Hence, we can test consistency by checking if such a point $\mathbf{p}$ exists.

In practice, we densely sample the surface of the reference object and project these sampled points $\mathbf{p}_1, \dots, \mathbf{p}_k$ into the input images to form set of reference observation vectors $\mathbf{V}_{\mathbf{p}_1}, \dots, \mathbf{V}_{\mathbf{p}_k}$. We use this database of reference observation vectors to determine if a given target observation vector is consistent with some normal. Formally, a voxel ν with observation vector $\mathbf{V}_\nu$ is *orientation-consistent* if there exists a point $\mathbf{p}_i$ such that

$$\frac{1}{d} \|\mathbf{V}_\nu - \mathbf{V}_{\mathbf{p}_i}\|^2 < \epsilon \quad (3)$$

for some user-determined threshold ϵ . Division by d , the number of dimensions in the vectors, gives the average squared error. We note that d may differ for different observation vectors: when a target voxel ν is occluded or in shadow for

some camera, the corresponding dimensions of the target and reference observation vectors are excluded.

We note some special features of orientation-consistency. First, unlike diffuse photo-consistency, orientation-consistency is not limited to diffuse BRDFs. In fact, the technique handles arbitrary isotropic BRDFs. Moreover, by finding the exact reference point $\mathbf{p}_i$ that minimizes Equation (3), we can assign the corresponding surface normal to each consistent voxel. Later, we can render these normals, as in bump mapping [19], to reveal much more visual detail than the voxel geometry itself contains.

3.3 Orthographic Ordinal Visibility

Our algorithm depends crucially on being able to determine for which cameras a voxel is visible. Seitz and Dyer showed that visibility could be determined if the voxels are traversed so that occluding voxels are always processed before those they occlude. This implies an *ordinal visibility constraint* on camera placement for perspective projection [6]. We use the same approach, but adapted for orthographic projection.

Suppose we have n cameras pointing in directions $\mathbf{c}_1, \dots, \mathbf{c}_n$, where each of the $\mathbf{c}_i$ are unit vectors. Our constraint on camera placement is that there exists a vector $\mathbf{v}$ such that $\mathbf{v} \cdot \mathbf{c}_i > 0$ for each camera direction $\mathbf{c}_i$ (Fig. 1). Informally, we can think of a plane with normal $\mathbf{v}$ separating the cameras from the object, and we iterate over voxels by marching in the direction of $\mathbf{v}$. More formally, the order is defined so that we visit point $\mathbf{a}$ before point $\mathbf{b}$ if $(\mathbf{b} - \mathbf{a}) \cdot \mathbf{v} > 0$.




Fig. 1. Orthographic Ordinal Visibility: A plane with normal $\mathbf{v}$ separates the camera from the scene. Each point $\mathbf{a}$ is processed before any point $\mathbf{b}$ that it occludes.

To see why this inequality processes the points in the correct order, suppose that some point $\mathbf{a}$ occludes a point $\mathbf{b}$ for camera direction $\mathbf{c}_i$. Then $\mathbf{b} - \mathbf{a} = \alpha \mathbf{c}_i$ with $\alpha > 0$. Applying a dot product with $\mathbf{v}$ we get $(\mathbf{b} - \mathbf{a}) \cdot \mathbf{v} = \alpha \mathbf{c}_i \cdot \mathbf{v}$. Since $\mathbf{c}_i \cdot \mathbf{v}$ is positive by assumption, $(\mathbf{b} - \mathbf{a}) \cdot \mathbf{v} > 0$ and $\mathbf{a}$ is processed before $\mathbf{b}$, by definition of the ordering.

3.4 Single Material Results

We now describe our experimental setup and present results. Our target object was a soda bottle, and our reference object was a snooker ball. Both objects were spray-painted with a shiny green paint to make the materials match. We then photographed the objects under three different lighting conditions from a fixed camera, while varying the object orientation with a turntable. The photographs were taken with a zoom lens at 135mm, and the light sources were placed at least five meters away. This setup approximates an orthographic camera and directional lighting, which are the assumptions of orientation-consistency. The camera was calibrated using a freely-available toolbox [20].




Fig. 2. (a) Reference spheres shown for the 3 different illuminations. (b) Corresponding input images for one of 30 object orientations. (c) Voxel reconstruction using three lights. (d) Rendered normals obtained from one light source. (e) Rendered normals obtained from three light sources.

Our input consisted of 30 object orientations with 3 illumination conditions each, for a total of 90 input images. Using fewer than 3 lights adversely affects the recovered normals, though it does not seem to affect the geometry. We also note that, by symmetry, the appearance of the sphere does not change as it rotates on the turntable. Therefore we need only one image of the sphere per illumination condition.

