

Binocular Helmholtz Stereopsis

Todd E. Zickler[†] Jeffrey Ho[‡] David J. Kriegman[‡] Jean Ponce[§] Peter N. Belhumeur[¶]
zickler@yale.edu jho@cs.ucsd.edu kriegman@cs.ucsd.edu ponce@cs.uiuc.edu belhumeur@cs.columbia.edu

[†]Electrical Engineering, Yale University, New Haven, CT 06511

[‡]Computer Science and Engineering, University of California at San Diego, La Jolla, CA 92093

[§]Computer Science, University of Illinois at Urbana-Champaign, Urbana, IL 61801

[¶]Computer Science, Columbia University, New York, NY 10027

Abstract

Helmholtz stereopsis has been introduced recently as a surface reconstruction technique that does not assume a model of surface reflectance. In the reported formulation, correspondence was established using a rank constraint, necessitating at least three viewpoints and three pairs of images. Here, it is revealed that the fundamental Helmholtz stereopsis constraint defines a nonlinear partial differential equation, which can be solved using only two images. It is shown that, unlike conventional stereo, binocular Helmholtz stereopsis is able to establish correspondence (and thereby recover surface depth) for objects having an arbitrary and unknown BRDF and in textureless regions (i.e., regions of constant or slowly varying BRDF). An implementation and experimental results validate the method for specular surfaces with and without texture.

1 Introduction

Helmholtz stereopsis has recently been introduced as a multi-view technique for estimating surface shape *without requiring an assumed model of reflectance*. Thus, unlike most existing methods, it enables the dense reconstruction of scenes that contain surfaces with unknown, spatially varying, and arbitrary reflectance functions (BRDFs). In Helmholtz stereopsis, this is accomplished through the use of images that are collected in reciprocal pairs [21]. The first image in a reciprocal pair is acquired under a single point light source, and the second image is acquired after interchanging the camera and light source positions. Figure 1 shows a reciprocal pair, and Fig. 2 illustrates the acquisition geometry. Reciprocal pairs have the unique property that the relation between intensities at corresponding image points depends only on surface shape and is independent of reflectance. This property follows directly from the symme-




Figure 1. Two rectified images of a painted, plastic mannequin head acquired as a Helmholtz reciprocal pair. Note the prominent specularities.

try of every bidirectional reflectance distribution function (BRDF) and is known as Helmholtz reciprocity [9, 15]. So while specular highlights generally move over the surface for changes of viewpoint under fixed lighting, in reciprocal pairs they correspond to the projection of fixed surface points and essentially become features for determining correspondence.

As originally formulated in [14, 21], stereo correspondence is established using a rank constraint that requires at least three reciprocal image pairs, and in turn this implies at least three camera locations and six images. (In fact, 36 images were used to great effect in [21].) Once correspondence is established, the same constraints can be used to estimate the normal field of the surface without the need for differentiation.

In this paper, we re-examine the constraint arising from a single reciprocal pair of images and reveal that it defines a partial differential equation (PDE). We show that this PDE can be solved to provide an accurate reconstruction of surface shape. This new reconstruction technique, *binocular*

Helmholtz stereopsis, offers the following two significant advantages over conventional stereopsis.

1. Binocular Helmholtz stereopsis is able to reconstruct surfaces with arbitrary and unknown BRDF's (including very specular surfaces), whereas conventional dense stereo correspondence is predicated on a constant brightness assumption (i.e., that the BRDF is Lambertian).
2. Binocular Helmholtz stereopsis is able to establish correspondence in textureless regions, whereas conventional stereo can only "guess" correspondence in such regions using a regularization or smoothing process.

The skeptical reader might turn to the reconstructions in Figures 5–7 of a specular mannequin shown in Fig. 1 and notice the lack of texture on the forehead.

Additionally, by using only two images, binocular Helmholtz stereopsis is faster, simpler, and cheaper to implement than the multinocular Helmholtz stereo technique reported in [21], and so it may be possible to apply this new binocular technique within a much broader range of applications.

