Reflectance and Shape from a Rotating Object

Jiping Lu Jim Little

Technical Report 95-9

April 1995

Laboratory for Computational Intelligence

Department of Computer Science

The University of British Columbia

Vancouver BC Canada V6T 1Z4

e-mail: {jplu, little}@cs.ubc.ca

Abstract

In this paper we show that the reflectance function of a rotating object illuminated under a collinear light source (where the light source lies on or near the optical axis) can be estimated from the image sequence of the object and applied to surface recovery. We first calculate the 3D locations of some singular points from the image sequence, and extract the brightness values of these singular points during the object rotation to estimate the surface reflectance function. Then we use the estimated reflectance function for surface recovery from the images of the rotating object. Two subprocedures are used in surface recovery. The first subprocedure computes the depth around a point of known depth and surface orientation by using first-order Taylor series approximation. The other computes the surface orientation of a surface point from its image brightness values in the two different images by applying the estimated reflectance function. Starting from surface points of known depth values and surface orientations and iteratively applying the two subprocedures, the surface depth and orientation are recovered simultaneously over the whole object surface. The experimental results on real image sequences of both matte and specular surfaces show that the technique is feasible and robust.

Key words: Physics based vision, Reflectance, Shape recovery, Multiview integration.

1 Introduction

An important task in computer vision is to recover surface shape from one or more 2D images. Based on physical properties of the reflectance of a surface, shading in images can be used for surface recovery. A lot of work, for example, shape from shading [2, 7], photometric stereo [14, 15, 16], and photometric sampling [4], has been done along this line. In order to use shading information, the reflectance function of the surface under recovery must be known. In most cases, the reflectance function is assumed to be a certain type [10, 4]. The most commonly used assumption is Lambertian reflectance [12, 7] because of its simplicity. In photometric stereo [15, 16], 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 recovery.

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 [13, 11, 5]. The empirical photometric stereo 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 procedure must be accurate otherwise it will introduce errors in surface recovery. The calibration process may become difficult or impossible when we do not have a calibration object of the same reflectance function as the object being imaged.

We attempt to compute the surface reflectance directly from 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 or detected. The images are taken by a fixed camera as the object rotates around a vertical axis. 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. Under the illumination of a collinear light source, all the components of reflectance at a surface point, such as the specular component, diffuse component and other components [11] are functions of the incident angle i, which is defined as the angle between the illuminant direction and the surface normal at a surface point. Thus, in this case, the total image brightness of a point on a uniform surface is a function of the incident angle i. This fact makes it easy to estimate the reflectance function and apply the reflectance function for surface recovery. The estimate of the reflectance function is based on a set of singular points. A singular surface point, of which the normal is opposite to the viewing direction, is also a singular point of brightness value in an image because the incident angle at a singular point is zero. Under the assumed orthographical projection, the x, y coordinates of a singular point can be directly found from an image by searching for a point of local maximum brightness. The z coordinate of a singular point can be found on the contours of the object image taken after the object has rotated by 90 degrees. Once we known the 3D location of a singular point, we can track and record the 3D location, the incident angle and its brightness value over the image sequence taken during the 90 degree rotation. From the incident angles and the corresponding image brightness values recorded, we build the reflectance function of the surface. For most surfaces the reflectance function under a collinear light source is strictly monotonic so we can compute its inverse. Using the inverse function, from the brightness of an image point we estimate the incident angle i for the corresponding 3D point. As the surface normal has only two degrees of freedom, two images of a 3D point will be sufficient to determine the surface normal with a two-way ambiguity if the object rotation angles are known when the two images are taken. The ambiguity can be removed by using continuity constraint and some 3D points of known positions on the surface during the recovery process.

To get depth and surface orientation (i.e., the surface normal at a point), we use two images. One image is the image taken before the rotation of the object. The other image is taken after the object has rotated by a certain angle. The depth and surface orientation are computed at every point in the first image. Two subprocedures are used to compute the depth and surface orientation. The first subprocedure computes the depth values of the neighbor points of an image point of known depth and surface orientation. For an image point of known depth and surface orientation, the depth values in the small neighborhood of this point are estimated by first-order Taylor series approximation. The second subprocedure computes the surface orientation for an image point using the depth value at this image point. From the depth value of an image point in the first image, its brightness value in the other image is found. By using the reflectance function, the two brightness values, and the rotation angle when the second image is taken , the surface orientation for the image point is determined. As depth and surface orientation at a singular point are known, we apply the first subprocedure to a singular point to compute the depth values for its neighbors. After the depth values of the neighbor points have been obtained, we apply the second subprocedure to the neighbor points to compute surface normals for these points. Applying the first subprocedure to the neighbor points, we expand depth over a larger area, and applying the second subprocedure again, we compute surface normals for the expanded area. In this way, taking singular points as starting points and iteratively applying the two subprocedures, we spread the computation over the entire image. The computation for depth basically is an integration process so it is robust against image noise. The computation for surface orientation can be done by look-up table just as in photometric stereo [15, 16].

Section 2 introduces the assumptions and the experimental setting for our work. Section 3 presents a method for building the reflectance function from an image sequence of a rotating object under a collinear light source. Section 4 describes the process for recovering surface shape and orientation by using the reflectance function obtained. Section 5 shows some experimental results on real image sequences. The final section discusses experimental results and future work.

2 Assumptions and Experimental Conditions

The imaging geometry is shown in Fig. 1. The object is on a turntable whose rotation angle can be controlled or detected. The Y axis coincides with the rotation axis of the turntable. The light source and the camera are fixed and point in the same direction. 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. Thus a surface point (x, y, z) is projected on an image point (x, y). To obtain the projection of the rotation axis in the images, a vertical black line on a board is aligned with the rotation

