

# Reflectance Function Estimation and Shape Recovery from Image Sequence of a Rotating Object

Jiping Lu                      Jim Little  
Laboratory for Computational Intelligence  
Department of Computer Science  
The University of British Columbia  
2366 Main Mall, Vancouver B.C., Canada V6T 1Z4

## Abstract

In this paper we describe a technique for surface recovery of a rotating object illuminated under a collinear light source (where the light source lies on or near the optical axis). We show that the surface reflectance function can be directly estimated from the image sequence without any assumption on the reflectance property of the object surface. From the image sequence, the 3D locations of some singular surface points are calculated and their brightness values are extracted for the estimation of the reflectance function. We also show that the surface can be recovered by using shading information in two images of the rotating object. Iteratively using the first-order Taylor series approximation and the estimated reflectance function, the depth and orientation of the surface can be recovered simultaneously. The experimental results on real image sequences of both matte and specular surfaces demonstrate that the technique is feasible and robust.

## 1 Introduction

Shading in images can be used for surface recovery, for example, in shape from shading [2, 7], photometric stereo [13, 14, 15], and photometric sampling [4]. In order to use shading information, the reflectance function of the surface under recovery must be known. The most commonly used assumption is Lambertian reflectance [11, 7] because of its simplicity. In photometric stereo [14, 15], the reflectance function is computed from a calibration object of known shape whose surface is made of the same material as the surface of the object under consideration.

However, for most real objects, the surface reflectance is not Lambertian. The Lambertian assumption is only valid in some limited cases and limited lighting and viewing conditions [12, 10, 6]. Empirical photometric stereo [14, 15] requires that the calibration object and the object under recovery have the same reflectance function and be illuminated and viewed under the same conditions. The calibration process may become difficult or impossible when we cannot find a suitable calibration object.

We compute the surface reflectance directly from an image sequence of a rotating object and then use the surface reflectance function to recover the orientation

and the scaled surface depth. The rotating object is on a turntable whose rotation angle can be controlled. A *collinear* light source, which points in the same direction as the camera viewing direction, and lies on or very near the optical axis, gives a uniform radiance over the object. The images are taken when the object rotates on the turntable. Under the illumination of a collinear light source, the image brightness of a surface point is a function only of the incident angle  $i$  which is the angle between the illuminant direction and the surface normal.

The estimate of the reflectance function is based on a set of singular surface points whose surface normals are in the viewing direction. The 3D coordinates of the singular points are computed from image points of the maximum brightness values and the corresponding contour points in the image sequence. Given the 3D locations of these singular points, their brightness values in the image sequence are extracted. The brightness values and the corresponding incident angles during the rotation are used to build the reflectance function.

Two subprocedures are used in surface recovery. The first subprocedure computes the depth around an image point of known depth and surface orientation by using first-order Taylor series approximation. The second procedure finds the orientation of a surface point of known depth value. This subprocedure computes the projections of the surface point on two different images and calculates the orientation of the surface point from the brightness values in the two images. Starting from a surface point of known depth and orientation and iteratively applying the first and second subprocedure, the computation on the depth and surface orientation can be expanded over the object surface. The calculation for surface orientation can be done by look-up table just as in photometric stereo [14, 15].

Section 2 introduces the experimental setting for our work. Section 3 presents a method for building the reflectance function. Section 4 describes surface recovery process. Section 5 shows some experimental results on real image sequences. The final section discusses experimental results and future work. The technical report [3] provides an extended discussion of this material.




Figure 1: Experimental setup

## 2 Experimental Conditions

The imaging geometry is shown in Fig. 1. The object is on a turntable whose rotation angle can be controlled. The Y axis coincides with the rotation axis of the turntable. Both the camera viewing direction and the light illuminant direction are aligned with the Z axis. The light source is a distant light source with uniform radiance over time and illuminated area. Since the camera is far away from the object, orthographical projection is used. To obtain the projection of the rotation axis in the images, a vertical black line on a white board is aligned with the rotation axis and then identified from the image of the board. Images are taken when the object rotates around the Y axis in the direction from the X axis to the Z axis.