Fig. 2(a) shows input images of the reference sphere for the 3 different lighting conditions. Fig. 2(b) shows corresponding input images of the bottle. The voxel reconstruction can be seen in Fig. 2(c). Finally, Fig. 2(d) and 2(e) show the model with the recovered normals. The improvement of 2(e) over 2(d) is achieved by using multiple illuminations.

The ability to recover normals is one of the strengths of our algorithm as it reveals fine-grained surface texture. The horizontal creases in the label of the bottle are a particularly striking example of this: the creases are too small to be captured by the voxels, but show up in the normals. Fig. 3 shows different views rendered from the reconstruction.




Fig. 3. Views of the bottle reconstruction.

4 Generalizations

While the technique described in the previous section makes voxel coloring possible for general BRDFs, it suffers from several limitations. First, our consistency function does not take into account shadows that the object casts upon itself. In addition, the model works only in the restrictive case that the target object has a single BRDF shared by the reference object. In this section, we show how these assumptions are relaxed.

4.1 Handling Shadows

Unlike previous work in multiview stereo, our framework allows the illumination to vary arbitrarily from image to image. Given this setup, it is almost inevitable that, in at least some views, a complex object will cast shadows on itself. We address this phenomenon by treating shadowed voxels as if they were occluded, that is, by excluding the shadowed dimensions from the observation vector. Thus, we need a method to determine if a voxel is in shadow.

We note that shadows occur when a voxel is not “visible” to a light source, and we can therefore compute shadows exactly as we do occlusions. We assume that the light directions are calibrated. Just as each image has an *occlusion mask* aligned with the camera (see [6]), we now say that each image also has a *shadow mask* aligned with the light. Voxels can be projected onto this mask as if the light were a camera. The occlusion and shadow masks are initially empty. When a voxel is deemed consistent, it is projected onto all occlusion masks and all shadow masks. The set of pixels to which the voxel projects, its *footprint* is marked to exclude these regions from future computation; subsequent voxels projecting onto marked regions are considered occluded or in shadow, as the case may be. Of course, this technique implies that voxels casting shadows must be processed before those which they shadow, but this is equivalent to the corresponding visibility requirement, and the same theory applies; both the cameras *and* lights must be placed so that their directions satisfy the ordinal visibility constraint of Section 3.3.

Note that this technique of detecting shadows is completely independent of the choice of consistency test, and can be implemented with other forms of voxel coloring. In particular, by using perspective projection and the conventional ordinal visibility constraint [6], point light sources can be handled.

4.2 Varying BRDFs

Our second generalization relaxes the assumption that the target and reference objects consist of the same single material. Instead, we use a *basis* of reference objects with BRDFs related to that of the target object. As in [5], we assume that the colors observed on the target object can be expressed as a linear combination of observation vectors from the reference objects.

As in Section 3.4, we use spheres as reference objects. Suppose that $\mathbf{p}_1, \dots, \mathbf{p}_k$ is a set of points *in sphere surface coordinates*, so that we may talk about the same point on multiple spheres. For every point $\mathbf{p}_i$ there are s observation vectors $\mathbf{V}_{\mathbf{p}_i}^1, \dots, \mathbf{V}_{\mathbf{p}_i}^s$, one for each of the s spheres, which we concatenate into an *observation matrix*:

$$\mathbf{W}_{\mathbf{p}_i} = [\mathbf{V}_{\mathbf{p}_i}^1, \dots, \mathbf{V}_{\mathbf{p}_i}^s]. \quad (4)$$

We assume a voxel is consistent if it can be explained by some normal and some material in the span of the reference spheres. Formally, a voxel ν with observation vector $\mathbf{V}_\nu$ is orientation-consistent with respect to the reference spheres if there exists a point $\mathbf{p}_i$ and a *material index* $\mathbf{m}$ such that

$$\frac{1}{d} \|\mathbf{V}_\nu - \mathbf{W}_{\mathbf{p}_i} \mathbf{m}\|^2 < \epsilon \quad (5)$$

for a user-specified threshold ϵ . As in Equation (3), d is the number of dimensions left after deleting all occluded and shadowed images from the observation vectors. Note that while we find the best point $\mathbf{p}_i$ by linearly searching through our database of samples, the best material index $\mathbf{m}$ can be directly computed for each $\mathbf{p}_i$ using the pseudo-inverse (+) operation:

$$\mathbf{m} = (\mathbf{W}_{\mathbf{p}_i})^+ \mathbf{V}_\nu. \quad (6)$$

In summary, the algorithm is as follows: for each target voxel, we test consistency by iterating over all possible source points $\mathbf{p}_i$. For each source point, we compute the optimal material index $\mathbf{m}$. If any pair $(\mathbf{p}_i, \mathbf{m})$ satisfies Equation (5), then the voxel ν is added to the reconstruction; otherwise it is discarded.

