2008－CVIM－162 (84)
2008／3／11

Manifold-Based Photometric Stereo with Cast Shadows

Lulu Chen (Univ. of Tokyo) Takahiro Okabe (Univ. of Tokyo)
Imari Sato (NII) Yoichi Sato (Univ. of Tokyo)

This paper presents a photometric stereo method for determining an object's shape from its appearance manifold by estimating the similarity of appearances observed at points under varying illumination. Assuming no cast shadows on the object's surface, we show that for a pair of surface points, the similarity of their observed intensities under varying illumination is closely related to the similarity of their corresponding surface normals. However, if the object is concave, there are cast shadows on the surface which alter the observed intensity of surface points and result in incorrect similarity estimation. After using a similarity measure, we find and combine similar vectors to interpolate in areas of cast shadow to create new observation values, which means that we remove cast shadows in input images. Then the object's surface normals can be estimated from the observation vector space using a dimensionality reduction technique. Unlike most previous shape reconstruction methods, our method does not require any particular reflectance model which makes it applicable to a wide variety of object materials.

1 Introduction

3D shape reconstruction of an object by photometric stereo is a central topic in computer vision [19]. The appearance of an object is determined by several factors including illumination, the viewing geometry, the surface shape and the reflectance properties of the object surface. Changing any one of these factors should change the object's appearance. Generally these factors are nonlinearly related to the object's appearance, thus it's an inverse problem to estimate each of them. Currently most previous methods estimate some of these factors by assuming that some other factors are already known. Based on different assumptions, different work addresses issues such as 3D shape reconstruction, reflectance property analysis, and illumination estimation.

Horn introduced the first photometric techniques of shape from shading in the early 1970s [9], and subsequently there have been many novel approaches to solve this problem. Zhang and Tsai made a good survey on shape from shading methods [22]. The classic photometric stereo approach presented in [9, 21] recovers the shape of a Lambertian object from multiple images of the object taken under known light directions. Then some work focuses on extending this approach to more applications, such as photometric stereo methods for non-Lambertian surfaces [10, 14, 4, 6] and uncalibrated photometric stereo [3, 7, 1].

The research above is ideal when a reflectance model for the object is known. However, in practical situations object materials are more complex and have varying reflectance properties on the object surface [5]. According to this idea, Hertzmann and Seitz used calibration objects with known shape as examples instead of computing a reflectance map [8]. They take images of one or more example objects with similar materials and known geometry under the same illumination conditions. Then they estimate the object's surface normals by finding the points with the same intensities on the example objects. In the previous approach, they assume that the example objects are made of the same material and need to be calibrated. Recently they improved their approach by using only a small number of example objects without calibration and extending the application to non-uniform surface materials. However, materials which can represent the complex real-world materials well still need to be sought. In a different approach, Koppal and Narasimhan proposed a clustering-based technique to find iso-surface normal clusters of scene points without requiring knowledge about materials and lighting [11]. Their approach shows how effective it is to analyze the temporal variation in the appearance of a scene to obtain meaningful geometric structure. However, the limitation of their method is that the illumination must be continuous with time and they need another shape reconstruction technique such as uncalibrated photometric stereo to complete the 3D shape reconstruction.

In our previous study [16], we proposed a novel photometric stereo method based on the similarity of the appearance changes observed at points on its surface under varying illumination. Assuming a convex object under distant illumination and orthographic projection, we prove that for a pair of surface points, the similarity of their observed intensity sequence under varying illumination is closely related to the similarity of their corresponding surface normals. We estimate the similarity of the appearances in the observation vector space, which contains a manifold that preserves the surface normals' geometry. Thus, we use a dimensionality reduction technique to extract the in-

trinsic geometry from the observation vector space as captured in geodesic distances on a manifold.

However, if the object is concave, there are cast shadows on the surface which alter the observed intensities of surface points and result in incorrect similarity estimation in the appearance manifold. The cast shadow elements among the observation vectors are like a big missing part with wrong information. To overcome this problem, we calculate a similarity metric between observation vectors on the high-dimensional manifold that can find the similarities of the observation vectors even when there are cast shadow elements. We then combine similar vectors by interpolating new observation values in areas of cast shadow which means that we remove cast shadows in input images. Then the interpolated images will be used to estimate the surface normals' structure.

