The 4-Source Photometric Stereo Under General Unknown Lighting

Chia-Ping Chen[1,2] and Chu-Song Chen[1,3]

[1] Institute of Information Science,
Academia Sinica, Taipei, Taiwan
[2] Department of Computer Science and Information Engineering,
National Taiwan University, Taipei, Taiwan
[3] Graduate Institute of Networking and Multimedia,
National Taiwan University, Taipei, Taiwan
{cpchen, song}@iis.sinica.edu.tw

Abstract. Many previous works on photometric stereo have shown how to recover the shape and reflectance properties of an object using multiple images taken under a fixed viewpoint and variable lighting conditions. However, most of them only dealt with a single point light source in each image. In this paper, we show how to perform photometric stereo with four images which are taken under distant but general lighting conditions. Our method is based on the representation that uses low-order spherical harmonics for Lambertian objects. Attached shadows are considered in this representation. We show that the lighting conditions can be estimated regardless of object shape and reflectance properties. The estimated illumination conditions can then help to recover the shape and reflectance properties.

1 Introduction

Photometric stereo methods recover the shape and reflectance properties of an object using multiple images under varying lighting conditions but fixed viewpoint. Most works on this problem assumed that lighting comes from a single source, generally a point source or a controlled, diffused source of light. Woodham [1] first introduced photometric stereo for Lambertian surfaces assuming known albedos and known lighting directions. The method was based on the use of the so-called reflectance maps in the form of look-up tables. Three images were used to solve the reflectance equation for recovering surface gradients and albedos of a Lambertian surface. Coleman and Jain [2] used four images to detect and exclude highlight pixels. They used four combinations of three light sources to compute four albedo values at each pixel. Presence of specular highlight will make the computed albedos different, indicating that some measurement should be excluded. Barskey and Petrou [3] showed that the method in [2] is still problematic if shadows are present, and generalized it to handle color images.

Moses [4] and Shashua [5] have pointed out that one can only recover the scaled surface normals up to an unknown linear transformation when each image

of the target object is lit by a single point source with unknown intensity and direction. Hayakawa [6] used [4,5]'s result in a factorization framework to handle many images. These results assumed there is no shadow. In [7,8,9], it has been shown that integrability reduces the ambiguity to a generalized bas-relief transformation, which allows the recovery of a surface up to a stretch and shear in the z-direction.

By considering non-Lambertian surfaces, Tagare and deFigueriredo [10] developed a theory of photometric stereo for the class of m-lobed reflectance map. Kay and Caelly [11] continued their work and investigated the problem from a practical point of view. They applied nonlinear regression to a larger number of input images. Solomon and Ikeuchi [12] extended the method in [2] by separating the object into different areas. The Torrance-Sparrow model was then used to compute the surface roughness. Nayar et al. [13] used a hybrid reflectance model, and recovered not only the surface gradients but also parameters of the reflectance model. In these approaches, the models used are usually somewhat complex, and more parameters need to be estimated.

Hertzmann and Seitz [14] used a reference object and presented an approach to compute surface orientations and reflectance properties. They made use of orientation consistency to establish correspondence between the unknown object and a known reference object. In many cases, however, obtaining a reference object for correspondence can be very difficult. Goldman et al. [15] further extended this method so that the reference object is no longer needed but assumed that objects are composed of a small number of fundamental materials.

Previous work on this problem shows a progression towards lighting conditions that are less constrained, but most of them still focused on recovering structure based on the assumption of a single point source in each image. For complicated lighting environments, Basri et al. [16] and Ramamoorthi et al. [17] have provided a new way to describe the effect of general lighting on a Lambertian object. Their results showed that only the low frequency components of lighting have a significant effect on the reflectance function of a Lambertian object. These components are represented as low-order spherical harmonics. They showed that the set of images produced by a convex Lambertian object under arbitrary lighting can be well approximated by a low dimensional linear set of images. This set is 4D for a first-order approximation, 9D for a second-order approximation.

Basri and Jacobs [18] used this representation to handle the photometric stereo problem under general lighting. They assumed that the zero- and first-order harmonics, which correspond to the albedos and surface normals scaled by albedos, will show up in the space spanned by the principal components obtained by performing SVD on input images. Their method can reconstruct object shape well when a large number of images are available.