The surface of the object is assumed to be piecewise continuous and differentiable. The surface orientation is defined as  $(p, q, 1)$  with  $p = \partial z(x, y)/\partial x$  and  $q = \partial z(x, y)/\partial y$ . When the object rotates, the coordinates and the orientation of the surface points on the object change. Let  $(x, y, z)$  be a 3D surface point on the object and  $(p, q, 1)$  be the surface orientation of this point. After an  $\alpha$  degree rotation, the 3D location of this point  $(x_\alpha, y_\alpha, z_\alpha) = (x \cos \alpha - z \sin \alpha, y, x \sin \alpha + z \cos \alpha)$  and the surface orientation  $(p_\alpha, q_\alpha, 1)$  of this point is determined by

$$p_\alpha = \frac{p \cos \alpha + \sin \alpha}{\cos \alpha - p \sin \alpha}, \quad (1)$$

$$q_\alpha = \frac{q}{\cos \alpha - p \sin \alpha}. \quad (2)$$

We also assume the reflectance of the object surface is uniform. In the general case, the image brightness of a 3D point under a distant light source is determined by the reflectance function  $R(i, e, g)$  [13], where the incident angle  $i$  is the angle between the incident ray and the surface normal, the emergent angle  $e$  is the angle between the emergent ray and the surface normal, and the phase angle  $g$  is the angle between the incident and emergent rays. Under a collinear light source, as shown in Fig. 2, the phase angle  $g$  becomes zero and the incident angle  $i$  becomes the same as the emergent angle  $e$ . In this case, all the components of




Figure 2: Under collinear light, the image brightness only depends on angle  $i$

the reflectance such as the specular component, diffuse component and other components [10] are functions of the incident angle  $i$  only. Thus for the surface point  $(x, y, z)$ , its image brightness value can be written as

$$E(x, y) = R(i(x, y)) \quad (3)$$

where  $i(x, y)$  is the incident angle at point  $(x, y, z)$ . We assume the function is strictly monotonic and its inverse exists. This is true for most surfaces.

## 3 Estimating the Reflectance Function

The estimate of the reflectance function is based on singular surface points whose surface orientations are the same as the viewing direction. These singular points also have singular brightness values in the image because the incident angles of these surface points is zero. Let point  $(x_s, y_s, z_s)$  be one of these points. The values of  $x_s$  and  $y_s$  can be directly found from the image by searching for a point of local maximum brightness. To determine  $z_s$ , we look at the image taken after the object has rotated by 90 degrees. In that image, the singular point is projected on an occluding contour point. It is easy to find the location of the contour point since its y coordinate is already known. The value  $z_s$  of the singular point is the horizontal distance from the contour point to the image of the rotation axis. Considering generic surfaces, we assume that some singular points whose images are available in the first image and will not be occluded during 90 degree rotation.

Given the depth values of these singular points in the first image, we can extract their brightness values from the image sequence. For the singular point  $(x_s, y_s, z_s)$ , after the object has rotated by an angle  $\theta_i \leq 90^\circ$ , the 3D location  $(x_i, y_i, z_i)$  of the singular point is  $(x_s \cos \theta_i - z_s \sin \theta_i, y_s, x_s \sin \theta_i + z_s \cos \theta_i)$ . The image brightness  $E(x_i, y_i)$  of this point can be directly obtained from image point  $(x_i, y_i)$ . The incident angle for the singular point is the same as the object rotation angle  $\theta_i$ . From the brightness  $E(x_i, y_i)$  and the corresponding incident angle  $\theta_i$ , we can build the reflectance function for the surface.

The dark line in Fig. 4 shows the reflectance function  $E = R(i)$  obtained from an image sequence of a