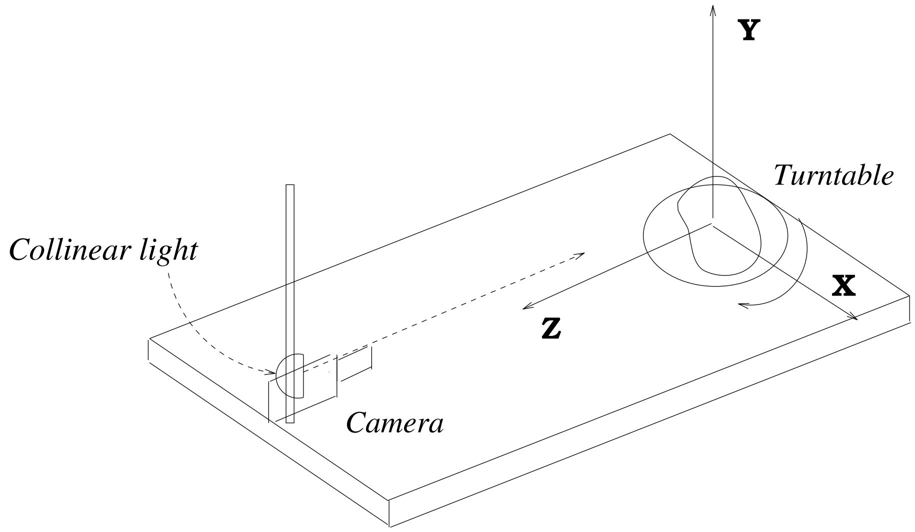

Figure 1: Experimental setup

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 α degree rotation, the 3D location $(x_\alpha, y_\alpha, z_\alpha)$ of this point is $(x_\alpha,\ y_\alpha,\ z_\alpha) = (x\cos\alpha - z\sin\alpha,\ y,\ x\sin\alpha + z\cos\alpha)$ and the surface orientation $(p_\alpha, q_\alpha, 1)$ of this point is determined by

$$p_\alpha = \frac{p\cos\alpha + \sin\alpha}{\cos\alpha - p\sin\alpha}, \tag{1}$$

$$q_\alpha = \frac{q}{\cos\alpha - p\sin\alpha}. \tag{2}$$

In surface recovery, the depth value z_α and surface orientation are recovered at every point in an image. In our experiment, the depth and the surface orientation in the first image, which is taken before the object rotation, are recovered.

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 $R(i, e, g)$ [14]. As shown in Fig. 2, the incident angle i is the angle

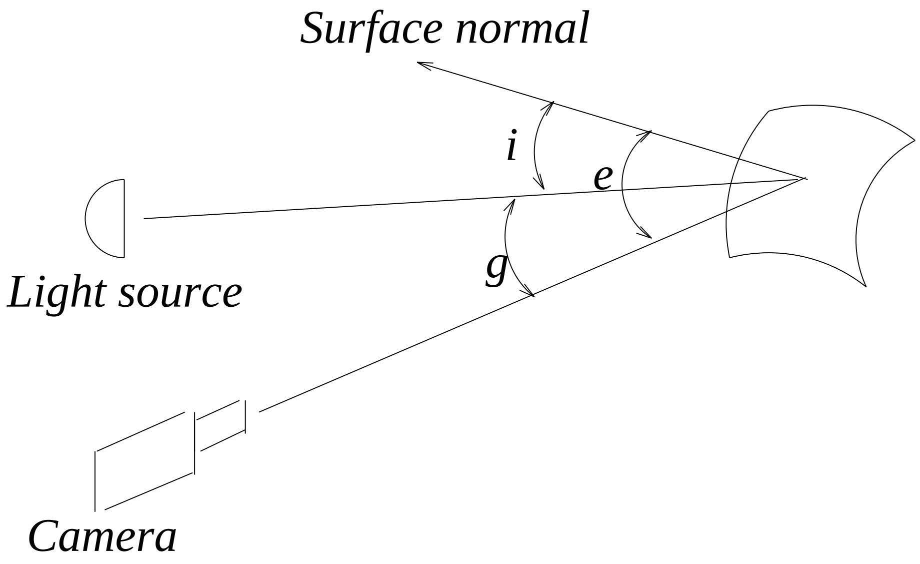

Figure 2: The image brightness is a function of the angles i, e and g

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. 3, 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 the reflectance such as the specular component, diffuse component and other components defined in Tagare's paper [11] are functions of the incident angle i only. Thus for the surface point (x, y, z), its image brightness value can be written as

$$E(x, y) = R(i(x, y)) \tag{3}$$

where $i(x, y)$ is the incident angle at point (x, y, z).

The reflectance function has maximum brightness value when $i = 0$ and minimum brightness value when $i = \pi/2$. We assume the function is strictly monotonic. This assumption is true for most surfaces. The most important aspect of the reflectance function $R(i)$ is that it is a function of one variable. This makes the relation between brightness value and surface orientation very simple. The simplicity of the reflectance function makes it much easier to estimate the reflectance function and recover the shape of the object.

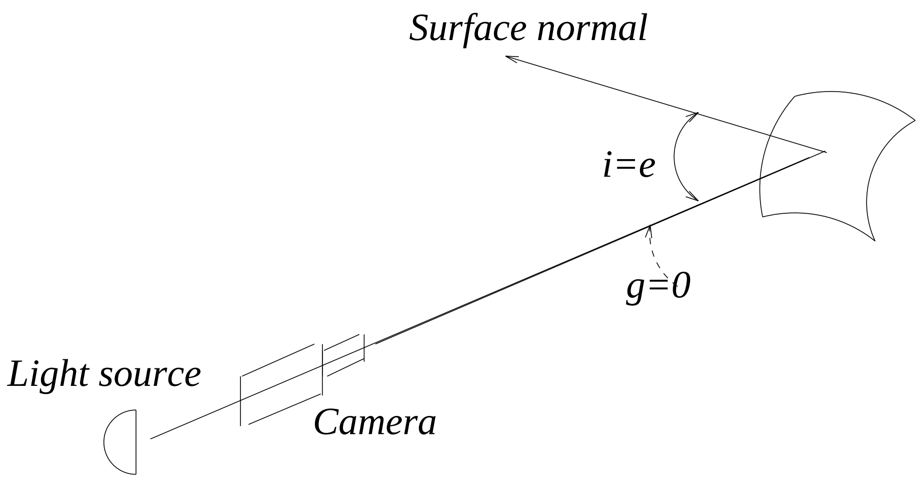

Figure 3: Under collinear light, the image brightness only depends on angle i

3 Estimating the Reflectance Function

The estimate of the reflectance function is based on singular points of surface orientation (0, 0, 1). These singular points also have singular brightness values in an 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. After the rotation, the incident angle of the original singular point is 90 degrees, and the 3D location of this point is $(-z_s, y_s, x_s)$. The fact that the incident angle of a surface point is 90 degrees tells us that the image of this point is on the contour of the object image. Finding this image point is not difficult since its y coordinate is already known and the tangent line of this point along the contour is parallel to the Y axis. Considering generic surfaces, we assume there are some singular points whose images are available in the first image and will not be occluded during 90 degree rotation. Given the 3D locations of these singular points in the first image, we can track and record the 3D locations, the incident angles, and corresponding brightness values of these points over the image sequence. For the singular point (x_s, y_s, z_s), after the object has rotated by an angle $\theta_i \leq 90°$, the 3D location of this point $(x_i,\ y_i,\ z_i) =$