4.3 Multiple Material Results

We now present experimental results for the generalized algorithm. For the experiments, we photographed two spheres, one matte gray, and the other specular black. We decomposed the gray sphere into a red, green, and blue diffuse basis; the black sphere was used to handle specularities. To calibrate the lights for the shadow computation, we photographed a mirrored ball and estimated the reflection of the light direction by computing the centroid of the brightest pixels.

We reconstructed two target objects. The first was a porcelain cat. Fig. 4(a) shows an input photograph. Fig. 4(b) is the voxel reconstruction, and Fig. 4(c) shows the reconstruction rendered with the recovered normals. Other views of the reconstruction are shown in Fig. 4(d). The black diagonal line across the




Fig. 4. Cat model. (a) Input image. (b) Voxel reconstruction. (c) Rendered with normals. (d) New views.

front of the cat in Fig. 4(d) is an artifact caused by insufficient carving. The true surface lies several voxels behind the recovered surface, but the algorithm has not been able to carve that far. Setting the consistency threshold lower would have prevented this artifact, but caused overcarving in other parts of the geometry. In both this model and the next, overestimation of the geometry added some noise to the normals, because the algorithm was trying to fit normals to points not on the surface. We discuss this issue further in Section 5.

The second target object was a polished marble rhinoceros. As with the cat, Fig. 5(a) shows an input photograph; Fig. 5(b) is the voxel reconstruction, and Fig. 5(c) shows the reconstruction with the recovered normals. Note that the normals reveal the three chisel marks on the side of the rhino in Fig. 5(c). Finally, Fig. 5(d) shows additional views.

We now highlight two aspects of the technique. First, to show that our algorithm can correctly carve past the visual hull, Fig. 6 shows a close-up of the hind legs of the rhino model. Looking, in particular, at the gap between the legs, we can see that the geometry is better approximated by our algorithm (Fig. 6(c)) than by the visual hull alone (Fig. 6(a)), even for this relatively untextured re-




Fig. 5. Rhino Model. (a) Input image (one of 120). (b) Voxel reconstruction. (c) Rendered with normals. (d) New views.




Fig. 6. Detail of the hind legs of the rhino model. Note the gap between the legs. (a) Visual Hull. (b) Photograph not from input sequence. (c) Reconstructed model.

gion. Note that the comparison photograph, Fig. 6(b), could not have been used as input, as it would have violated the ordinal visibility constraint. To carve this area, we used a lower consistency threshold than that used in Fig. 5, which resulted in overcarving of other parts of the model. We expect that an adaptive threshold technique like [9] could address this problem.

The second aspect we highlight is the shadowing technique. Specifically, Fig. 7 provides a visualization of the shadows. Fig. 7(a) shows one input image from the rhino sequence, and Fig. 7(b) shows the shadow voxels computed for that image. Note the shadows cast by the right ear, and those on the right fore and hind legs.

These show that we are getting a good approximation of the shadows. However, the shadows would improve further if the geometry were better estimated.




Fig. 7. (a) Input image. (b) Shadows (dark regions) computed for this image.

5 Discussion and Future Work

This paper presented a novel volumetric reconstruction algorithm based on orientation-consistency. We assume only that the cameras and lights are distant, and separated from the object by a plane. We showed how to integrate orientation-consistency into voxel coloring; other multiview techniques would have been possible as well. Our ability to solve for normals within the voxel framework yields a dramatic improvement in the representation of fine-scale surface detail. Although our experimental setup is designed for isotropic BRDFs, a more general setup could be used for anisotropic BRDFs, as was shown in [5] for the photometric case.

We also showed how voxel coloring can be adapted to handle cast shadows if the lights are calibrated. As a step toward this end, we adapted the ordinal visibility constraint to the orthographic case. We view our treatment of shadows as a first step toward integrating non-local lighting phenomena into traditional reconstruction techniques.

The main difficulty we encountered is carving past the visual-hull. Two issues may be simultaneously contributing to this problem. First, voxel coloring is hampered by trying to find a *consistent* model, which can be weaker than minimizing an error metric on the model. Second, our algorithm works better for a single material than for multiple materials, possibly because the target materials are not sufficiently well modeled by the reference materials. As a result, the consistency threshold must be set high, and too few voxels are carved. Better results may be possible using orientation-consistency in conjunction with more recent multi-view stereo techniques such as [9] or [11]. Another option would be to use better reference objects, but the good results obtained in [5] indicate that

the linear combination model need not perfectly match the target material. We also believe that better results could be obtained using more highly-textured surfaces, because voxels not on the surface are more clearly inconsistent for highly-textured surfaces.