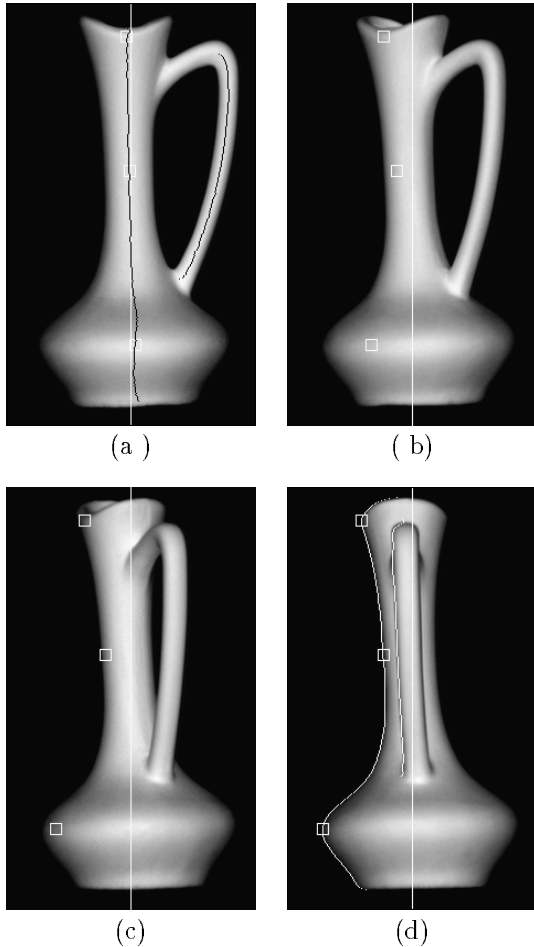


Figure 3: Tracking the singular points over the image sequence

vase. Nineteen images of the vase are taken during a 90 degree rotation. The rotation angle between two consecutive images in the image sequence is 5 degrees. The images in Fig. 3 are four images of the rotating vase. The images (a), (b), (c) and (d) are the images taken after rotation by 0, 30, 60, and 90 degrees, respectively. The white line in the middle of each image is the virtual image of the rotation axis of the object. Three singular points are extracted from the image sequence. The center of the small square box in each image denotes the image of a singular surface point. The average of the brightness values of the three surface points is used to build the reflectance function. The estimated function are shown in Fig. 4. Only the three singular points on the body of the vase are tracked and used to estimate the reflectance function. The singular points on the handle are not used for the estimation because interreflection make the brightness values of these points much bigger than these points on the body part. However, the 3D locations of these

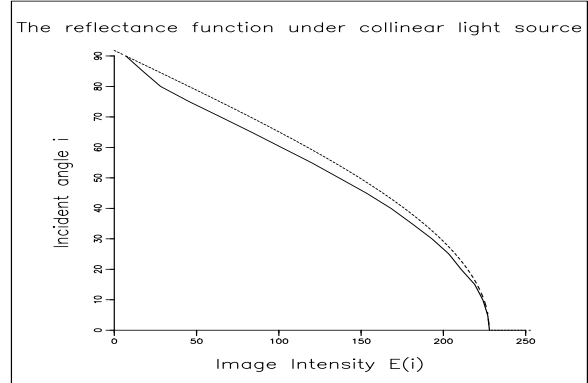


Figure 4: The estimated reflectance function of the vase

singular points can be easily found and can be used in surface recovery.

The vase is made of clay and the surface of the vase is considered a matte surface. The dotted line in Fig. 4 shows the reflectance function of an ideal Lambertian surface with the same maximum brightness value as that of the estimated function. Although the surface of the vase is considered a matte surface its surface reflection is not exactly Lambertian.

#### 4 Surface Recovery

Surface recovery can be done by using any two images in the image sequence of the rotating object after the reflectance function has been obtained. Without losing generality, we use the first image and the image taken after the object has rotated by a certain angle. The depth and surface orientation are computed at every point in the first image. Surface recovery uses two subprocedures to compute the depth and surface orientation. The first subprocedure does local expansion of depth using surface orientation by first-order Taylor series approximation. For an image point  $(x, y)$ , if the depth  $z$  and the surface orientation  $(p, q, 1)$  are known, the depth  $z'$  of an image point  $(x + \delta x, y + \delta y)$  in the small neighborhood of the image point  $(x, y)$  is calculated by  $z' = z + \delta z = z + p\delta x + q\delta y$ . The second subprocedure computes the projections of a surface point of known depth value on the two images and determines surface orientation (up to a two-way ambiguity) from image brightness values. The following calculation shows how we derive orientation from image brightness.