In this paper, we also consider images produced by general lighting conditions that are not known ahead of time, and we require only four images of the target object. The starting point of our method is the spherical harmonic

representation. We first estimate the lighting condition without knowing the object shape with the 4D approximation. Then we couple our refinement techniques for surface normal, albedo, and lighting condition into an iterative process, where improved shape results in improved lighting and albedo estimation and vice versa. We confirm experimentally that these optimization procedures produce good results in spite of the approximations made by the low-order spherical harmonic representations. Existing techniques can then be used to translate these normals into an integrable surface, if desired. We also present some experiments to illustrate the potential of our method.

We formulate our problem and show the spherical harmonic representation used by our method in Section 2. Lighting estimation method without knowing object shape is described in Section 3. The shape and reflectance reconstruction process is shown in Section 4, where we also propose an iterative optimization algorithm to obtain a more robust solution. Section 5 shows experimental results of both synthetic and real images, while conclusions and future work are made in Section 6.

2   Problem Formulation

The inputs of our method are four images of a static object taken at a fixed pose but under different illuminations. The lighting conditions, shape, and reflectance properties are all unknown. From these inputs, we seek to estimate the lighting conditions, and to reconstruct the shape and reflectance properties.

We assume that the surface of the target object has Lambertian reflectance. The only parameter of this model is the albedo of each point on the object, which describes the fraction of the light reflected. We also assume that this object is illuminated by distant light sources, so that the directions and intensities of light sources are the same for all points of this object. We do not model the effects of cast shadows and interreflections.

The output of our method are albedos, surface normals of each pixel, and the lighting conditions of four input images. We can then reconstruct the surface by integrating the normal field. We can also render novel images under new lighting conditions or new viewpoints.

2.1   Modeling Reflection

We use the symbol I_i to represent the intensity of a certain pixel in input image i, $i = 1 \ldots 4$. With distant light source assumption, the intensities of a Lambertian object under general lighting conditions can be represented as follows:

$$I_i = \rho \int L_i(l) \max(l \cdot N, 0) \, dl, \tag{1}$$

where ρ and N are the albedo and surface normal of this pixel, l is the unit vector indicating the direction of incoming light, and $L_i(l)$ is the radiance intensity from direction l in image i. The integral is over all possible lighting directions.

In [16, 17], the authors viewed this process of light reflection as a convolution, where the incident illumination signal is filtered by the reflective properties of the surface, which is the Lambertian kernel in our case. They also showed that the Lambertian kernel acts like a low-pass filter, preserving only the lowest frequency components of the lighting. According to their results, the effects of general lighting on a Lambertian object can be represented by spherical harmonics:

$$I_i = \rho \sum_{n=0}^{\infty} \sum_{m=-n}^{n} \sqrt{\frac{4\pi}{2n+1}} k_n l_{nm} Y_{nm}(N), \tag{2}$$

where Y_{nm} are the surface spherical harmonics, k_n and l_{nm} are the coefficients of harmonic expansions of the Lambertian kernel and lighting, respectively. It has been proved in [16] that for any distant and isotropic lighting, at least 98% of the resulting function can be captured by the second-order spherical harmonic approximation. A first-order approximation captures at least 75% of the reflectance. These bounds are not tight, and in fact many common lighting conditions yield significantly better approximations. For example, under a point source illumination the first- and second-order harmonics approximate the reflectance function to 87.5% and 99.22% respectively.

Then, we can relate the intensity quadruple of the same pixel in four input images to the spherical harmonics with the second-order approximation:

$$I = \begin{bmatrix} I_1 \\ I_2 \\ I_3 \\ I_4 \end{bmatrix} = \rho L_{4\times9} H_{9\times1} = \rho L_{4\times9} \begin{bmatrix} 1 \\ N_x \\ N_y \\ N_z \\ 3N_z^2 - 1 \\ N_x N_y \\ N_x N_z \\ N_y N_z \\ N_x^2 - N_y^2 \end{bmatrix}, \tag{3}$$

where I is the stack of intensity quadruple, each row of $L_{4\times9}$ is the lighting configuration of each image, and $H_{9\times1}$ is the spherical harmonics which can be decided analytically when the surface normal $N = [N_x, N_y, N_z]^T$ is known. Note that we omit additional constant factors since they do not change the space spanned by these bases.