Helmholtz stereopsis in general is related to a small set of fairly recent reconstruction techniques (others are [13, 14]) that use both changing viewpoint and illumination to reconstruct surfaces with arbitrary and unknown BRDF. Lu and Little [13] used the term *photogeometric* to describe their technique, which seems like an appropriate term for the entire class of methods. Helmholtz stereopsis differs from these methods, however, in that it both provides direct surface normal estimates and can handle a BRDF that varies over the surface. Conventional photometric stereo can also be applied to surfaces with a non-Lambertian BRDF [10], and it is possible to reconstruct a surface from only two images [17]. However, the reflectance map (i.e., BRDF) must be known *a priori* or be of a known parametric form [11, 19]. On the other hand, conventional stereo has been augmented to handle specularities treating them as an outlier process or using more than two views [3, 5].

In the next section, we show how the Helmholtz constraint arising from a reciprocal pair leads to a differential equation that, given some initial correspondence, can be integrated across corresponding epipolar lines. (See also the recent work of Tu and Mendonca [20].) In Sec. 3, binocular Helmholtz stereo is recast using a functional that does not require initial conditions; this functional is optimized using a multipass, dynamic programming algorithm. Experimental results shown in Sections 2 and 3 demonstrate that accurate reconstruction is possible for non-Lambertian surfaces. Conclusions and future directions are considered in Sec. 4.




Figure 2. The binocular Helmholtz stereo setup. First an image is acquired with the scene illuminated by a single point source as shown on the left. Then, a second image is acquired after the positions of the camera and light source are exchanged as shown on the right.

2 A PDE for Helmholtz Stereopsis

Consider the imaging geometry shown in the left half of Fig. 2 in which a scene is illuminated by an isotropic point source and observed by a perspective camera. Let $\mathbf{o}_l$ and $\mathbf{o}_r$ denote the positions of the camera and light source, respectively. We also denote by $\mathbf{p}$ and $\hat{\mathbf{n}}$ a point on the surface and its associated unit normal vector. The unit vectors $\hat{\mathbf{v}}_l = \frac{1}{|\mathbf{o}_l - \mathbf{p}|}(\mathbf{o}_l - \mathbf{p})$ and $\hat{\mathbf{v}}_r = \frac{1}{|\mathbf{o}_r - \mathbf{p}|}(\mathbf{o}_r - \mathbf{p})$ denote the directions from $\mathbf{p}$ to the camera and light source, respectively. Given this system, the image irradiance at the projection of $\mathbf{p}$ is

$$e_l = f_r(\hat{\mathbf{v}}_r, \hat{\mathbf{v}}_l) \frac{\hat{\mathbf{n}} \cdot \hat{\mathbf{v}}_r}{|\mathbf{o}_r - \mathbf{p}|^2} \quad (1)$$

where $\hat{\mathbf{n}} \cdot \hat{\mathbf{v}}_r$ gives the cosine of the angle between the direction to the light source and the surface normal, $\frac{1}{|\mathbf{o}_r - \mathbf{p}|^2}$ is the $1/r^2$ fall-off from a unit-strength, isotropic point light source, and f_r is the BRDF.

Now, consider the reciprocal case shown on the right of Fig. 2 in which the light source is positioned at $\mathbf{o}_l$, and the camera observes $\mathbf{p}$ from $\mathbf{o}_r$. Because of Helmholtz reciprocity, we have that $f_r(\hat{\mathbf{v}}_r, \hat{\mathbf{v}}_l) = f_r(\hat{\mathbf{v}}_l, \hat{\mathbf{v}}_r)$. This allows us to eliminate the BRDF, and obtain the Helmholtz stereo constraint first introduced in [14]:

$$\left(e_l \frac{\hat{\mathbf{v}}_l(\mathbf{p})}{|\mathbf{o}_l - \mathbf{p}|^2} - e_r \frac{\hat{\mathbf{v}}_r(\mathbf{p})}{|\mathbf{o}_r - \mathbf{p}|^2} \right) \cdot \hat{\mathbf{n}} = 0 \quad (2)$$

Note that Eq. 2 is a first order, nonlinear partial differential equation in the point coordinates $\mathbf{p}$ and their derivatives as expressed through the normal $\hat{\mathbf{n}}$. To solve this, we will first consider an imaging situation in which the PDE is simpler and for which the results are more transparent, and then we will impose the epipolar geometry.