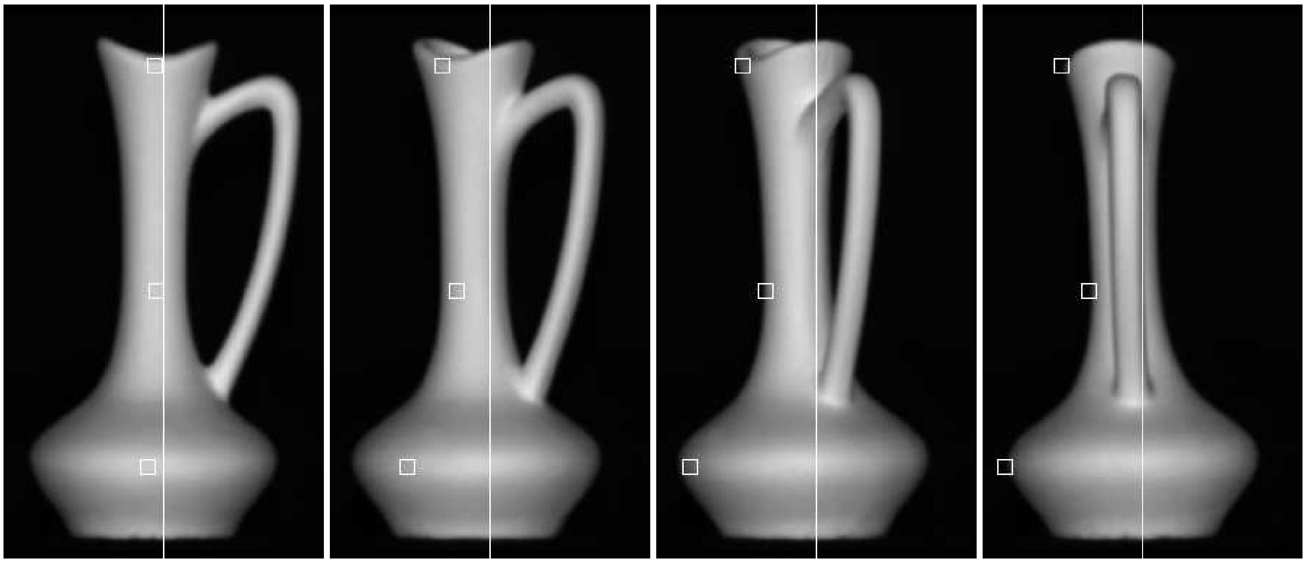

(a) (b) (c) (d)

Figure 4: Tracking the singular points over image sequence

$(x_s \cos\theta_i - z_s \sin\theta_i,\ y_s,\ x_s \sin\theta_i + z_s \cos\theta_i)$. The image brightness $E(x_i, y_i)$ of this point can be directly obtained from image point (x_i, y_i) in the image which is taken after a θ_i degree rotation of the object. It is easy to show that the incident angle for point (x_i, y_i, z_i) is θ_i. From the brightness $E(x_i, y_i)$ and the corresponding incident angle θ_i, we can build the reflectance function for the surface. The dark line in Fig. 5 shows the reflectance function $E = R(i)$ obtained from an image sequence of a rotating vase. Each image is taken after the object has rotated by another 5 degrees and nineteen images are taken during a 90 degree rotation. The images in Fig. 4 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. The centers of the small square boxes in each image denote the points tracked over the image sequence. The 3D positions of the three singular points are obtained from the first image and the last image in the sequence. So the image brightness values are sampled from the three points after each successive 5 degree rotation. The average of the brightness values of the three points is used to build the reflectance function. The estimated function

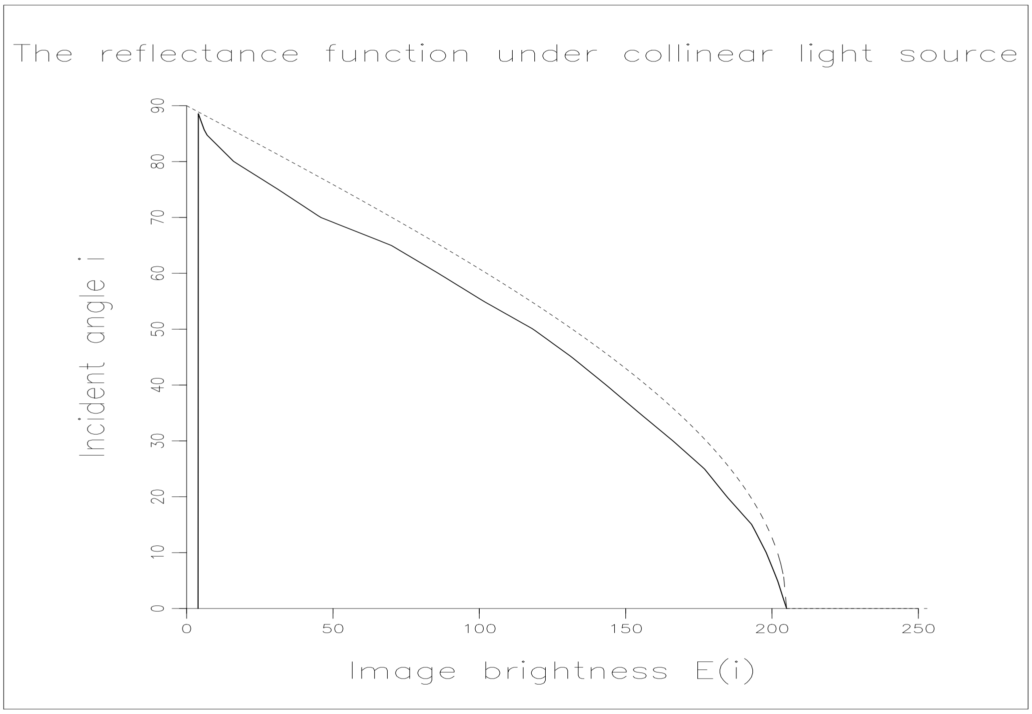

Figure 5: The reflectance function obtained from the image sequence of the rotating vase

are shown in Fig. 5. 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 the interreflection make the brightness values of these points much bigger than these points on the body part. Looking at the image after the 90 degree rotation, we can see that the images of the singular points on the body lie on the occluding contour against the background and their brightness values are the same as the background brightness value but the images of the singular points on the handle lie on the occluding contour against the body of the vase and their brightness value are much higher than the background brightness value. Thus using the singular points on the handle to estimate the reflectance function will introduce errors. However, the 3D locations of these singular points can be easily found and can be used in the surface recovery.