3 Lighting Estimation

In an uncalibrated photometric stereo problem, the lighting condition L, albedo ρ, and surface normal N are all unknown, making it highly unconstrained. However, with first-order approximation of spherical harmonics, we show in this section that, regardless of object shape and reflectance, the lighting conditions are constrained and can be estimated up to a subgroup of the 4×4 linear transformation, called *Lorentz* transformations (also in [18]).

When first-order approximation of spherical harmonics is adopted, the intensity quadruple can be rewritten from (3) as follows:

$$I = \rho L_{4\times4} H_{4\times1} = \rho L_{4\times4} \begin{bmatrix} 1 & N_x & N_y & N_z \end{bmatrix}^T . \tag{4}$$

When L is known, the albedo and surface normal can be recovered simply by inverting this equation as:

$$\rho H = \begin{bmatrix} \rho & \rho N_x & \rho N_y & \rho N_z \end{bmatrix}^T = L^{-1} I . \tag{5}$$

However, in a photometric stereo problem under general lighting, L is usually unknown. In the following, we show how to estimate L with this first-order approximation.

Since N is the unit surface normal, $N^T N = 1$. Let $J = \text{diag}\{-1, 1, 1, 1\}$, it can be easily verified that

$$(\rho H)^T J (\rho H) = \left(I^T L^{-T} \right) J \left(L^{-1} I \right) = I^T B I = 0, \tag{6}$$

where $B = L^{-T} J L^{-1}$ is a 4×4 symmetric matrix. Equation (6) indicates that intensity quadruples are constrained to lie on a quadratic surface regardless of what the albedos and surface normals are:

$$\begin{aligned}
0 = {} & B_{11} I_1^2 + B_{22} I_2^2 + B_{33} I_3^2 + B_{44} I_4^2 + \\
& 2B_{12} I_1 I_2 + 2B_{13} I_1 I_3 + 2B_{14} I_1 I_4 + \\
& 2B_{23} I_2 I_3 + 2B_{24} I_2 I_4 + 2B_{34} I_3 I_4 .
\end{aligned} \tag{7}$$

This equation has only ten unknowns, which follows from the fact that the matrix B is symmetric, and we can list one equation per pixel.

The ten unknowns of B can be determined even when the lighting conditions L is unknown and even when there is no object point where the albedo ρ and surface normal N are given. All that is required is having sufficient measured intensity quadruples. A standard linear least-squares method can be used to estimate these ten unknowns. This least-squares estimation should be robust since the number of image pixels are usually large and each measured intensity quadruple contributes useful information. However, because this linear system is homogeneous, B can only be solved up to an unknown scale. This ambiguity comes from (4). For any scalar $s > 0$,

$$I = (\frac{\rho}{s})(sL)H, \tag{8}$$

which means that we cannot distinguish between brighter surfaces lit by a dimmer illumination or darker surfaces lit by a brighter illumination. In the remainder of this paper, we will therefore ignore this scale ambiguity.

Thus, empirical measurements determine the matrix B. The constraint that B imposes on the lighting condition L can be interpreted when expressed in terms of B^{-1}:

$$B^{-1} = (L^{-T} J L^{-1})^{-1} = L J L^T, \tag{9}$$

where $J^{-1} = J$. Because B^{-1} is also symmetric, B^{-1} can be factorized as follows by Symmetric Schur Decomposition:

$$B^{-1} = QJ\Lambda Q^T, \tag{10}$$

where Q contains the orthogonal eigenvectors and $\Lambda = \text{diag}\{|\lambda_1|, \lambda_2, \lambda_3, \lambda_4\}$ is the eigenvalue matrix. Without lost of generality, we order Λ and Q so that the negative eigenvalue λ_1 is the first, and we move the negative sign into J to ensure that every element of Λ is positive. Then we can decide that $L = \sqrt{\Lambda}Q$. If there is only one positive eigenvalue we reverse the sign of B^{-1}.

When there is significant noise, or when the assumptions do not strictly hold (Lambertian surfaces, distant light source, etc.), the eigenvalues of B may not have the proper signs. In that case we resort to an iterative optimization to find L that minimizes the Frobenous norm $||B^{-1} - LJL^T||$.