Unlike most previous shape reconstruction methods, we estimate an object's shape from only a set of images of the object, and we don't require any particular reflectance models, reference objects or light directions as apriori knowledge, which makes our method applicable to a wide range of materials. Also, our method is capable of estimating a dense surface normal distribution by analyzing similarities of appearance changes, as opposed to the finite number of iso-normal clusters estimated in [11].

Our method will be introduced through section 2. Section 2.1 analyzes the similarity measure of appearance changes. Section 2.2 and 2.3 explain the influence of cast shadows and give the algorithm to solve the problem, respectively. Section 2.4 describes the main idea and steps of our shape reconstruction method based on the dimensionality reduction technique. Section 3 shows some experimental results of interpolated images and 3D shape reconstruction. The conclusion is presented in section 4.

2 Proposed Method

2.1 Manifold-based Photometric Stereo [16]

Considering a set of images of an object which is captured under n varying illumination directions, let I_p^i be the intensity of each surface point p ($p = 1, ..., m$) under the ith illumination, or the ith image. Then for each surface point p, we can obtain an intensity sequence vector $\mathbf{I}_p$, which is known as its *observation vector*:

$$\mathbf{I}_p = \left[I_p^1, I_p^2, ..., I_p^n\right]^T \tag{1}$$

On the observed object surface we will obtain m observation vectors, each of which can be considered a point in an n-dimensional vector space. This high-dimensional vector space contains the input appearance manifold, which preserves the appearance changes observed at the surface points caused by changing illumination. To illustrate the similarity of the observation vectors is related to the similarity of the surface normals, we will analyze the similarity measure as follows.

Based on the assumptions above, a normalized observation vector of pixel p is redefined as below:

$$\mathbf{I}_p = \frac{[I_p^1, I_p^2, \ldots, I_p^n]^T}{\sqrt{\sum_{i=1}^{n} (I_p^i)^2}} \tag{2}$$

We assume no local effects like cast shadows and interreflections. Then the brightness of point p under the i-th illumination is of the form below:

$$I_p^i = \rho_p g(\theta'_{ip}, \theta'_{vp}, \phi'_{ip} - \phi'_{vp}) \tag{3}$$

where ρ_p is the reflection coefficient of point p, $(\theta'_{ip}, \phi'_{ip})$ and $(\theta'_{vp}, \phi'_{vp})$ are the incident and viewer directions defined in the local coordinates according to point p. g is the function for the amount of the incident light from the direction $(\theta'_{ip}, \phi'_{ip})$ reflected on the point p to the viewer direction $(\theta'_{vp}, \phi'_{vp})$.

As defined by the global coordinate frame, $\mathbf{n}_p$ is the surface normal of surface point p, $\mathbf{l}_{ip}$ and $\mathbf{v}_p$ are the incident direction of the i-th illumination and the viewer direction respectively. Since $\cos\theta'_{ip} = \mathbf{n}_p^T\mathbf{l}_{ip}$, $\cos\theta'_{vp} = \mathbf{n}_p^T\mathbf{v}_p$, and $\cos(\phi'_{ip} - \phi'_{vp}) = [\mathbf{l}_{ip}^T\mathbf{v}_p - (\mathbf{n}_p^T\mathbf{l}_{ip})(\mathbf{n}_p^T\mathbf{v}_p)]/[\sin\cos^{-1}(\mathbf{n}_p^T\mathbf{l}_{ip})\sin\cos^{-1}(\mathbf{n}_p^T\mathbf{v}_p)]$, (3) can be modified into:

$$I_p^i = \rho_p g'(\mathbf{n}_p, \mathbf{l}_{ip}, \mathbf{v}_p) \tag{4}$$

As distant light and orthographic projection are assumed, the incident direction and the viewer direction are constant among the different surface points, that is, $\mathbf{l}_{ip} = \mathbf{l}_i$ and $\mathbf{v}_p = \mathbf{v}$. Then we rewrite (4) as:

$$I_p^i = \rho_p g'(\mathbf{n}_p, \mathbf{l}_i, \mathbf{v}) \tag{5}$$

Thus, from (5) we can see that, ρ_p is an independent reflectance coefficient, then the difference in appearance $\mathbf{I}_p$ and $\mathbf{I}_q$ of two surface points p and q is caused by the differences in their surface normals $\mathbf{n}_p$ and $\mathbf{n}_q$. Accordingly, we can say that the similarity between each pair of observation vectors $\mathbf{I}_p$ and $\mathbf{I}_q$ is closely related to the similarity between their corresponding surface normals.