The function we actually used for surface recovery is the inverse function $i = R^{-1}(E)$. Since the reflectance function obtained is strictly monotonic, its inverse exists. The inverse function is linearly interpolated for every integer brightness value ranging from 0 to 255. The vase is made of clay and the surface of the vase is considered a matte surface. The dotted line in Fig. 5 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

After the reflectance function has been obtained, surface recovery can be done by using any two images in the image sequence of the rotating object. In this paper, the first image and the image taken after the object is rotated by a certain angle are used. 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 $(x + \delta x, y + \delta x)$ 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 determines surface orientation (up to a two-way ambiguity) from image brightness values in two images. 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 α 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

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

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

$$p_1 = \frac{p_0 \cos\alpha + \sin\alpha}{\cos\alpha - p_0 \sin\alpha}, \tag{6}$$

and

$$q_1 = \frac{q_0}{\cos\alpha - p_0 \sin\alpha}. \tag{7}$$

Substituting p_1 and q_1 in Equation 5, we get

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

The equation can be simplified to

$$\cos i_1 = \cos i_0 (\cos\alpha - p_0 \sin\alpha). \tag{9}$$

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

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

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

The geometric explanation for the solution of p_0 and q_0 is shown in Fig. 6. Here we assume $image_0$ and $image_1$ are taken from two different viewing directions with the object fixed. We denote the viewing direction for $image_0$ as vector v_0 and the viewing direction for $image_1$ as vector v_1. All the surface normals with incident angle i_0 to the viewing direction vector v_0 form a cone. All the surface normals with incident angle i_1 to the viewing direction v_1 form another cone. The intersection of the two cones is two vectors. The two vectors are symmetric about the XZ plane. Only one of the two vectors overlaps with the surface normal. So there is a two-way ambiguity in the solution for surface orientation which is caused by the two solutions for q_0 in Equation 11.

The surface recovery procedure starts at the image points whose depth and surface orientations are known. These starting points could be the singular points we used to compute the reflectance function. For each starting point of known depth, we use the first subprocedure to compute the depth for the neighbors of the starting point. For each

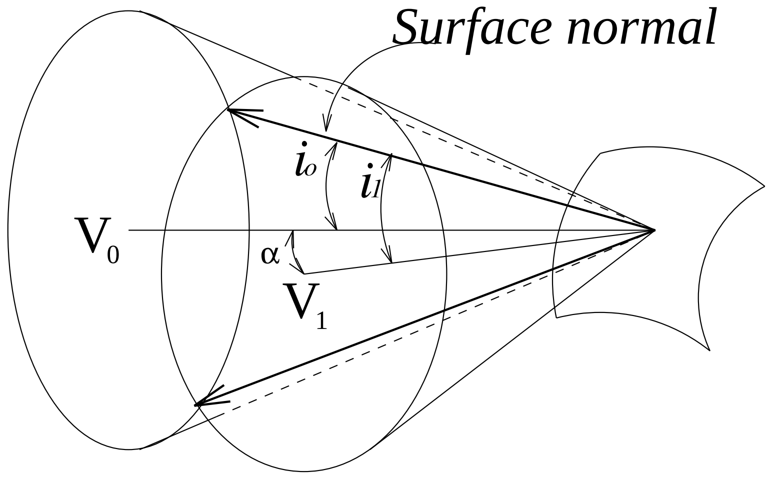

Figure 6: The solution for surface orientation

neighbor point, using the depth value computed by the first subprocedure, we compute the location of its correspondence and the brightness value of its correspondence in the other image. Then we use the second subprocedure with the two brightness values found to compute surface orientation for every neighbor point. Using the first subprocedure on the neighbor points, we expand the depth over a larger area. Using 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.

When we compute a new depth value z' in the y direction, we have $z' = z + q\delta y$. Thus the 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. 7(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. 7(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 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. 7(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