Previous work using orientation-consistency [5] yielded results with more detail and less noise than ours. We believe the explanation is that the results in [5] were run at a higher resolution (one normal per pixel) and required estimating fewer parameters. A key advantage of our work, however, is that we create full object models as opposed to a single depth map. While, in principle, multiple depth maps created from photometric stereo methods such as [5] could be merged together into a full object model, we expect that global distortions incurred from normal integration errors would make such a merging procedure difficult. An additional advantage over [5] is that we properly account for cast shadows.

In general, this work suggests that bridging the gap between photometric techniques such as orientation-consistency and multiview techniques such as voxel coloring is a promising avenue in geometric reconstruction. We have found that recovering normals yields much finer surface detail than is possible with stereo methods alone, but an important open problem is finding better ways of merging geometric and photometric constraints on shape.

Acknowledgements. This work was supported in part by NSF grants IIS-0049095 and IIS-0113007, an ONR YIP award, a grant from Microsoft Corporation, the UW Animation Research Labs, an NSERC Discovery Grant, the Connaught Fund, and an NSF Graduate Research Fellowship. Portions of this work were performed while Aaron Hertzmann was at the University of Washington.

References

1. Carceroni, R.L., Kutulakos, K.N.: Multi-view scene capture by surfel sampling: From video streams to non-rigid 3d motion, shape and reflectance. *International Journal of Computer Vision* **49** (2002) 175–214
2. Jin, H., Soatto, S., Yezzi, A.: Multi-view stereo beyond Lambert. In: *Proceedings of the 9th International Conference on Computer Vision*. (2003)
3. Yang, R., Pollefeys, M., Welch, G.: Dealing with textureless regions and specular highlights—a progressive space carving scheme using a novel photo-consistency measure. In: *Proceedings of the 9th International Conference on Computer Vision*. (2003)
4. Zickler, T.E., Belhumeur, P.N., Kriegman, D.J.: Helmholtz stereopsis: Exploiting reciprocity for surface reconstruction. *International Journal of Computer Vision* **49** (2002) 215–227
5. Hertzmann, A., Seitz, S.M.: Shape and materials by example: A photometric stereo approach. In: *Conference on Computer Vision and Pattern Recognition*. (2003) 533–540
6. Seitz, S.M., Dyer, C.R.: Photorealistic scene reconstruction by voxel coloring. *International Journal of Computer Vision* **35** (1999) 151–173

7. Kutulakos, K.N., Seitz, S.M.: A theory of shape by space carving. *International Journal of Computer Vision* **38** (2000) 199–218
8. Kolmogorov, V., Zabih, R.: Multi-camera scene reconstruction via graph cuts. In: 7th European Conference on Computer Vision. (2002) 82–96
9. Broadhurst, A., Drummond, T., Cipolla, R.: A probabilistic framework for space carving. In: Proceedings of the 8th International Conference on Computer Vision. (2001) 388–393
10. Slabaugh, G.G., Culbertson, W.B., Malzbender, T., Stevens, M.R., Schafer, R.W.: Methods for volumetric reconstruction of visual scenes. *International Journal of Computer Vision* **57** (2004) 179–199
11. Faugeras, O.D., Keriven, R.: Complete dense stereovision using level set methods. In: 5th European Conference on Computer Vision (1998). (1998) 379–393
12. Zheng, J., Murata, A.: Acquiring a complete 3D model from specular motion under the illumination of circular-shaped light sources. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **22** (2000) 913–920
13. Bonfort, T., Sturm, P.: Voxel carving for specular surfaces. In: Proceedings of the 9th International Conference on Computer Vision. (2003)
14. Savarese, S., Perona, P.: Local analysis for 3d reconstruction of specular surfaces - part ii. In: 7th European Conference on Computer Vision. (2002) 759–774
15. Silver, W.M.: Determining shape and reflectance using multiple images. Master's thesis, MIT, Cambridge, MA (1980)
16. Woodham, R.J.: Photometric method for determining surface orientation from multiple images. *Optical Engineering* **19** (1980) 139–144
17. Wolff, L.B., Shafer, S.A., Healey, G.E., eds.: *Physics-based vision: Principles and Practice, Shape Recovery*. Jones and Bartlett, Boston, MA (1992)
18. Savarese, S., Rushmeier, H., Bernardini, F., Perona, P.: Shadow carving. In: Proceedings of the 8th International Conference on Computer Vision. (2001)
19. Blinn, J.F.: Simulation of wrinkled surfaces. In: *Computer Graphics (Proceedings of SIGGRAPH)*. Volume 12. (1978) 286–292
20. Bouguet, J.Y.: Camera calibration toolbox for matlab.
<http://www.vision.caltech.edu/bouguetj/calib.doc/> (2004)