Let the distances of the light source and camera to the scene be large with respect to the relief of the scene, and

let the camera field of view be narrow. Under these conditions, the cameras can be modeled by scaled orthographic projection, and the vectors $\hat{\mathbf{v}}_l(\mathbf{p})$ and $\hat{\mathbf{v}}_r(\mathbf{p})$ can be taken as constant over the scene. As well, the denominators $|\mathbf{o}_l - \mathbf{p}|^2$ and $|\mathbf{o}_r - \mathbf{p}|^2$ can each be taken as constant over the scene. The ratio $\frac{|\mathbf{o}_l - \mathbf{p}|}{|\mathbf{o}_r - \mathbf{p}|}$ can be easily determined when calibrating a Helmholtz stereo rig, and here we take this ratio to be 1. Under these assumptions, Eq. 2 reduces to

$$(e_l \hat{\mathbf{v}}_l - e_r \hat{\mathbf{v}}_r) \cdot \hat{\mathbf{n}} = 0, \quad (3)$$

where $\hat{\mathbf{v}}_l$ and $\hat{\mathbf{v}}_r$ are constants determined during calibration.

We now impose the epipolar constraint to provide a solution to Eq. 3. Without loss of generality, establish a coordinate system for the left camera with a rotation matrix that is the identity, so that $\hat{\mathbf{v}}_l = (0, 0, -1)$. Let the coordinates of points in the world be expressed in this system as (x, y, z) where (x, y) are the image coordinates, and z is the depth. We will consider the surface to be the graph of a C^1 function $z(x, y)$. Furthermore, consider the pair of images to be rectified, so that the second camera's orientation can be expressed as

$$\mathbf{R}_r = [\mathbf{w}_r \mathbf{u}_r \mathbf{v}_r]^T = \begin{bmatrix} w_{r,1} & 0 & w_{r,3} \\ 0 & 1 & 0 \\ v_{r,1} & 0 & v_{r,3} \end{bmatrix}. \quad (4)$$

So, a point (x, y, z) will project to (x, y) in the left image and to $(w_{r,1}x + w_{r,3}z, y)$ in the right image (i.e., the scan lines are epipolar lines). The disparity is then given by $w_{r,1}x + w_{r,3}z(x, y) - x$.

Expressing the depth as $z(x, y)$ and noting that the unnormalized surface normal is $(\frac{\partial z}{\partial x}, \frac{\partial z}{\partial y}, -1)$, we can write the constraint as

$$e_l(x, y) - e_r(x_r, y)(v_{r,1} \frac{\partial z}{\partial x} - v_{r,3}) = 0 \quad (5)$$

where $x_r = w_{r,1}x + w_{r,3}z(x, y)$, and this holds for all y . Rewriting this, we have

$$\frac{\partial z}{\partial x} = - \frac{e_l(x, y) + e_r(x_r, y)v_{r,3}}{v_{r,1}e_r(x_r, y)} \quad (6)$$

This can be numerically integrated as

$$z(x, y) = \int_{x_0}^x \frac{\partial z}{\partial x} dx + z(x_0, y). \quad (7)$$

In other words, for each epipolar line y , this integral can be independently evaluated to provide the depth across the epipolar line. Note that there is no search for correspondence over some disparity space, as correspondence is determined as a byproduct of integration. The only open issue is: "For each epipolar line y , what is the boundary condition $z(x_0, y)$?" There are two ways to look at this issue.




Figure 3. A family of reconstructions for one epipolar line of the specular cylinder shown in the bottom row of Fig. 4. The family arises from different initial conditions $z(x_0, y)$ when integrating Eq. 7. The thick (red) curve is the member of this family with the correct geometry, and is redrawn with a different scaling in the lower right of Fig. 4.

On the one hand, knowing $z(x_0, y)$ for some (x_0, y) on each epipolar line amounts to having the means to determine the depth or establish correspondence for one point on each epipolar line. Alternatively, one can view Eq. 7 as defining a one-parameter family of reconstructed curves along each epipolar line; elements of the family are indexed by different depth values at (x_0, y) . In Sec. 3, we will introduce a method for selecting a member of this family for each epipolar line.