In the observation vector space, all of the observed vectors lie on an intrinsic two-dimensional manifold which preserves the surface normals' geometry. Note that the surface normals are described by elevation and azimuthal angles. Surface normals are defined on a unit sphere in a three-dimensional space, so if we can find the three-dimensional embedding among the observation vector space, then we can obtain the relationships between surface normals. We use a dimensionality reduction technique to find the intrinsic three-dimensional geometry of the surface normals, which will be introduced later.

The form in (3) can represent a wide range of BRDFs, such as Lambertian, Oren-Nayar [15], Phong [2], Torrance-Sparrow [18], and isotropic Gaussian [20].

Some BRDFs are not available for this form in (3) exactly, however, from the results our method is still able to recover the object shapes reasonably with those BRDFs.

2.2 Problem with Cast Shadows

In the last section we analyzed the relationship between the appearance changes at points on the object surface and their corresponding surfaces normals. From (5) we concluded that the similarity of appearance changes is closely related to the similarity of surface normals. However, the analysis is valid only for convex objects with no cast shadows. If an object's shape is concave, the relationship between the two similarities does not hold.

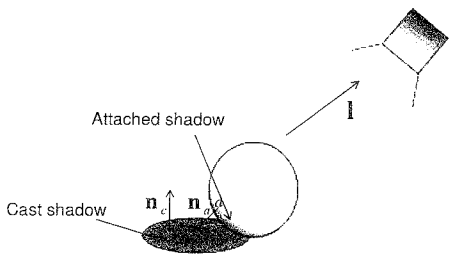

Figure 1: Cast shadows and attached shadows. $\mathbf{n}_a$, $\mathbf{n}_c$ are normal vectors of the points in attached shadow, cast shadow respectively; and $\mathbf{l}$ is the light direction vector.

Figure 1 shows a sphere with cast shadows and attached shadows under a light source. $\mathbf{n}_a$, $\mathbf{n}_c$ are normal vectors of the points in attached shadow and cast shadow respectively; and $\mathbf{l}$ is the light direction vector. We can see that for the attached shadow points on the sphere surface, which do not receive light, the directions of the surface normals are quite 'opposite' to the light direction. But the cast shadows are caused by light being interrupted by the sphere, so the directions of cast shadow pixels' surface normals are the facing toward the light. To explain this problem simply, we can assume that the surface of an object is Lambertian, and all shadow intensity values are zero. Then the observation vectors are defined as:

$$\mathbf{I}_i = \left[max\left(\rho_i {\mathbf{n}_i}^T \mathbf{l}_1, 0\right), ..., max\left(\rho_i {\mathbf{n}_i}^T \mathbf{l}_N, 0\right)\right]^T \quad (6)$$

Then for attached shadows and cast shadows:

- Attached shadows: $\mathbf{n} \cdot \mathbf{l} < 0$, so the intensity of attached shadow points will be zero.
- Cast shadows: $\mathbf{n} \cdot \mathbf{l} > 0$, so the intensity of cast shadow points should be $\rho(\mathbf{n} \cdot \mathbf{l})$.

However, both cast shadow and attached shadow points are zero. That means the attached shadows reflect the true relationship between the intensity and surface normal, but the cast shadows do not.

If the object is no longer convex, but concave and the input images are taken under numerous varying light directions, then among the input images there will be lots of cast shadows. The information about the surface normal contained in the pixel intensity is lost in areas of cast shadow, and the intensity does not obey Equation (6). This means that using similarity between observation vectors to find similarity between surface normals will cause incorrect estimation.

2.3 Shadow Interpolation Based on Similarity Measure

Considering the similarity measure between observation vectors, our key idea to solve the problem is to find similar points for each surface point, and then combine the similar vectors by interpolating new observation values in areas of cast shadow. Figure 2 shows an image of the key idea for solving the problem. For the figure we assume a video sequence where the object is Lambertian and the light moves smoothly. The large jumps are assumed to be caused by cast shadows. But actually, these assumptions are not necessary for the method.

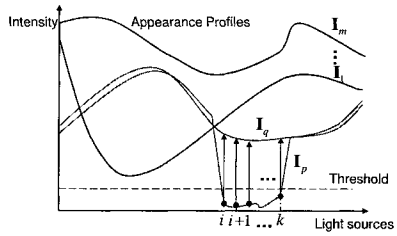

Figure 2: $\mathbf{I}_q$ is a similar observation vector to $\mathbf{I}_p$. The observation values of q can be used to recover the cast shadow area in p.