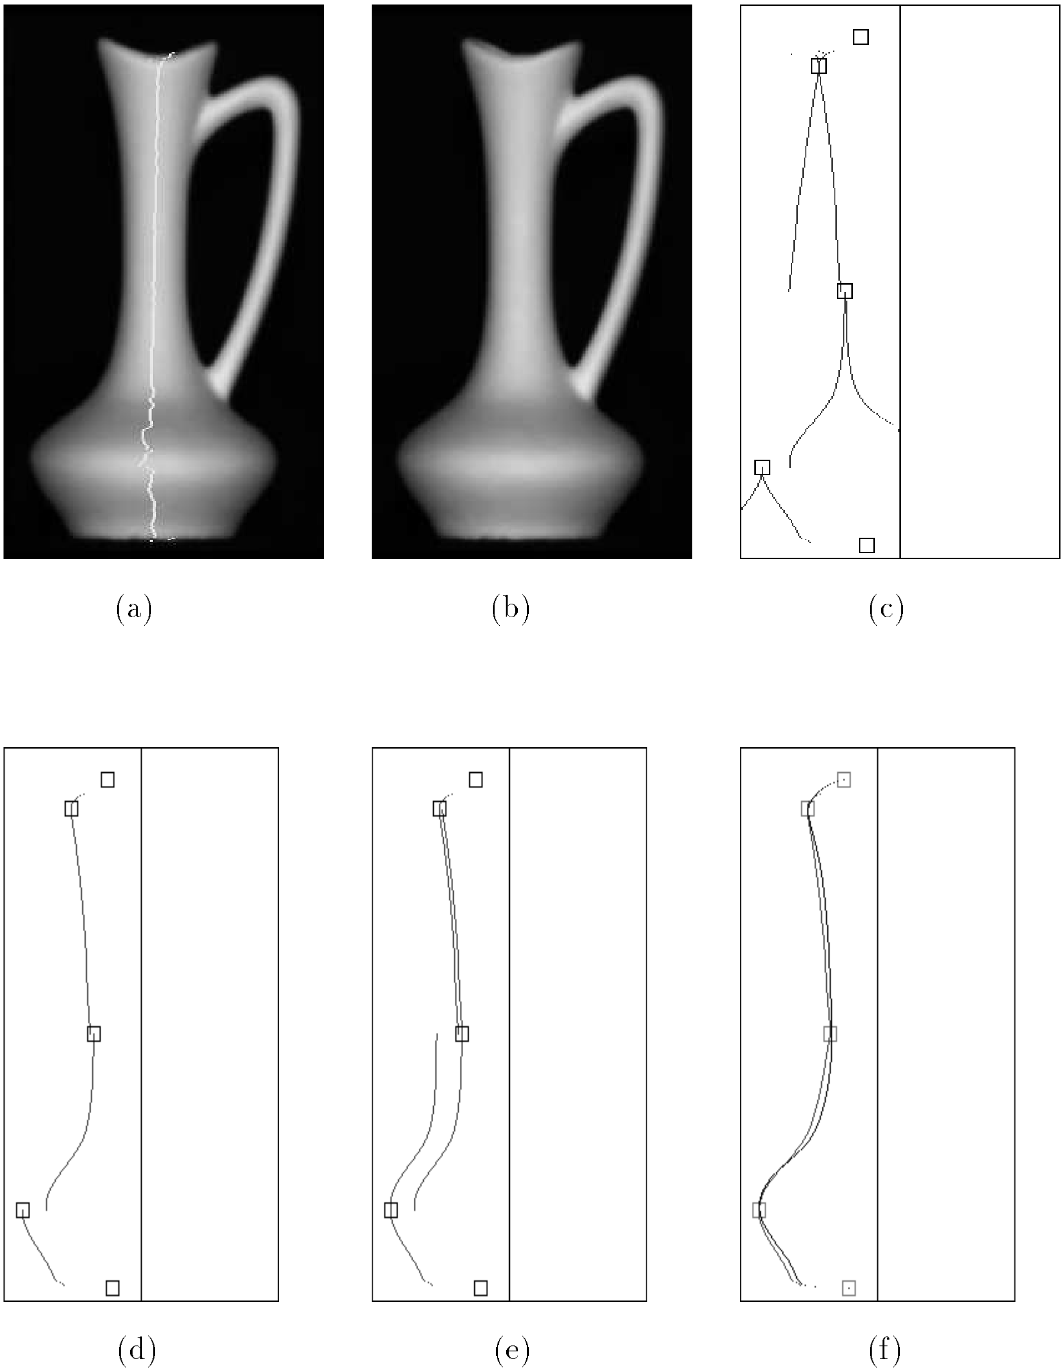

Figure 7: Computing depth on a curve and removing the ambiguity in the depth

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. 7(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

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

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. 7(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 our experiment, we use a calibrated image facility (CIF) [16] built in our lab to control the rotation of the object and the imaging condition. Although the camera we used is a 24-bit RGB camera, we only use one of the three B&W images. 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 to the object on a turntable. We considered putting a half-silvered mirror in front of the camera to make the viewing direction and illuminant direction precisely collinear, but it turned out to be unnecessary because we did not observe the effects caused by the small angle between the viewing direction and the illuminant direction. In practice, the radiance of the light source is not constant over 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 orthographical projection is assumed for a reasonable approximation. 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 effect. Four images from the image sequence of the vase are shown in Fig. 4. Images (a), (b), (c) and (d) are, respectively, the images taken after 0, 30, 60 and 90 degrees rotation of the vase. The estimated reflectance function of the vase is shown in Fig. 5 as dark line. We track the three singular points over the image sequence to estimate the reflectance function. The brightness value for the reflectance function is the average of the brightness values of the three points. The singular points in the last image can also be used for better estimation. For the time being, we only use singular points in the first image.

In surface recovery, the first image (Fig. 7(a)) and the image taken after a 10 degree rotation (Fig. 7(b)) are used. We first compute depth and surface orientation along a curve of $p = 0$ (see Fig. 7(a)). The curve passes through the three singular points which are used to estimate the reflectance function. 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. 7(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. Fig. 5 is the partially constructed surface of the vase. For each unrecovered area, we repeat the process of computing depth along the curve of $p = 0$, removing the depth ambiguity on the curve, then expanding the depth values in the x direction, until the depth values on the entire image are recovered.