We have implemented this method in Matlab using Runge-Kutta integration and empirically validated its effectiveness in the following experiment. We gathered reciprocal pairs of images of three cylinders made of a Lambertian material, a rough non-Lambertian material [18], and a specular plastic material. Images were acquired with a Kodak DCS 760 digital camera, and the scene was illuminated with a 150W halogen bulb. The camera system was geometrically calibrated, and the distance from the camera to the object was about two meters which satisfies the approximation needed in deriving Eq. 3. Figure 4 shows for each cylinder a pair of rectified images and a plot of the image intensity across a pair of central epipolar lines. Note that these curves are characteristic of these three material types.

For the epipolar line, a family of reconstructions can be obtained for a discrete set of initial depths (disparities), and Fig. 3 shows such a family for the specular cylinder. Since $z(x, y)$ must be C^1 , the integration cannot cross the occluding contour where there is a depth (zeroth-order) discontinuity. Within this family lies the correct reconstruction, which could be selected from a single correspondence or

some other means (e.g., smoothness, a shape prior, the dynamic programming method introduced in Sec. 3, etc.). The last column of Fig. 4 shows the reconstructed depth across one epipolar line overlaid on a circular cross section. In this experiment, the initial conditions were chosen manually. The RMS errors between the reconstructed curve and overlaid circle as a percentage of the cylinder's radius are 0.11%, 1.7%, and 0.94% respectively for the Lambertian, generalized Lambertian, and specular cylinders. Note that the reconstructed curve for the Lambertian cylinder is indistinguishable from the ground truth circle whereas there is a slight deviation for the specular cylinder.

3 Surface Reconstruction

As discussed in the previous section, we can solve the binocular Helmholtz PDE by integrating along corresponding epipolar lines. The recovery of the correct solution, however, requires a correct initial correspondence for starting the integration. This requirement may seem like a difficult obstacle for the method; yet, we show in this section that this can be easily overcome by applying a matching constraint common to traditional stereo algorithms.

Recall that traditional dense stereo algorithms must determine the correspondence for all points along corresponding epipolar lines in the left and right images. In contrast, binocular Helmholtz stereo needs only – in theory at least – to determine the correspondence of a single pair of points. This correspondence provides an initial condition which can then be integrated to the left and right along the epipolar line to determine all correspondences and establish depth for all points.

One could develop an algorithm that first identified the most prominent/salient feature (e.g., edges) for each pair of epipolar lines and then assigned correspondence between these features as way of creating anchor points for the integration described above. Even for surfaces with arbitrary BRDF, there will generically exist image intensity discontinuities (edges) corresponding to discontinuities in the albedo or BRDF across the surface, discontinuities in the surface normal (creases), and discontinuities in depth at occlusion boundaries [4]. Consider more closely the two images of the mannequin in Fig. 1. While the specular highlights clearly indicate that this surface is far from Lambertian, there are also BRDF discontinuities arising from the eyes, eyebrows, and lips which can be used as features.

While coupling a feature-based stereo algorithm with the integration method of Sec. 2 might work, it would rely on the identification of feature points and would break down ungracefully when the initial correspondences were incorrectly assigned. Instead, we set up the problem of depth recovery along epipolar lines as the minimization of a functional that includes a term for matching image features and a term for keeping the solution close to that dictated by the

integration. The functional is chosen to allow for efficient, global minimization using dynamic programming.

Pass 1: Along Epipolar Lines

As before, we assume that we have a reciprocal pair of orthographic images taken of a C^1 surface (a graph $z(x, y)$), that the camera/source positions are far from the surface, and that the images are rectified. In Sec. 2, a differential equation was derived for $\frac{\partial z}{\partial x}$ as a function of intensities at corresponding points along an epipolar line, here we modify the equation so that it is defined relative to a cyclopean coordinate system as in [1].