Let  $image_0$  be the image taken before the rotation and  $image_1$  be the image taken after  $\alpha$  degree rotation of the object. Let  $(x_0, y_0)$  and  $(x_1, y_1)$  be the projections of a 3D surface point in  $image_0$  and  $image_1$  respectively and their brightness values are  $E(x_0, y_0)$  and  $E(x_1, y_1)$ . Using the inverse reflectance function  $i = R^{-1}(E)$ , we obtain the incident angle  $i_0$  and  $i_1$  from  $E(x_0, y_0)$  and  $E(x_1, y_1)$ . Let the surface orientation of the 3D point be  $(p_0, q_0, 1)$  when  $image_0$  is taken and the surface orientation of the same 3D point be  $(p_1, q_1, 1)$  when  $image_1$  is taken. From the definition of incident angle and the transformation between the

object coordinates, we have

$$\cos i_0 = \frac{1}{\sqrt{p_0^2 + q_0^2 + 1}}, \quad (4)$$

$$\cos i_1 = \frac{1}{\sqrt{p_1^2 + q_1^2 + 1}}, \quad (5)$$

$$p_1 = \frac{p_0 \cos \alpha + \sin \alpha}{\cos \alpha - p_0 \sin \alpha}, \quad (6)$$

and

$$q_1 = \frac{q_0}{\cos \alpha - p_0 \sin \alpha}. \quad (7)$$

Substituting  $p_1$  and  $q_1$  in Equation 5, we get

$$\cos i_1 = \frac{1}{\sqrt{1 + \left(\frac{p_0 \cos \alpha + \sin \alpha}{\cos \alpha - p_0 \sin \alpha}\right)^2 + \left(\frac{q_0}{\cos \alpha - p_0 \sin \alpha}\right)^2}}; \quad (8)$$

The equation can be simplified to

$$\cos i_1 = \cos i_0 (\cos \alpha - p_0 \sin \alpha). \quad (9)$$

Solving Equation 9 for  $p_0$  and from Equation 4, we get

$$p_0 = \frac{1}{\tan \alpha} - \frac{\cos i_1}{\cos i_0 \sin \alpha}, \quad (10)$$

$$q_0 = \pm \sqrt{\frac{1}{\cos^2 i_0} - p_0^2 - 1}. \quad (11)$$

The surface recovery procedure starts at the image points whose depth and surface orientations are known. These starting points could be the singular image points we used to compute the reflectance function. For each starting point, we use the first subprocedure to compute the depth for the image points in its neighborhood. For each neighbor point, we used the second procedure to compute its projection on the other image from its depth computed by the first procedure and determine its surface orientation from its two brightness values in the two images. Applying the first subprocedure on the neighbor points, we expand the depth over a larger area. Applying the second subprocedure on the larger area, we compute surface orientation for the large area. By iterating the first and the second subprocedures, we spread the computation over the whole image to obtain the depth and surface orientation at the same time. The number of local operations in this process is linear in the number of pixels in the image. The surface recovery procedure can be described as follows.

### Surface recovery procedure

1. Take an image point of known depth and surface orientation in  $image_0$  as an input point.
2. Calculate the depth values for the image points in the small neighborhood of the input point with first-order Taylor approximation.

3. For each neighbor point, calculate its projection on  $image_1$  by using its depth value computed in the first step.
4. For each neighbor point, estimate its incident angle  $i_0$  and  $i_1$  from its brightness values in  $image_0$  and  $image_1$  by using the inverse of the reflectance function; then determine its surface orientation by using Equation 10-11.
5. Take each neighbor point as the input point and repeat steps 2-4 until the depth and orientation has been computed at every point in  $image_0$ .

When we compute a new depth value  $z'$  in the  $y$  direction, we have  $z' = z + q\delta y$ . Thus an ambiguity in depth occurs as there are two solutions of  $q$  obtained from Equation 11. There is no ambiguity when we expand the computation in the  $x$  direction as the solution for  $p$  is unique. The ambiguity in the depth caused by the ambiguity in  $q$  can be removed by the continuity constraint and some points of known depth values. We first select a non-horizontal curve of known depth values at some points in the first image. We compute the depth and surface orientation by iterating the first and the second subprocedures, and remove the depth ambiguity on the curve. Then we take points on the curve as starting points and compute depth in the  $x$  direction. We divide the curve into several paths with each path starting and ending at the points of known depth values. We assume that any value of  $q$  along a path is not zero. Two depth values are computed at each point and two sets of depth values will be found on a path. It is clear that only one of the two sets is the correct set and at the end of a path, only the depth value from the correct set matches with the known depth. By this fact and the continuity constraint, the correct depth values along a path can be determined. The computed depth values will not exactly match with the known depth value at the end point of a path, because of image noise and some other errors, but they are close enough so that the ambiguity can be removed.