We process this algorithm by two main steps as described below:

- Calculate distance metric: We calculate a distance metric between each pair of observation vectors to find their similarity by using a set of scales between each pair of observation vectors which is robust to changes caused by surface texture.
- Select similar observation vectors and interpolate cast shadows: From the distance metric, we select similar observation vectors and create new values by interpolating in areas of cast shadow.

1. **Calculate distance metric**

Considering the set of input images of an textured object captured under n illumination directions, let I_p^i be the intensity of each surface point p ($p = 1, ..., m$) under the ith illumination. Then because of the texture on the object surface, we need to normalize all the pixels by their observation vectors as in form (2), where n is the number of illumination directions.

For different observation vectors the quantity of cast shadows is very different. So using the general normalization as shown in form (2) above, the denominator may be very different between observation vectors containing cast shadows. Because of this we calculate a scale for each pair of observation vectors instead of the normalization form in 2 that is used to calculate the distance between these two vectors.

Points p and q are two points on the object surface, and $\mathbf{I}_p$ and $\mathbf{I}_q$ are their observation vectors respectively. For each pair of p and q we choose a scale s_{pq}. If p and q have similar surface normals, they will have similar surface appearance changes, accordingly, there will be $\mathbf{I}_p \approx s_{pq}\mathbf{I}_q$. Based on this idea, we find each scale s_{pq} by minimizing the distance between each pair of observation vectors $\mathbf{I}_p$ and $\mathbf{I}_q$. Then it becomes a minimization problem to solve the following equation (7). From the solution we can obtain a set of scales for all pairs of observation vectors.

$$s_{pq} = \arg\min_s \sqrt{\sum_{\substack{i=1 \\ I_p^i, I_q^i \neq shadow}}^{n} \left(I_p^i - sI_q^i\right)^2} \qquad (7)$$

Then Euclidean distances between observation vectors can be calculated as in (8):

$$D(p,q) = \sqrt{\sum_{\substack{i=1 \\ I_p^i, I_q^i \neq shadow}}^{n} \left(I_p^i - s_{pq}I_q^i\right)^2} \qquad (8)$$

We can see that for calculating the distance between each pair of observation vectors, we only use pairs of values where neither point is in shadow. i.e. we remove all the pairs where either point is in shadow, to calculate the distance. However, because a large amount of shadow exists, and the positions of the shadows in each observation vector are different, the distances directly obtained from (8) may be not accurate in reflecting the true distance relationship of surface pixels.

In Figure 3, we show an example of two points p and q which have quite different surface orientations from each other. And the shadow positions in their observation vectors are different or even almost opposite. Let L_p and L_q be the number of non-shadow points in vector $\mathbf{I}_p$ and vector $\mathbf{I}_q$. Then L_{pq} is the overlap of L_p and L_q. We can see that the L_{pq} we can actually use to compute their distance might be very short, and it can probably result in a very small distance for p and q. Accordingly, we calculate L_p, L_q, and L_{pq}. Then if the proportion between L_{pq} and the minimum of L_p and L_q is very small (the range of this proportion is from 0 to 1), we determine that these two points are far away and have a big dissimilarity from each other. Then we can avoid the incorrect distance relationship like showed in Figure 3.

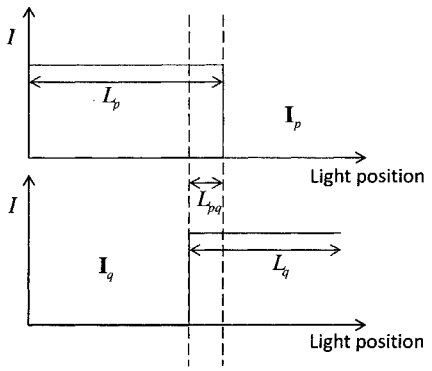

Figure 3: The problem of two very different appearance vectors exhibiting a high similarity.

2. **Select similar pixels and interpolate cast shadows**

After calculating the distances, we sort the vectors by distance from closest to farthest, and then we get a sorted list for each observation vector $\mathbf{I}_p$. We pick the top n closest pixels for each p which has shadows in its observation vector. Our aim is to find similar pixels that share a similar surface normal with p, but which are not in cast shadow at the same time. For instance, we choose $n = 5$. Then among the 5 closest vectors, in each image, we use a mean value of the observation vectors which are not in shadows to interpolate the cast shadow pixel p.