Let $r(x, z)$ denote the following ratio of image measurements

$$r(x, z) = -\cot(\theta) \frac{e_l(x_l) - e_r(x_r)}{e_l(x_l) + e_r(x_r)} \quad (8)$$

where x is the cyclopean coordinate; θ is the half angle between viewing direction; $x_l = x \cos(\theta) + z \sin(\theta)$ and $x_r = x \cos(\theta) - z \sin(\theta)$ are coordinates on the left and right epipolar lines, respectively. For corresponding points, we have that

$$\frac{\partial z}{\partial x} = r(x, z). \quad (9)$$

The functional that we minimize has as its matching term the squared magnitude of the difference between the gradient in the left and right images. This matching term is coupled with a term measuring the squared distance between $r(x, z)$ and the partial derivative of the hypothesized solution z with respect to x . The functional is given as

$$E_h(z(x)) = \int \left(\frac{\partial z}{\partial x} - r \right)^2 dx + \alpha \int |\nabla e_l - \nabla e_r|^2 dx \quad (10)$$

where $e_l = e_l(x_l(x, z))$, $e_r = e_r(x_r(x, z))$, and α is a weighting term. For each epipolar line, we use dynamic programming to find a discrete approximation to $z(x)$ that minimizes $E_h(z(x))$. For n discrete values of x , we consider m possible depth values, and the computational cost of finding the global minimum of $E_h(z(x))$ is $O(nm^2)$; see [2] for details.

Pass 2: Across Epipolar Lines

If each epipolar line had an image feature such as an albedo edge, the gradient matching term would lock onto the edge and effectively provide the initial correspondence needed for the integration. We are not guaranteed, however, that each epipolar line will have an image feature that will allow the minimization to find the correct initial correspondence and, thus, the correct solution. Reconsider the images of the mannequin in Fig. 1. While there are strong image features for epipolar lines crossing the eyebrows, eyes and mouth regions, the forehead is smooth and textureless, whereas the nose has significant relief and little texture. However, we




Figure 4. Reconstruction of three real cylinders of three material types: The cylinder in Row 1 is approximately Lambertian, the rough cylinder in Row 2 has Oren-Nayar generalized Lambertian reflectance [18], and the plastic cylinder in Row 3 is highly specular. The first two columns show a rectified pair of images of the cylinder. The third column shows a plot of the image intensities across one epipolar line in the left (blue, solid) and right (red, dashed) images. The fourth column shows the reconstructed shape (thick, blue) across the epipolar line superimposed on a circular cross section.

expect that in general other epipolar lines will have image features, and we can use these to determine the best solutions for epipolar lines that do not. To operationalize this, we define and minimize a second stage functional across epipolar lines (in a manner different than that used in [16]).

The idea is to compute as the output from Pass 1 a family of n solutions minimizing Eq. 10 such that the endpoints of the solutions vary over a range of possible z values. If the family of solutions is big enough (i.e., if our sampling of the range of z values is fine enough), then the correct solution should be well represented by one member from the family of solutions. Note that this should hold whether or not the epipolar line has an image feature. To choose the correct solution for each epipolar line, we simply choose the collection of solutions (one for each line) that produces the most coherent surface according to a simple smoothness criteria imposed across epipolar lines.

More precisely, let $x = x_e$ denote the end of the epipolar line. Let $z(x_e, y) = z_e$ denote $z(x, y)$ at the endpoint $x =$

x_e for scanline y . For each y and for each ending point z_e in the range of possible z values, we compute a solution

$$\hat{z}(x, y | z(x_e, y) = z_e) = \arg \min_{(z(x, y) | z(x_e, y) = z_e)} E_h(z(x, y)). \quad (11)$$

In other words, $\hat{z}(x, y | z(x_e, y) = z_e)$ is the solution along epipolar line y that minimizes Eq. 10 subject to the constraint that $z(x_e) = z_e$. Thus for each y , the family of solutions is indexed by the value of the ending point z_e . Note that this family differs from the one arising in Sec. 2 and shown in Fig. 3. And, within this family there should be a z_e and a corresponding solution $\hat{z}(x, y | z(x_e, y) = z_e)$ that is close to the correct solution. We denote the family of solutions over all epipolar lines by $\hat{Z} = \{\hat{z}(x, y | z(x_e, y) = z_e), \forall z_e, \forall y\}$.