At this point we have recovered a valid L. However, there is still an unsolved ambiguity. For any matrix C that satisfies $CJC^T = J$, $B^{-1} = (LC)J(LC)^T = LJL^T$. This set of transformations forms the Lorentz group [19]. Because the symmetric quadratic form $CJC^T = J$ gives ten quadratic equations in the 16 unknown components of C, a Lorentz transformation has six degrees of freedom. We can resolve this ambiguity, for example, if we know the surface normals and albedos of two points, or we can remove this ambiguity by enforcing surface integrability as in [9].

In sum, the lighting condition can be estimated using the first-order approximation of spherical harmonics, regardless of the albedo and surface normal. Although not very accurate, this first-order approximation suffices for good initial lighting conditions. It helps us reconstruct shape and reflectance properties of the target object and can be refined afterwards.

4   Shape and Reflectance Reconstruction

In the previous section we show how to estimate the initial lighting conditions without knowing the albedos and surface normals of the target object. With this initial estimation $L_{4\times4}$, we can easily recover the albedo ρ and surface normal N by (5). However, the results could not be accurate enough because only first-order spherical harmonics are used. In this section, we apply the estimated lighting conditions and surface normals from the first-order approximation as initial, and iteratively refine lighting conditions, albedos, and surface normals by incorporating second-order spherical harmonics into an optimization process.

4.1   Refine Lighting Estimation

When the surface normals are reconstructed to some extent, they can really help the re-estimation of lighting conditions even if they are not very accurate. In the following, we show how to refine the lighting conditions with the second-order spherical harmonics when surface normals are available.

Given a surface normal N and two image intensities I_s and I_t, its albedo can be estimated using Equation 3 as:

$$\rho_s = \frac{I_s}{L_s H_{9\times 1}} \quad \text{or as} \quad \rho_t = \frac{I_t}{L_t H_{9\times 1}}, \tag{11}$$

where L_s and L_t are the sth and tth rows of $L_{4\times 9}$, respectively. By expanding $\rho_s = \rho_t$, we obtain:

$$\begin{bmatrix} -I_s & I_t \end{bmatrix}_{1\times 2} \begin{bmatrix} L_t \\ L_s \end{bmatrix}_{2\times 9} H_{9\times 1} = 0, \tag{12}$$

where $H_{9\times 1}$ can be derived analytically given surface normal N, and I_s, I_t are the observed intensities in two input images. This forms a linear equation where only the lighting conditions L_s and L_t are unknown. Four observed intensities in the four input images yield three independent equations for every pixel. With P pixels, we have $3P$ equations to solve the 36 unknown coefficients of $L_{4\times 9}$. Since usually $3P \gg 36$, we can solve $L_{4\times 9}$ effectively using least-squares methods with the help of known surface normals.

4.2   Refine Reflectance

Once we have a better estimation of the lighting conditions, we can improve the albedo estimation with known surface normals. This can be easily done with the following equation:

$$\rho L_{4\times 9} H_{9\times 1} = \rho \begin{bmatrix} L_1 H \\ L_2 H \\ L_3 H \\ L_4 H \end{bmatrix} = \begin{bmatrix} I_1 \\ I_2 \\ I_3 \\ I_4 \end{bmatrix}. \tag{13}$$

The optimal albedo which has the least square error can be derived as:

$$\rho = \sum_{i=1}^{4} (L_i H I_i) / \sum_{s=1}^{4} (L_i H)^2 . \tag{14}$$

4.3   Refine Surface Normals

With refined lighting conditions and albedos, we can further improve the estimation of surface normals. For each pixel, we seek to find the best unit surface normal N that minimizes the following energy function:

$$E(N) = ||\rho L_{4\times 9} H_{9\times 1} - I||^2 = \sum_{i=1}^{4} (\rho L_i H_{9\times 1} - I_i)^2 . \tag{15}$$

Instead of using conventional convex minimization over the continuous surface normal field and enforcing unity constraint explicitly on N, we adopt discrete optimization over a fixed number of available unit normals. We seek an optimal labeling which minimizes (15)

$$v^* = \arg\min_v E(\hat{N}_v), \tag{16}$$

where $\{\hat{N}_v | v = 1, 2, \ldots V\}$ is the set of available unit directions. Then the surface normal N is refined as $\hat{N}_{v^*}$. To produce uniform sampling of normal directions, we start with an icosahedron, and perform subdivision on each face 5 times recursively [20]. Totally, 10168 points are uniformly sampled on the unit sphere.