Since the locations and the depth values of singular points can be easily found from images, we usually select a curve which passes through several singular points for removing the ambiguity. In our experiment we select a curve  $p = 0$  because the depth values on this curve correspond to an occluding boundary in the image taken after the object is rotated by 90 degrees. This boundary gives us the true depth values on the curve so that we can check the accuracy of the computed depth values. Fig. 5(a) and (b) are two images of a vase used for surface recovery. The image in (a) is taken before rotation, and the image in (b) is taken after a 10 degree rotation. The white curve on the body part of the vase in (a) is the curve of  $p = 0$ . The curve should be smooth and continuous but it is not because of image noise, non-uniform albedo and other facts. The straight line in the middle of Fig. 5(c)-(f) denotes the line of depth=0. The horizontal distance from a point to the line is the depth value. The center of the small boxes in the diagrams represents the true depth value used for removing ambiguity. These true




Figure 5: Computing depth on a curve and removing the ambiguity in the depth

depth values are obtained from the occluding boundary of the image taken after the object has rotated by 90 degrees. The two sets of depth values computed for the curve are shown in (c). The correct set of depth values, after removal of the ambiguity, is shown in (d). Since the depth computed at the end point of a path will not be exactly the same as the known depth at that end point, the depth computed on a curve which consists of several paths will not be continuous (see Fig. 5(d)). The discontinuity on the depth in (d) is mainly caused by non-uniform albedo on the object surface. The discontinuity occurs at a point which is the starting point for a path and ending point for another path or the ending points for two paths.

A *distance-weighted averaging* method is used to make the depth value continuous on the curve. The idea is to make the depth value at the two ends of a path equal to the known depth values. For a path which connects two points  $(x_0, y_0)$  and  $(x_k, y_k)$  with  $y_0 \neq y_k$ , the computation of depth and orientation along the path can start from either of the two points. Thus two sets of depth values along a path can be obtained (see Fig. 5(e)). Let  $Z_i$ , ( $i = (0, k)$ ) be a set of depth values computed along the path starting from  $(x_i, y_i)$  and ending at  $(x_{k-i}, y_{k-i})$ . Let the known depth value at the point  $(x_i, y_i)$  be  $z_i$ . Let point  $(x_j, y_j)$ ,  $0 \leq j \leq k$ , be a point on the path and let  $z_{i,j}$  be a depth value in  $Z_i$  and computed at point  $(x_j, y_j)$ . Since  $(x_0, y_0)$  is the starting point for  $Z_0$  and

$(x_k, y_k)$  is the starting point for  $Z_k$ , we have  $z_{0,0} = z_0$  and  $z_{k,k} = z_k$ . In the general case,  $z_{0,k} \neq z_k$  and  $z_{k,0} \neq z_0$ . A new depth value  $\tilde{z}_j$  is calculated at point  $(x_j, y_j)$  by

$$\tilde{z}_j = \frac{(k-j)z_{0,j}}{k} + \frac{jz_{k,j}}{k} \quad (12)$$

It is easy to show that  $\tilde{z}_0 = z_{0,0} = z_0$  and  $\tilde{z}_k = z_{k,k} = z_k$ . After recomputing depth for some paths which cause discontinuity, we get continuous depth values on the curve. Furthermore, if the depth along a path before the averaging is smooth, the depth along the path after the averaging is also smooth. The recomputed depth values are consistent with the original depth values at several known 3D points, and the resulting depth values are close to the real depth values. In Fig. 5(f), the true depth and the recomputed continuous depth are represented in one image. The two sets of depth values are quite close.