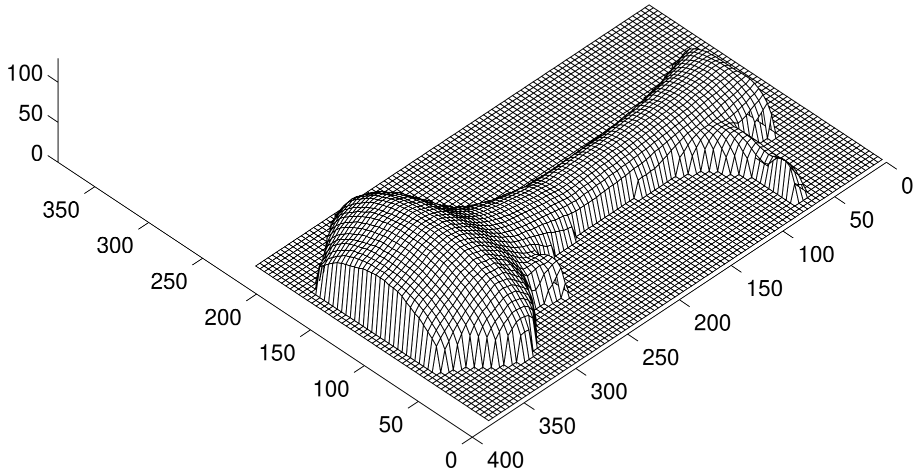

Figure 8: The surface height of the partially recovered vase

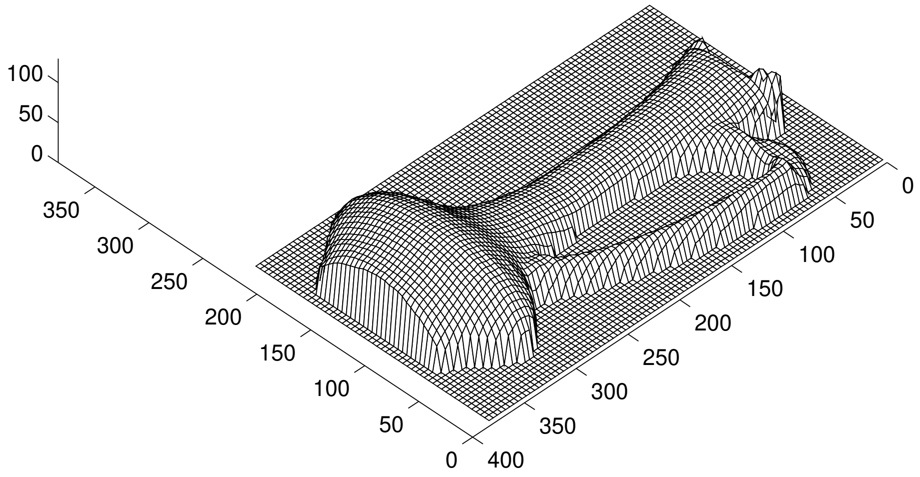

Figure 9: The surface height of the recovered vase

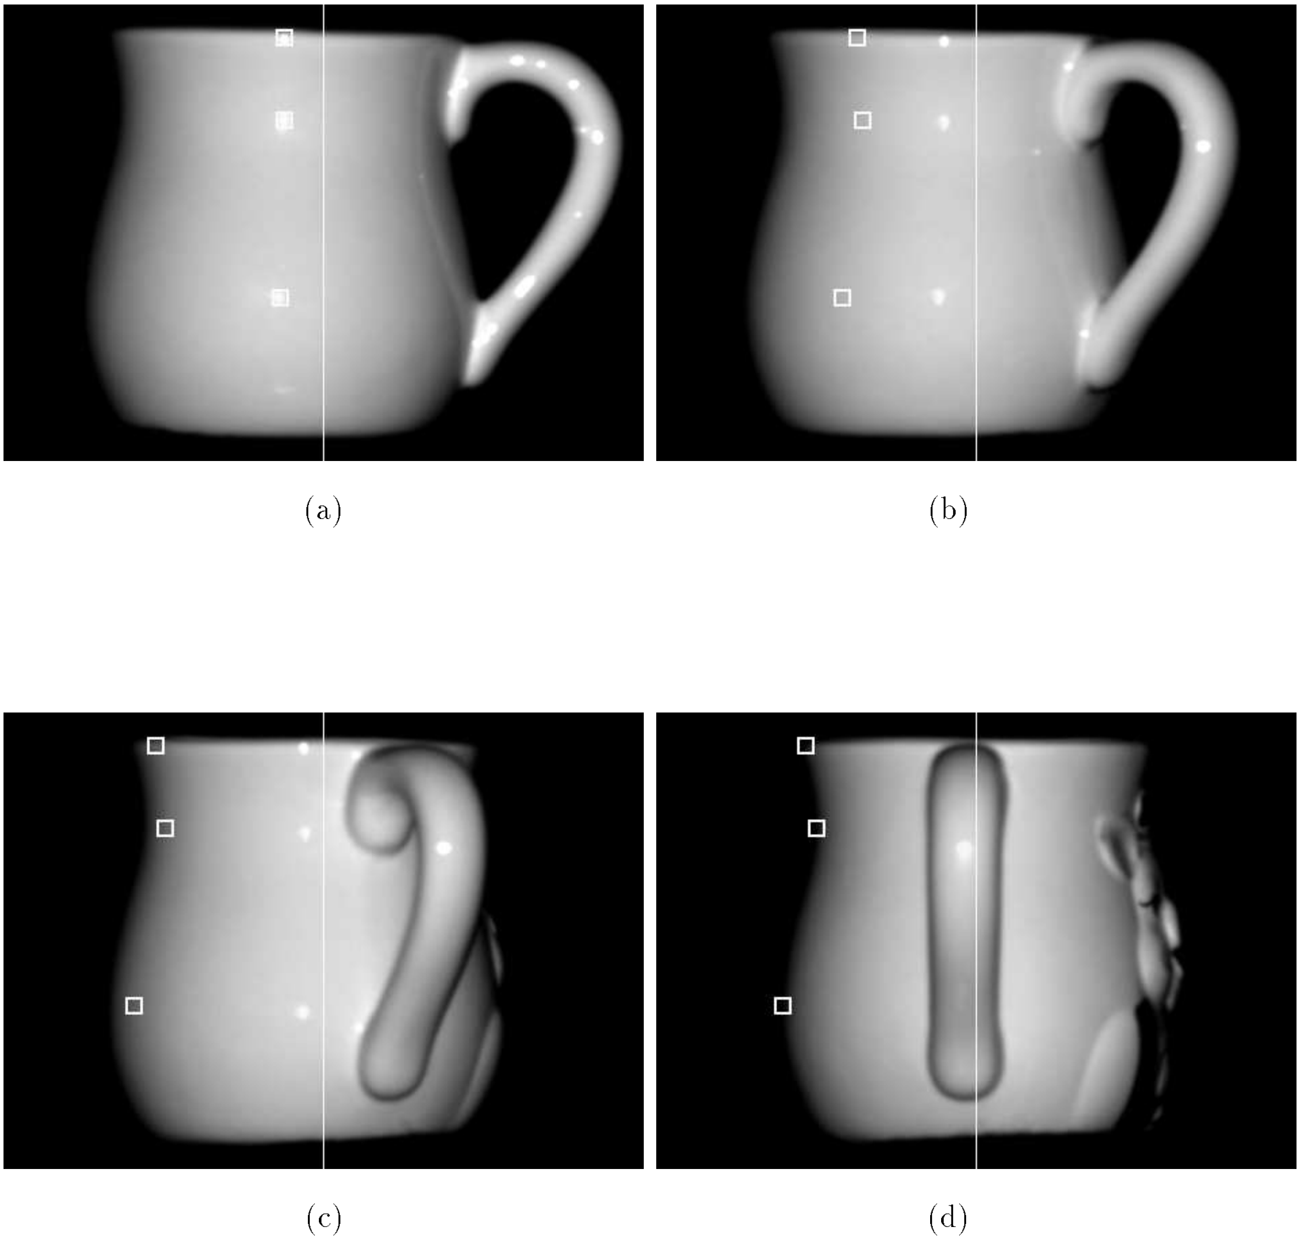

Figure 10: Tracking the singular points over the image sequence of a cup

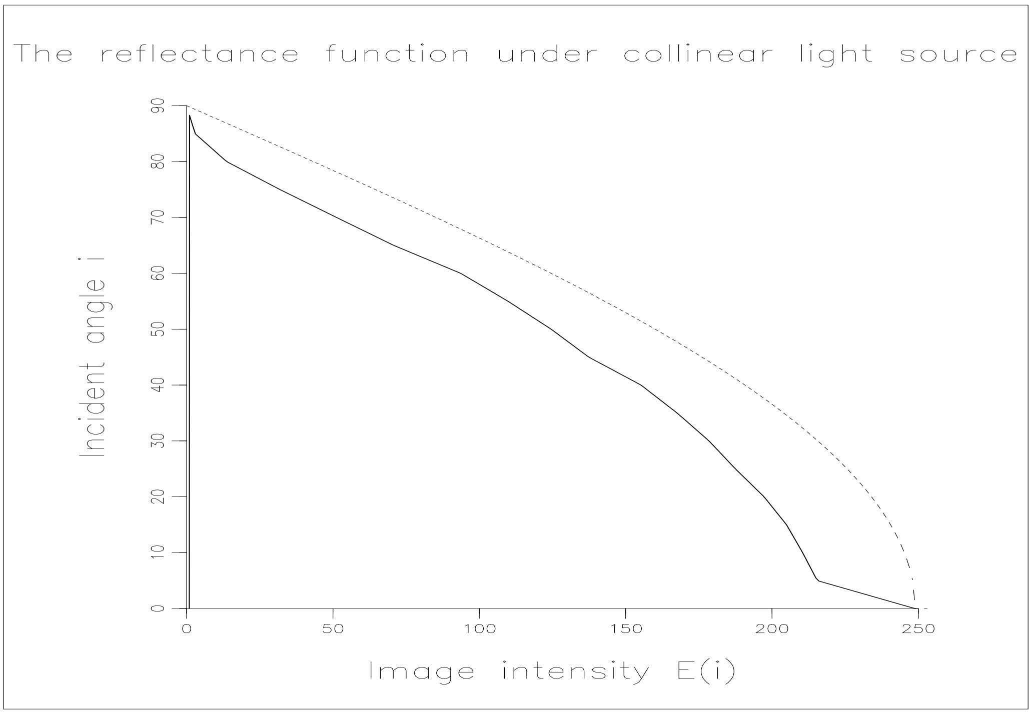

Figure 11: The reflectance function obtained from the image sequence of a rotating cup

During surface recovery, depth and orientation are computed at every pixel on the first image. The brightness values on the second image are interpolated between pixels as the projection of a 3D point on the second image may be located between pixels. The surface plot of the final recovered depth of the vase is shown in Fig. 9. 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 and surface orientation data.