However, we cannot directly use the value of the selected vectors to interpolate in $\mathbf{I}_p$, because different pixels may have different albedo. Considering the step of scale based normalization in (7), we find a scale for each pair of observation vectors preserving these differences. We combine each scale of each selected pixel respectively with p to recover the texture of p. So the procedure is described as bellow:

$$I_p^i = mean(s_{pa}I_a^i, s_{pb}I_b^i, s_{pc}I_c^i) \qquad (9)$$

where a, b, and c are the pixels which are not in shadows among the selected similar list for point p.

We interpolate the image pixels using the algorithm above to create a new set of input images, to remove any cast shadows that are present. This reveals the correct appearance which can be used for shape reconstruction.

2.4 Finding Surface Normals from Similarities of Appearance Profiles [16]

From the set of input images with m surface points taken under n illumination directions, the m sets of n-dimensional observation vectors are obtained which represent the appearance changes. Then the surface normals and heights of these surface points are estimated as described in the following steps:

- Estimate the intrinsic three-dimensional embedding of the observation vectors using a dimensionality reduction technique.
- Recover the surface normals and object heights. The occluding boundary is used as a reference to transform the output from Isomap to the object's true surface normal geometry.

Three-dimensional embedding estimation

Dimensionality reduction techniques are widely applied in many areas. The motivation is to find the intrinsic low-dimensional structures hidden in their high-dimensional observations. Principal Component Analysis (PCA) and Multi-Dimensional Scaling (MDS) are two classical techniques that can discover the true structure of data lying on or near a linear subspace of high-dimensional input space. In other words, PCA and MDS are not capable of estimating the true structure of the data sets which contain essential nonlinear structures. Tenenbaum et al. developed a nonlinear embedding technique called isometric feature mapping (Isomap) [17]. Isomap builds on MDS [13] but seeks to preserve the intrinsic geometry of the data, as captured in the geodesic manifold distances between all pairs of data points instead of the Euclidean distance which is used in MDS and PCA. The algorithm of Isomap has three main steps:

1. Form the neighborhood relation based on the distances $d_x(i,j)$ between pairs of points i,j in the input space X, two simple methods, radus ϵ or K nearest neighbors, are used to calculate the neighborhood;
2. Estimate the geodesic distances $d_M(i,j)$ between all pairs of points by computing their shortest path distances;
3. Apply classical MDS to the matrix of graph distances $D_G = d_G(i,j)$, constructing an embedding in a d-dimensional space that best preserves the original manifold's intrinsic geometry.

Tenenbaum et al. shows that Isomap can find low-dimensional embeddings from an input high-dimensional appearance space. We use Isomap to extract the intrinsic three-dimensional embedding which preserves the surface normals' geometry from our input observation vector space.

3D shape reconstruction

The output of Isomap is a three dimensional embedding presenting the geometry of surface normals. However, this geometry only reflects a relative relationship between surface normals. Because there are distance-preserving transformations such as translations, reflections, and rotations in a three-dimensional space, the estimated three-dimensional embedding does not necessarily correspond to the true surface normals' geometry.

We use a transformation step to obtain the true surface normals' geometry. Assuming that the shape of the object is smooth and it has an occluding boundary, we use the occluding boundary points as a reference to do the transformation. A sobel filter is applied to estimate the normals of the occluding boundary points. Then based on the relative relationship between surface normals from Isomap, the true surface normals' geometry can be obtained by the occluding boundary points.

From the object's surface normals, any photometric method can be applied to estimated the height field of the object to complete the shape reconstruction. The classic technique is relaxation method which was introduced by Horn [9]. Recently, Kovesi presented a novel 3D shape reconstruction technique by using shapelets which is robust and simple to implement [12]. We use this technique to complete the shape reconstruction.

3 Experiment Results

We used sets of synthetic images to do the evaluation of the performance of our system. The images of an object with different surface reflectance properties were generated under 300 distant light sources that were randomly distributed around the object. The synthetic object is a crater shape which is obviously concave in the middle part. So, in the sets of input images there are plenty cast shadows. Versions of the shape with different reflectance properties were used: diffuse ((1) uniform Lambertian and (2) textured Lambertian), (3) specular (Phong), and (4) diffuse and specular (textured Lambertian + Phong). For each input image, the size is 64 pixels $\times$ 64 pixels and there are about 2828 object surface pixels.