## 5 Experimental Results

In experiments, we use a calibrated image facility (CIF) [15] in our lab to control the rotation of the object and the imaging condition. We use a DC powered beamed light source and mount it on the top of the camera. The light source and the camera point in the same direction at the object on a turntable. In practice, the radiance of the light source is not constant over the illuminated area. We use a uniform white board to calibrate the non-uniform illumination. Since the distance from the camera to the object is far bigger than the size of the object, the camera is set to telephoto mode and orthographic projection is used. The actions of rotating an object and taking images are well synchronized by a computer. Nineteen images of a vase are taken with each successive image taken after a successive 5 degree rotation. The total rotation for the image sequence is 90 degrees.

The images are smoothed with a Gaussian filter of  $\sigma = 1$  to filter image noise and quantization effects. Four images from the image sequence of the vase are shown in Fig. 3. The estimated reflectance function of the vase is shown in Fig. 4 as a dark line. In surface recovery, the first image (Fig. 5(a)) and the image taken after a 10 degree rotation (Fig. 5(b)) are used. We first compute depth and surface orientation along a curve of  $p = 0$  (see Fig. 5(a)). The depth ambiguity on the curve caused by the ambiguity in  $q$  is removed by the method described in the previous section. The continuous depth values are obtained by the distance-weighted averaging. In Fig. 5(f), it is overlapped with the true depth value on the curve. Starting from points on the curve, we expand the computation on the depth and surface orientation in the  $x$  direction by  $z' = z + p\delta x$  until we reach the background. This process may not reach some areas, such as handle part, in the image. Then we expand the computation in the  $y$  direction to reach the unrecovered areas. For each of these areas, we repeat the process of computing depth along the curve of  $p = 0$ , removing the depth ambiguity on the curve and expanding the computation in the  $x$  direction.




Figure 6: The height plot of the recovered surface




Figure 7: The image sequence of a cup

The surface plot of the final recovered depth of the vase is shown in Fig. 6. Surface plots are displayed with Matlab by using the depth values calculated on the first image. We did not do any smoothing or regularization on the depth data.

We also experimented with a porcelain cup. The reflectance of the cup presents a strong specular reflection. There are peak brightness values at singular points on the surface of the cup. Fig. 7 contains four images of the cup. Image (a) to (d) are the images of the cup taken after the 0, 30, 60, 90 degrees rotation. Three singular points are tracked to get the surface reflectance function. In Fig. 8, the dark line denotes the reflectance function of the cup, and the dotted line is the Lambertian reflection with the same maximum brightness value as that of the estimated reflectance function. The difference between the two reflectance




Figure 8: The surface reflectance function of the cup




Figure 9: The height plot of the recovered cup

functions is obvious. The final recovered surface of the cup is shown in Fig. 9.

## 6 Discussion and Future Work

The experimental results are promising. Since we don't assume the reflectance function has any particular form, the reflectance estimation method in this paper can be applied to any surface of isotropic reflection. Though we did not do experiments on synthetic images with added noise, we did perturbations on the depth at the starting point and on the projection of the rotation axis. The perturbations on the starting point only affect the surface points near the starting point and do not change the surface which is far from the starting point. Shifting the rotation axis 3 or 4 pixels does not make much difference on the final results.

The surface of the vase is not as uniform as we expected. On the surface of the vase, the reflectance on the middle body part is stronger than that on the lower body part. The depth discontinuity on the curve of  $p = 0$  on the body (see Fig. 5(d)) is mainly caused by this non-uniform reflection. The depth difference

at the point of the discontinuity point is large. The surface reflectance of the cup is quite uniform so the estimated reflectance function is accurate for the whole surface of the cup. The depth values computed on the curve of  $p = 0$  is quite close to the real depth value. The depth difference at the point of discontinuity is small. For both cup and vase, the joints between the body and the handle have been successfully recovered. The joints can not be recovered by the techniques which just using occluding contours in the image sequence.

From analysis on the image sequence and the surface recovered, we know that the errors on the reconstructed surface mainly come from three sources: the error in estimation of the reflectance function, the non-uniform albedo over the object surface and the interreflection on the object. To reduce the error in the estimation of the reflectance function, we can use more singular points or use other surface points whose surface orientation and 3D location can be computed by other cues such as contours [16]. To reduce the error caused by non-uniform albedo, we can extract the reflectance function for a local area of relatively constant albedo and use the local reflectance function for the local surface recovery. This idea can be applied to an object whose surface has different colors or different albedos. This is our current area of research. Reducing the error caused by interreflection in the general case is very hard [5]. So far we do not have effective methods for removing interreflection on surfaces.