Nevertheless, we still have the problem of determining the correct solution from $\hat{Z}$ for each epipolar line y . Con-




Figure 5. Two views of a reconstructed plastic mannequin created using binocular Helmholtz stereopsis.

sider the following functional defined over the entire image

$$E_v(z(x, y)) = \iint \left(\frac{\partial z}{\partial y} \right)^2 dx dy. \quad (12)$$

We can find the correct choice from $\hat{Z}$ for each epipolar line y by minimizing the above functional. We do not consider all possible solutions $z(x, y)$, rather we limit the search of solutions to the family of solutions $\hat{Z}$ given by Pass 1. Thus, the optimization in Eq. 12 is effectively being done over the endpoints z_e for each y . We take as our solution to Pass 2,

$$\hat{z}(x, y) = \arg \min_{\hat{Z}} E_v(z(x, y)). \quad (13)$$

As in Pass 1, we use dynamic programming to find a discrete approximation to the $\hat{z}(x, y)$ that minimizes $E_v(z(x, y))$. The computational cost of this dynamic programming step is $O(m^2l)$ where m is the number of endpoints (depth values), and l is the number of epipolar lines.

Note that this formulation has an inherent asymmetry, as the second pass considers a range of ending points and not a range of starting points. We correct this by re-running this two stage process in reverse. Specifically, we run Pass 1 and Pass 2 across the data to find a collection of optimal ending points $\hat{z}_e = \hat{z}(x_e, y)$ for each epipolar line. We then re-run Pass 1 in reverse (i.e., from right to left), fixing the endpoint such that $z(x_e, y) = \hat{z}_e$ for each y . At this stage, for each y we now have a family of solutions indexed by the value of the beginning point z_b . The overall solution is then chosen by re-running Pass 2 to select the optimal starting points.

We should point out that this algorithm has only one real parameter: the weighting α of the image gradient term used in optimization of Eq. 10. The optimization in Eq. 12 is parameter free. This optimization does not smooth the solution along the epipolar lines, rather it chooses the solutions which together form the surface that is smoothest across the epipolar lines.

Here, we present results on a reciprocal pair of images of a mannequin head shown in Fig. 1. In Fig. 5, we display two

views of a mesh of the surface reconstructed using our two pass dynamic programming method. For this reconstruction, the brightness in the left and right images was normalized by the maximum brightness in both images, and the value of $\alpha = 0.1$. In Fig. 6, we display a depth map image in which light corresponds to near and dark to far. Finally, in Fig. 7, we display a single view of the mesh with the left image texture mapped onto it. Notice that the method is unhampered by the specularities and is able to both “lock onto” the features such the eyes, eyebrows, and lips, but also provide good reconstructions in textureless regions such as the forehead.

4 Discussion

This paper introduces binocular Helmholtz stereopsis and shows that the constraint arising from reciprocal images is a partial differential equation which can be readily solved. The implementation validates that the method can be applied to objects with arbitrary and unknown BRDF, and unlike conventional stereo it is able to accurately reconstruct shape in textureless regions.

There are still many ways to improve and enhance binocular Helmholtz stereopsis. First, it is straightforward to extend the implementation to perspective projection and nearby light sources, though the PDE becomes more complex. Second, it is assumed throughout that the depth function $z(x, y)$ is C^1 , yet it should be possible to permit depth discontinuities. As described in [14, 21], the equivalent to half-occluded regions corresponding to depth discontinuities in conventional stereo are shadowed regions in reciprocal image pairs. Hence, it should be possible to augment Eq. 10 to introduce depth discontinuities as in [1, 6, 8], using shadows as a cue for half occlusion.

Finally, though an advance brought forth in this paper is to reduce Helmholtz stereopsis from requiring at least three camera positions [14, 21] to binocular imaging, it would be worthwhile to re-examine multinocular Helmholtz stereopsis by directly considering the set of differential equations arising from each reciprocal pair. It may then be possible to



Figure 6. A depth map from the cyclopean viewpoint of the mannequin face in which light corresponds to near and dark to far.



Figure 7. A texture mapped reconstruction of the mannequin face, rendered with a Phong reflectance model.

eliminate the smoothness constraint between epipolar lines used in the second pass of Sec. 3. Perhaps, multinocular Helmholtz reconstruction can be formulated to exploit advantages found through photoconsistency and space carving [12] or level set methods [7].