In Figure 4, we show the interpolation results in a single image for four different reflection models. For each reflectance model, we have chosen one image from the results to show the recovery of the cast shadow area. On the top left is one of 300 input images, on the top right are the corresponding shadow regions (cast shadows are in black color and attached shadows are in grey color), on the bottom left is interpolation result image and on the bottom right is the corresponding ground truth image. We can see that for each reflectance model, we find the cast shadow area and remove it correctly. Also, the textures and specularities

on the object surface are recovered well. In Figure 6 we show more interpolation results for the reflectance models excluding the Lambertian model. The left column for each model, from top to bottom, shows nine of the input images, interpolated image results and their corresponding ground truth images.

After doing interpolation, we used the set of new generated images as input to do the normal estimation and 3D shape reconstruction. We compare our results with the ground truth which is shown in Figure 5. The ground truth of the shape obtained from the true surface normals is illustrated in two ways: on the left is the full three-dimensional shape and on the right is a cross section showing the shape of the middle cast shadow area. Correspondingly , in Figure 6 we show the cross section of the reconstructed result when interpolation is not used, a cross section of our result, and a 3D view of our result. We can see in Figure 6 that the middle cast shadow area in the result without interpolation is obviously incorrect compared with the ground truth in Figure 5, because the cast shadows cause the incorrect estimation of surface normals. But after removing the area of cast shadows by interpolation, the concave shape of the middle part was recovered well compared to the ground truth.

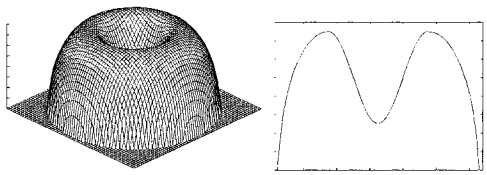

Figure 5: The ground truth shape in 3D (left) and cross section (right).

We analysed the errors in the interpolated result by calculating the difference between the obtained intensities of each pixel in each image and the corresponding ground truth. For the ground truth, we know the cast shadow positions in each image, and the true intensity values of them calculated by the known surface normals. In Table 1 below, the errors in intensity values are given. The range of the intensities is 0–255. We calculate four different error values: Mean pixel error, mean cast shadow pixel error, mean cast shadow pixel error by using data without interpolation and mean non-cast shadow pixel error. From the results in the table we can see that after interpolation, compared with the ground truth at each pixel in each image, the errors of cast shadow pixels are around 3, which is small, especially compared to the error without interpolation which is around 60. The errors of non-cast shadows are only around 0.1 which shows that our algorithm does not influence the other pixels, including attached shadows.

We also compute the error in the surface normals' direction by comparing our results and the true surface normals. The errors are listed as angles in Table 2. We calculate the errors for the four different reflectance models as we do in calculating the interpolation errors. And the errors are given as mean normal error, mean cast shadow normal error, and mean non-cast shadow normal error. From Table 2 we can see that the errors for the diffuse reflectance models are about 9° and for specular reflection models are about 5°, which are reasonably small. For cast shadow pixels, the normals' errors are similar to that of non-cast shadow pixels.

Reflection Models	Mean error	CS error	CS error (original)	non-CS error
Lambertian	0.20	3.20	75.9	0.08
Textured Lambertian	0.15	2.55	58.7	0.06
Specular	0.24	2.61	59.9	0.15
Textured Specular	0.15	1.80	41.8	0.09

Table 1: Interpolation error in intensities calculated in different parts of the object surface

Reflection Models	Mean normal error	CS error	non-CS error
Lambertian	9.1	8.4	9.2
Textured Lambertian	9.1	8.3	9.2
Specular	5.6	3.8	5.8
Textured Specular	4.9	3.0	5.1

Table 2: Normal estimation error in angles (degrees) calculated in different parts

To illustrate the contribution of using interpolation to remove the cast shadows, in Table 3 we show the error of surface normals' estimation using the sets of data without interpolation. The same with Table 2, the results are showed in degree and the parts are correspondingly to those in Table 2. Comparing the Table 3 to Table 2, it's obvious to see that the normal estimation becomes more accurate by using interpolated data, not only in the cast shadow area, but also in non-cast shadow area. It indicates that the interpolated values can accurately reveal the similarity of the appearance changes, and our method can successfully extract the intrinsic surface normals' distribution and estimate the surface normals' directions from the observation vector space for objects with different reflectance properties on the surface.