We also experimented with a porcelain cup. Nineteen images are taken during 90 degree rotation. The rotation angle between successive images is 5 degrees. The reflectance of the cup presents a strong specular reflection. There is a peak brightness value at singular points on the surface of the cup. Fig. 10 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 surface reflectance function. In Fig. 11, the dark line denotes the reflectance function of the cup, and the dotted line is the Lambertain reflection with the same maximum brightness value as that of the estimated reflectance function. The difference between the two reflectance functions is obvious. Two images (see Fig. 12(a) and

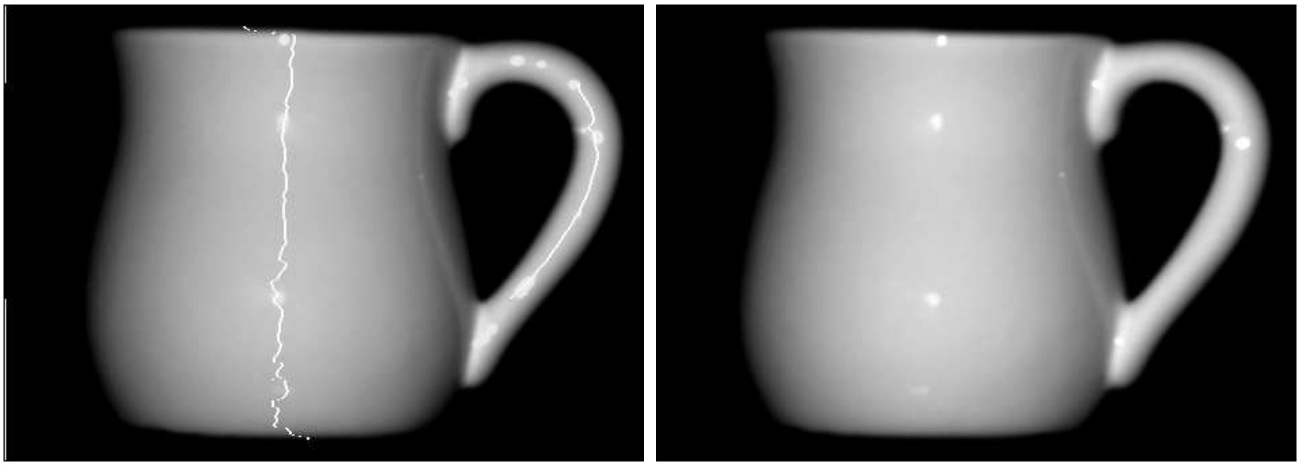

(a) (b)