Besides shading, there are other kinds of cues such as contour and stereo available from an image sequence of a rotating object. Contours in an image sequence of a rotating object can be used to compute the location and orientation of some surface points [16]. These surface points can be alternatives when singular points are not available for extracting reflectance function. One limitation in our work is that the estimation of the surface reflectance function relies on a set of singular points on the object surface. Our current work attempts to use contours to remove this limitation. Stereo information can provide the 3D location of starting points for our surface recovery procedure. The stereo information can also remove the ambiguity in the  $q$  component of surface orientation. The integration of all the cues is not an easy task [8, 9, 17, 1]. Extending our work to more complicated surfaces will require integrating other cues.

Another extension of our work is surface recovery by rotating the object more than 90 degrees. In this way, we can get more singular points and obtain a more accurate estimate of the reflectance function. We can also construct the whole object by integrating depth recovered from different views. The intended application of our work is automatic modeling. Beside modeling the shape of an object, we also want to model the color of an object.

## 7 Acknowledgments

We would like to thank Yuji Iwahori, Esfandiar Bandari and Jeffery Beis for their useful comments. The author's email address are jplu@cs.ubc.ca and lit-

tle@cs.ubc.ca. This research was supported by grants from the Natural Sciences and Engineering Research Council of Canada and the Networks of Centres of Excellence Institute for Robotics and Intelligent Systems, Projects A-1 and ISDE-6.

## References

- [1] P. Fua and Y. Leclerc. Using 3-dimensional meshes to combine image-based and geometry-based constraints. In *ECCV-94*, pages 282–291, May 1994.
- [2] B. K. P. Horn and M. J. Brooks, editors. *Shape from Shading*. MIT Press, Cambridge, MA, 1989.
- [3] J. Lu and J. J. Little. Reflectance and shape from a rotating object. TR-95-09, UBC-CS, Vancouver, BC, 1995.
- [4] S. K. Nayar, K. Ikeuchi, and T. Kanade. Determining shape and reflectance of hybrid surfaces by photometric sampling. *IEEE ICRA*, 6(4):418–431, 1990.
- [5] S. K. Nayar, K. Ikeuchi, and T. Kanade. Shape from interreflections. *Int. J. Comp. Vision*, 6(3):173–195, 1991.
- [6] M. Oren and S. K. Nayar. Seeing beyond Lambert's law. In *ECCV-94*, pages 269–280, May 1994.
- [7] A. P. Pentland. Local shading analysis. *IEEE-PAMI*, 6:170–187, 1984.
- [8] T. Poggio, E. Gamble Jr., and J. J. Little. Parallel integration of vision modules. *Science*, 242(4877):436–440, October 1988.
- [9] R. Szeliski. Shape from rotation. In *CVPR-91*, pages 625–630, June 1991.
- [10] H. D. Tagare and R. J. deFigueiredo. A theory of photometric stereo for a class of diffuse non-Lambertian surfaces. *IEEE-PAMI*, 13(2):133–152, 1991.
- [11] L. B. Wolff. Shape understanding from Lambertian photometric flow fields. In *CVPR-89*, pages 46–52, June 1989.
- [12] L. B. Wolff. Diffuse reflection. In *CVPR-92*, pages 472–478, June 1992.
- [13] R. J. Woodham. Photometric method for determining surface orientation from multiple images. *Opt. Eng.*, 19:139–144, 1980.
- [14] R. J. Woodham. Determining surface curvature with photometric stereo. In *IEEE Conf. Robotics & Automation*, pages 36–42, Scottsdale, AZ, 1989.
- [15] R. J. Woodham. Gradient and curvature from photometric stereo including local confidence estimation. *J. Opt. Soc. Amer.*, 11(11):3050–3068, Nov. 1994.
- [16] J. Y. Zheng. Acquiring 3-D models from sequences of contours. *IEEE-PAMI*, 16(2):163–178, 1994.
- [17] J. Y. Zheng and F. Kishino. Verifying and combining different visual cues into a 3D model. In *CVPR-92*, pages 777–780, June 1992.