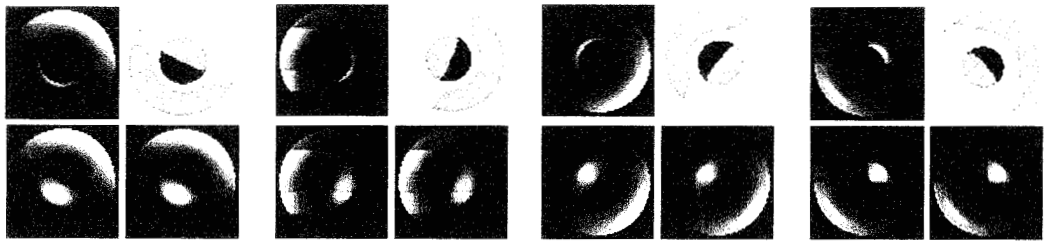

Figure 4: Interpolation results shown in single images: (a) Lambertian model, (b) textured Lambertian model, (c) specular model, and (d) textured specular model.

Reflection Models	Mean normal error	CS error	non-CS error
Lambertian	11.2	10.9	11.3
Textured Lambertian	11.1	10.7	11.2
Specular	9.4	10.5	8.9
Textured Specular	7.8	10.5	6.5

Table 3: Normal estimation error when using data without interpolation

4 Conclusion

We presented a novel photometric stereo method for recovering an object's shape directly from a set of input images of the object. It is based on the similarity of points on the appearance manifold, which contains observed appearances at points under varying illumination. We don't require any particular reflectance models or illumination properties apriori. Also, based on the similarity measure, we remove the influence caused by cast shadows in the input images. Considering the similarity measure, we find a set of similar observation vectors for each vector that contains cast shadows, and then combine the similar vectors' values to interpolate new values in areas of cast shadows. We reveal that for a pair of surface points, the similarity of their observed appearance on the object's surface under varying illumination is closely related to the similarity of their corresponding surface normals. Then we use a dimensionality reduction technique to estimate the intrinsic surface normals' geometry.

References

[1] P. N. Belhumeru, D. J. Kriegman, and A. L. Yuille. The bas-relief ambiguity. In *Int'l J. Computer Vision*, volume 35(1), pages 33–44, 1999.

[2] P. T. Bui. Illumination for computer generated pictures. *Commun. ACM*, 18(6):311–317, 1975.

[3] M. K. Chandraker, F. K., and D. J. Kriegman. Reflections on the Generalized Bas-Relief Ambiguity. *CVPR*, 1:788–795, 2005. ISSN 1063-6919.

[4] E. N. Coleman and R. Jain. Obtaining 3-dimensional shape of textured and specular surfaces using four-source photometry. *Computer Graphics and Image Processing*, 18:309–328, 1982.

[5] K. J. Dana, B. Ginneken, S. K. Nayar, and J. J. Koenderink. Reflectance and texture of real-world surfaces. *ACM Trans. Graph.*, 18(1):1–34, 1999. ISSN 0730-0301.

[6] A. S. Georghiades. Incorporating the Torrance and Sparrow Model of Reflectance in Uncalibrated Photometric Stereo. In *ICCV '03: Proceedings of the Ninth IEEE International Conference on Computer Vision*, pages 816–825, Washington, DC, USA, 2003. IEEE Computer Society.

[7] H. Hayakawa. Photometric stereo under a light source with arbitrary motion. *Journal of the Optical Society of America. A, Optics, image science, and vision*, 11:3079–3089, 1994.

[8] A. Hertzmann and S. M. Seitz. Example-based photometric stereo: shape reconstruction with general, varying BRDFs. *IEEE Trans. Pattern Anal. Mach. Intell.*, 27(8):1254–1264, 2005.

[9] B. Horn. *Robot Vision*. MIT Press, Cambridge, MA, 1986.

[10] K. Ikeuchi. Determining Surface Orientations of Specular Surfaces by Using the Photometric Stereo Method . *IEEE TRANS. PATTERN ANALY. AND MACH. INTELLIG.*, PAMI-3(6): 661–669, 1981.