Acknowledgements

P.N. Belhumeur and T. Zickler were supported by the National Science Foundation under grants PECASE IIS-9703134, ITR IIS-00-85864, EIA-02-24431, IIS-03-08185, and KDI-99-80058. D.J. Kriegman and J. Ho were supported by the NSF under grants EIA-00-04056, CCR-00-86094 and IIS-03-08185. J. Ponce was supported by the UIUC Campus Research Board and the NSF under grant IIS-03-08087.

References

- [1] P. Belhumeur. A Bayesian approach to binocular stereopsis. *Int. Journal of Computer Vision*, 19, 1996.
- [2] R. Bellman. *Dynamic Programming*. Princeton University Press, 1957.
- [3] D. Bhat and S. Nayar. Stereo and specular reflection. *Int. J. Computer Vision*, 26(2):91–106, 1998.
- [4] T. Binford. Generic surface interpretation: Observability model. In *Proc. of the 4th International Symposium on Robotics Research*, Santa Cruz, CA, August 1987.
- [5] A. Blake and G. Brelstaff. Specular stereo. In *Proc. Int. Joint Conf. on Art. Intell.*, pages 973–976, 1985.
- [6] I. J. Cox, S. Hingorani, B. M. Maggs, and S. B. Rao. Stereo without disparity gradient smoothing: a Bayesian sensor fusion solution. In D. Hogg and R. Boyle, editors, *British Machine Vision Conference*, pages 337–346. Springer-Verlag, 1992.
- [7] O. Faugeras and R. Keriven. Complete dense stereovision using level set methods. In *Proc. European Conf. on Computer Vision*, 1998.
- [8] D. Geiger, B. Ladendorf, and A. Yuille. Occlusions in binocular stereo. In *Proc. European Conf. on Computer Vision*, Santa Margherita Ligure, Italy, 1992.
- [9] H. v. Helmholtz. *Treatise on Physiological Optics*, volume 1. Dover (New York), 1925.
- [10] B. Horn. *Robot Vision*. MIT Press, Cambridge, Mass., 1986.
- [11] K. Ikeuchi and B. Horn. Numerical shape from shading and occluding boundaries. *Artificial Intelligence*, 17:141–184, 1981.
- [12] K. Kutulakos and S. Seitz. A theory of shape by space carving. *Int. J. Computer Vision*, 38(3):197–216, July 2000.
- [13] J. Lu and J. Little. Reflectance and shape from images using a collinear light source. *Int. Journal of Computer Vision*, 32(3):1–28, 1999.
- [14] S. Magda, T. Zickler, D. Kriegman, and P. Belhumeur. Beyond Lambert: Reconstructing surfaces with arbitrary BRDFs. In *Proc. Int. Conf. on Computer Vision*, pages II: 391–398, 2001.
- [15] F. Nicodemus, J. Richmond, J. Hsia, I. Ginsberg, and T. Limperis. Geometric considerations and nomenclature for reflectance. Monograph 160, National Bureau of Standards (US), 1977.
- [16] Y. Ohta and T. Kanade. Stereo by intra- and inter-scanline search. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 7(2):139–154, 1985.
- [17] R. Onn and A. Bruckstein. Integrability disambiguates surface recovery in two-image photometric stereo. *Int. J. Computer Vision*, 5(1):105–113, 1990.
- [18] M. Oren and S. Nayar. Generalization of the Lambertian model and implications for machine vision. *Int. J. Computer Vision*, 14:227–251, 1996.
- [19] H. Tagare and R. deFigueiredo. A theory of photometric stereo for a class of diffuse non-lambertian surfaces. *IEEE Trans. Pattern Anal. Mach. Intelligence*, 13(2):133–152, February 1991.
- [20] P. Tu and P. Medonca. Surface reconstruction via Helmholtz reciprocity with a single image pair. In *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, June 2003.
- [21] T. Zickler, P. Belhumeur, and D. Kriegman. Helmholtz stereopsis: Exploiting reciprocity for surface reconstruction. In *Proc. European Conf. on Computer Vision*, pages III: 869–884, May 2002.