One problem of the above estimation is that the surface normals are refined one by one independently. The generated normal field is not guaranteed to be consistent with a 3D surface, thus we enforce integrability by solving a Poisson equation to obtain a least-squares surface reconstruction, and subsequently the normals are recomputed by differentiate this surface. To compute a 3D surface from the estimated surface orientations, given the normal field $N(x, y)$, we solve the height field $z(x, y)$ that minimizes

$$\Psi(z) = \sum_{x,y} \left(\frac{\delta z(x, y)}{\delta x} + \frac{N_x(x, y)}{N_z(x, y)} \right)^2 + \left(\frac{\delta z(x, y)}{\delta y} + \frac{N_y(x, y)}{N_z(x, y)} \right)^2. \qquad (17)$$

This amounts to integrating the normal field. The minimization gives rise to a large but sparse system of linear equations. The normals are then recomputed from this surface approximation. This step can be viewed as projecting the possibly non-integrable normal field into the subspace of feasible normal fields.

4.4   Iterative Algorithm

We couple our lighting, albedo, and shape estimation techniques described above into an iterative process, where improved shape estimation leads to improved lighting and albedo estimation and vice versa.

Initialization. Get an initial estimation of lighting condition $L_{4\times 4}$ as described in Section 2, then recover the surface normal N by (5).

Step 1. Refine lighting conditions $L_{4\times 9}$ with the help of the estimated surface normal N by solving (12).

Step 2. Refine albedos with currently estimated surface normals and lighting conditions according to (14).

Step 3. Search the optimal surface normal for each pixel by (16). Integrate the normal field to get an approximated surface according to (17) and then recompute the surface normals again.

Termination. Step 1-3 are iterated until the estimated components no longer change or the specified maximum iteration number is reached.

Each step in our iterative algorithm is guaranteed to monotonically decrease the reconstruction error $|\rho LH - I|$ between reconstructed images and input images, except the enforcement of integrability. However, this integration step only makes little changes to the surface in our experience and the reconstruction error does not change much after a few iterations. Therefore, our optimization algorithm is likely to find a solution near a local optimum.

5   Experiments

We now present experiments to evaluate our method. Because we use the second-order approximation of spherical harmonics, certain built-in errors persist even in the ideal case. So we first describe experiments on synthetic data to verify some basic properties of our method.

We generate four images of a diffuse unit sphere with uniform albedo under four different lighting conditions. We then use our method to recover the surface normals of this sphere. Since we have ground truth, we can resolve the Lorentz ambiguity by some known surface normals. Figure 1 shows the four input images of this unit sphere. The recovered surface normals are shown in Figure 2, where positive values are shown in green, and negative values are shown in red. The mean error between the recovered and real surface normals is 0.12 degrees. This experiment tells us that our method will produce good results in ideal situation, that is, in the absence of sensing error, cast shadow, specularity and any other source of noise.

We have also run our method on real images from Yale Face Database B [21]. To reduce the effect that may be caused by cast shadows, only frontally illuminated images are used. Each image is lit by a single point source; thus we average pairs of images to simulate complicated lighting conditions. This time we enforce the integrability constraint to resolve the Lorentz ambiguity. Figure 3(a) shows the four input images. The recovered albedos and surface normals are shown in Figure 4. Note that there are some cast shadows, noises, unreliable pixels that have been saturated in the four input images. But the results of our method are still quite satisfactory. Noticeable artifacts occur in the eyes which exhibit highly specular reflection, in the side of the nose where there are cast shadows that we do not model, as well as in the outline of eyebrow because of alignment error since this target object is not really static.
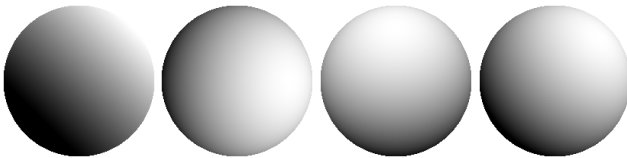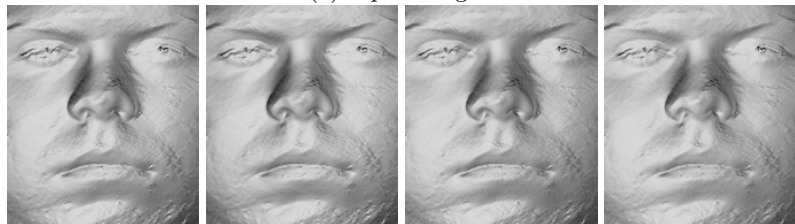