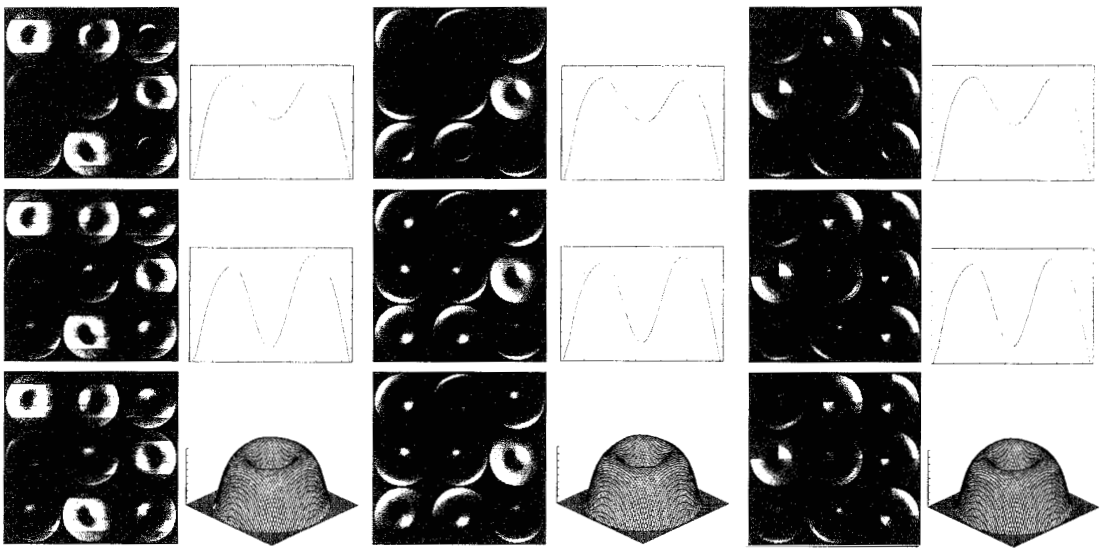

(a) (b) (c)

Figure 6: Interpolation and 3D shape results for each model: (a) textured Lambertian model, (b) specular model, and (c) textured specular model. Images on left: input images, interpolated results, ground truth. Shapes on right: cross section of reconstruction without interpolation (top), with interpolation (middle) and and 3D shape (bottom). Compare the reconstructed shapes (right and right-bottom) with the ground truth shape (Figure 5)

[11] S. J. Koppal and S. G. Narasimhan. Appearance Clustering: A Novel Approach to Scene Analysis. In *IEEE Conference on Computer Vision and Pattern Recognition*, June 2006.

[12] P. Kovesi. Shapelets correlated with surface normals produce surfaces. In *ICCV 2005. Tenth IEEE International Conference on Computer Vision,2005*, volume 2, pages 994–1001, 2005.

[13] K. V. Mardia, J. T. Kent, and J. M. Bibby. *Multivariate analysis.* Probability and Mathematical Statistics, London: Academic Press, 1979, 1979.

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

[15] M. Oren and S. K. Nayar. Generalization of the Lambertian model and implications for machine vision. *Int. J. Comput. Vision*, 14(3):227–251, 1995. ISSN 0920-5691.

[16] I. Sato, T. Okabe, Q. Yu, and Y. Sato. Shape Reconstruction Based on Similarity in Radiance Changes under Varying Illumination. In *Proc. IEEE Intfl Conf. Computer Vision (ICCV 2007)*, 2007.

[17] J. B. Tenenbaum, V. Silva, and J. C. Langford. A Global Geometric Framework for Nonlinear Dimensionality Reduction. In *Science*, volume 290, pages 2319–2323, 2000.

[18] K. E. Torrance and E. M. Sparrow. Theory for off-specular reflection from roughened surfaces. *Journal of the Optical Society of America (1917-1983)*, 57:1105–1114, sep 1967.

[19] E. Trucco and A. Verri. *Introductory Techniques for 3-D Computer Vision.* Prentice Hall, 1988.

[20] G. J. Ward. Measuring and modeling anisotropic reflection. *SIGGRAPH Comput. Graph.*, 26(2): 265–272, 1992. ISSN 0097-8930.

[21] R. J. Woodham. Analysing images of curved surfaces. *Artificial Intelligence*, 17(2):117–140, 1981.

[22] R. Zhang, P. Tsai, J. E. Cryer, and M. Shah. Shape from Shading: A Survey. *IEEE Trans. Pattern Anal. Mach. Intell.*, 21(8):690–706, 1999. ISSN 0162-8828.