Figure 12: Images used for surface recovery and curves of $p = 0$

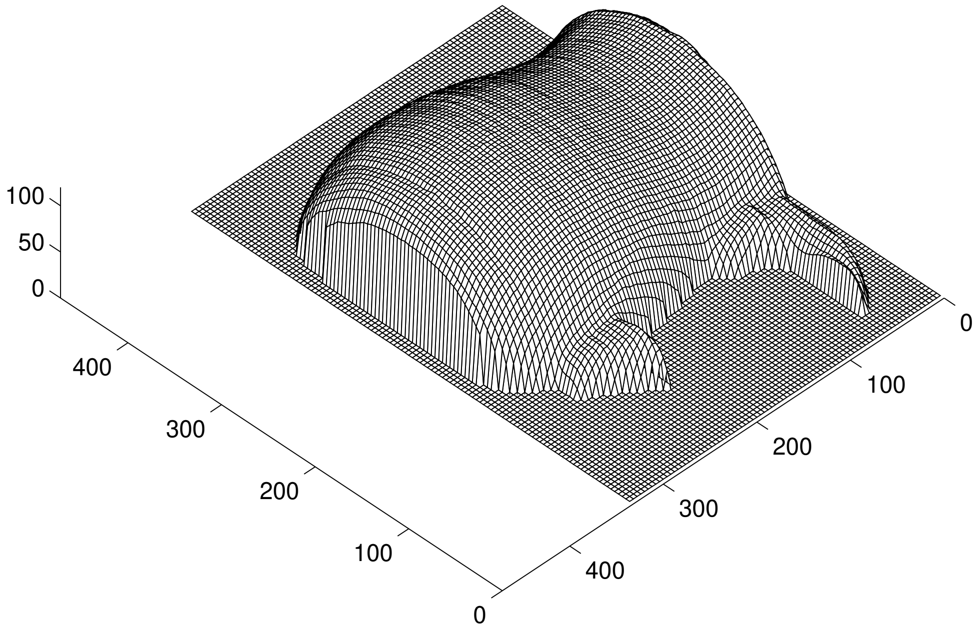

Figure 13: The height plot of the partially recovered cup

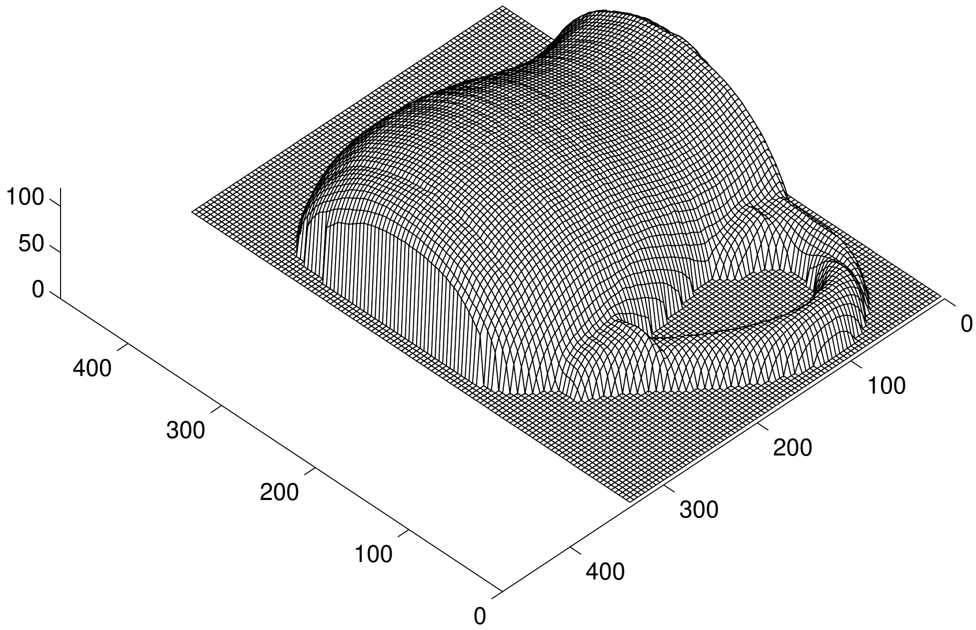

Figure 14: The height plot of the recovered cup

(b)) are used for surface recovery. The second image is the image taken after the object has made a 10 degree rotation. We start the computation on the white curve on the body of the cup. The depth ambiguity on the curve is removed by the continuity constraint and the known depth values at the singular points. The depth on the curve is made continuous by distance-weighted averaging. A partial surface (see Fig. 13) is constructed along the x direction from the curve of $p = 0$. The surface of the handle of the cup is constructed in the same way as we did for the body of the cup. The white curve on the handle in Fig. 12(a) is a curve of $p = 0$. The final recovered surface is shown in Fig. 14.

6 Discussion and Future Work

The results obtained show that the technique is feasible and robust for surface recovery. Since we don't assume the reflectance function has any particular form, the reflectance estimation method described in this paper can be applied to surface of isotropic reflection.

By increasing the number of images taken during rotation, the accuracy of the estimated reflectance function can be improved. The surface recovery procedure exploits the photometric constraint through the reflectance function and the binocular stereo constraint by using two images taken when the object has rotated by different angles. The procedure to get surface depth integrates surface orientation. From theory this procedure is robust against image noise. 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.

We did experiments on real image sequence for two objects. One is a pottery vase. The other is a porcelain cup. The surface of the pottery vase is a matte surface. 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. 7(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 joint between the body and the handle has been successfuly recovered. The joint can not be recovered by just using extremal contours because it can not be seen as extremal 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 image cues such as contours [17]. 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 [3]. So far we do not have effective methods for removing interreflection on surfaces with non-Lambertian reflectance.

Besides shading, there are other kinds of cues such as contour and stereo available from an image sequence of a rotating object. These cues can be used in different ways. One way to use the contour cue is to derive a local reflectance function from contours. In our work we use singular points to estimate the reflectance function and assume these points are available. It has been shown that contours in an image sequence of a rotating object can be used to compute the location and orientation of surface points [17]. These computed surface points can be an alternative when singular points are not available for extracting reflectance function. One way to use the stereo information is to get the 3D location of some surface points and use these points as starting points for our surface recovery procedure. The stereo information can also be used to remove the ambiguity in the q component of surface orientation. The integration of all the cues is not an easy task [8, 9, 6, 18, 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 little@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.G. Leclerc. Using 3-dimensional meshes to combine image-based and geometry-based constraints. In *Proc. 3rd European Conference on Computer Vision*, pages 282–291, May 1994.

[2] B. K. P. Horn and M. J. Brooks, editors. *Shape from Shading*. MIT Press, Cambridge, MA, 1989.

[3] S. K. Nayar, K. Ikeuchi, and T. Kanade. Shape from interreflections. *International Journal of Computer Vision*, 6(3):173–195, 1991.

[4] Shree K. Nayar, Katsushi Ikeuchi, and Takeo Kanade. Determining shape and reflectance of hybrid surfaces by photometric sampling. *IEEE Transactions on Robotics and Automation*, 6(4):418–431, 1990.

[5] Michael Oren and Shree K. Nayar. Seeing beyond Lambert's law. In *Proc. 3rd European Conference on Computer Vision*, pages 269–280, May 1994.

[6] Sharath Pankanti, Anil K. Jain, and M. Tuceryan. On integration of vision modules. In *Proc. IEEE Conf. Computer Vision and Pattern Recognition, 1994*, pages 316–322, June 1994.

[7] A. P. Pentland. Local shading analysis. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 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 *Proc. IEEE Conf. Computer Vision and Pattern Recognition, 1991*, pages 625–630, June 1991.

[10] Hemant D. Tagare and Rui J.P. deFigueiredo. Simultaneous estimation of shape and reflectance maps from photometric steroeo. In *Proc. 3rd International Conference on Computer Vision*, pages 340–343, 1990.

[11] Hemant D. Tagare and Rui J.P. deFigueiredo. A theory of photometric stereo for a class of diffuse non-Lambertian surfaces. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 13(2):133–152, 1991.

[12] Lawrence B. Wolff. Shape understanding from Lambertian photometric flow fields. In *Proc. IEEE Conf. Computer Vision and Pattern Recognition, 1989*, pages 46–52, June 1989.

[13] Lawrence B. Wolff. Diffuse reflection. In *Proc. IEEE Conf. Computer Vision and Pattern Recognition, 1992*, pages 472–478, June 1992.

[14] R. J. Woodham. Photometric method for determining surface orientation from multiple images. *Optical Engineering*, 19:139–144, 1980.

[15] R. J. Woodham. Determining surface curvature with photometric stereo. In *IEEE Conf. Robotics & Automation*, pages 36–42, Scottsdale, AZ, 1989.

[16] R. J. Woodham. Gradient and curvature from photometric stereo including local confidence estimation. *Journal of the Optical Society of America*, 11(11):3050–3068, November 1994.

[17] J. Y. Zheng. Acquiring 3-D models from sequences of contours. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 16(2):163–178, 1994.

[18] J. Y. Zheng and Fumio Kishino. Verifying and combining different visual cues into a 3D model. In *Proc. IEEE Conf. Computer Vision and Pattern Recognition, 1992*, pages 777–780, June 1992.