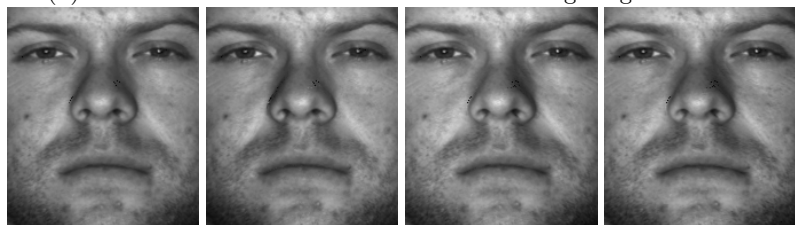**Fig. 1.** Four input images of an unit sphere

Fig. 2. Recovered surface normals of the unit sphere. Positive values are shown in green, and negative values are shown in red.
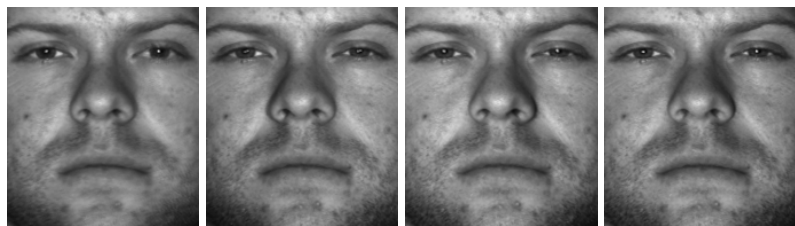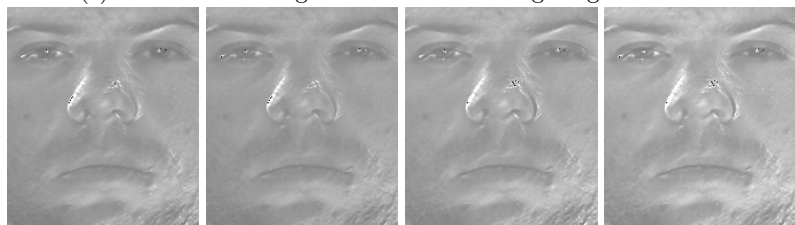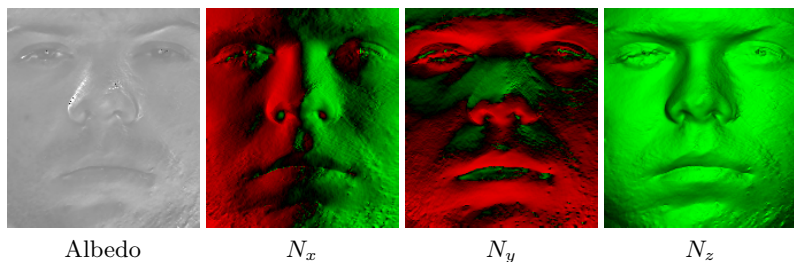
(a) Input images.

(b) Surface normals rendered under estimated lighting conditions.

(c) Re-rendered images under estimated lighting conditions.

(d) De-lighted images.

Fig. 3. Experiments on real images

Albedo N_x N_y N_z

Fig. 4. Recovered albedo and surface normals. Positive values are shown in green, and negative values are shown in red.

Fig. 5. Re-rendered images under novel lighting and viewpoints

Figure 3(b) shows the recovered surface normals rendered with uniform albedos under the estimated lighting condition of each input image. The re-rendered images with the estimated albedos reproduce the four input images quite well, as shown in Figure 3(c). The de-lighted images shown in Figure 3(d) exhibit much less structure than original ones, because shading effects that are accounted for changes in shape have been greatly attenuated. We also render novel views under different lighting conditions and viewpoints, as shown in Figure 5. These novels views are quite realistic, which indicates the usefulness of our method for many applications, such as recognition under novel lighting conditions and viewpoints.

6   Conclusions and Future Work

In this paper, we proposed a method that handles the 4-source photometric stereo problem under general unknown lighting conditions. We showed that the lighting conditions can be estimated regardless of the albedos and object shape when first-order approximation of spherical harmonics is adopted. Then we used this initial estimation in an iterative process, where second-order spherical harmonics were incorporated. During the optimization process, improved shape estimation leads to improved lighting and albedo estimation and vice versa. The effects of attached shadows are considered, which benefits from the spherical harmonic representation. The experimental results showed that our method can derive quite satisfactory results even when only four input images are used. Future work will be focused on introducing specularity models, prediction of cast shadows, and utilization of color information.

References

1. R.J. Woodham. Photometric Method for Determining Surface Orientation from Multiple Images. *Optical Engineering*, 19(1):139-144, 1980.
2. E.N. Coleman, Jr. and R. Jain. Obtaining 3-dimensional shape of textured and specular surfaces using four-source photometry. *CGIP*, 18(4):309-328, April 1982.
3. S. Barsky and M. Petrou. The 4-source photometric stereo technique for three-dimensional surfaces in the presence of highlights and shadows. *PAMI*, 25(10):1239-1252, October 2003.
4. Y. Moses. Face recognition: generalization to novel images. Ph.D. Thesis, Weizmann Institute of Science, 1993.
5. A. Shashua, On Photometric Issues in 3D Visual Recognition from a Single 2D Image. *IJCV* 21(1-2):99-122, 1997.
6. H. Hayakawa. Photometric stereo under a light source with arbitrary motion. *JOSA*, 11(11):3079-3089, 1994.
7. P.N. Belhumeur, D.J. Kriegman, A.L. Yuille. The Bas-Relief Ambiguity. *IJCV*, 35(1):33-44, 1999.
8. J. Fan, L.B. Wolff. Surface Curvature and Shape Reconstruction from Unknown Multiple Illumination and Integrability. *CVIU* 65(2):347-359, 1997.
9. A. Yuille, D. Snow, R. Epstein, P. Belhumeur. Determining Generative Models of Objects Under Varying Illumination: Shape and Albedo from Multiple Images Using SVD and Integrability. *IJCV*, 35(3):203-222, 1999.
10. H.D. Tagare and R.J.P. deFigueiredo. A theory of photometric stereo for a class of diffuse non-lambertian surfaces. *PAMI*, 13(2):133-152, February 1991.
11. G. Kay and T. Caelly. Estimating the parameters of an illumination model using photometric stereo. *GMIP*, 57(5):365-388, 1995.
12. F. Solomon and K. Ikeuchi. Extracting the shape and roughness of specular lobe objects using four light photometric stereo. *PAMI*, 18(4):449-454, April 1996.
13. S.K. Nayar, K. Ikeuchi, and T. Kanade. Determining shape and reflectance of hybrid surfaces by photometric sampling. *IEEE Trans. on Robotics and Automation*, 6(4):418-431, 1990.
14. A. Hertzmann and S.M. Seitz. Shape and materials by example: a photometric stereo approach. *CVPR*, 2003.
15. D.B. Goldman, B. Curless, A. Hertzmann, and S.M. Seitz. Shape and Spatially-Varying BRDFs From Photometric Stereo. *ICCV*, 2005.
16. R. Basri and D. Jacobs. Lambertian reflectance and linear subspaces. *PAMI*, 25(2):218-233, 2003.
17. R. Ramamoorthi and P. Hanrahan, On the relationship between radiance and irradiance: determining the illumination from images of a convex Lambertian object. *Journal of the Optical Society of America A*, 18(10):2448-2459, 2001.
18. R. Basri, D. Jacobs. Photometric Stereo with General Unknown Lighting. *CVPR*, 2001.
19. K. Kanatani. Geometric Computation for Machine Vision. *Oxford University Press*, 1993.
20. D.H. Ballard and C.M. Brown. Computer Vision. In *Prentice Hall*, 1982.
21. A.S. Georghiades and P.N. Belhumeur. From Few to many: Illumination cone models for face recognition under variable lighting and pose. *PAMI*, 23(6):643-660, 2001.