

Efficient Photometric Stereo on Glossy Surfaces with Wide Specular Lobes*

Hin-Shun Chung Jiaya Jia

Department of Computer Science and Engineering
The Chinese University of Hong Kong

{hschung, leojia}@cse.cuhk.edu.hk

Abstract

This paper presents a new photometric stereo method aiming to efficiently estimate BRDF and reconstruct glossy surfaces. Rough specular surfaces exhibit wide specular lobes under different lightings. They are ubiquitous and usually bring difficulties to both specular pixel removal and surface normal recovery. In our approach, we do not apply unreliable highlight separation and specularity estimation. Instead, an important visual cue, i.e. the cast shadow silhouette of the object, is employed to optimally recover global BRDF parameters. These parameter estimates are then taken into a reflectance model for robustly computing the surface normals and other local parameters using an iterative optimization. Within the unified framework, our method can also be used to reconstruct object surfaces assembled with multiple materials.

1. Introduction

Shape reconstruction from different lightings has long been a fundamental problem in machine vision. Early approaches [25] used photometric stereo to estimate surface normal and diffuse albedo assuming Lambertian surfaces. Recent methods [15, 20, 2, 9] were proposed to also deal with non-Lambertian surfaces. In order to tackle complex reflections, most previous approaches either require that the highlight regions in input images are relatively small and can be easily removed, or assume that the highlights in different input images do not overlap such that Lambertian rule can be applied to recovering the surface normals. To estimate the normal of one pixel, traditional photometric stereo uses at least three frames in which the corresponding pixels contain no specular components.

Specular component removal is not easy for rough specular surfaces. The micro-facets on these surfaces reflect light in different directions, producing wide specular lobes




Figure 1. The wide specular lobes on rough specular surfaces under directional light.

with highly blurred boundary under directional light. Two examples are shown in Figure 1. Moreover, it is usually not easy to determine whether the reflected light from the surface contains any specular components. An imprecise specularity segmentation makes the surface normal recovery error-prone. In order to remove the wide specular lobes on surfaces, Mallick *et al.* [13] proposed a data-dependent specular separation method using color information. It requires that the diffuse color of the surface and the color of the incident light are different. Georgiades [8] used a reflectance model including specular component to directly estimate surface normals and BRDF parameters. This method has to solve an optimization problem containing a large number of unknowns.

In this paper, we present an efficient photometric stereo method to robustly estimate both surface normals and BRDF parameters. Our method does not require specularity separation in order to recover surface normals, thereby is capable of automatically and precisely reconstructing convex object surfaces with different level of roughness from a small number of images. The object surface is allowed to have complex textures and colors without influencing the computation accuracy.

We adopt a two-step optimization to estimate the normals and all reflectance parameters. In the first step, we use the visual information from shadow to estimate the normals

*The work described in this paper was supported by a grant from the Research Grants Council of the Hong Kong Special Administrative Region, China (Project No. 412307) and a grant from SHIAE (No. 8115016).

of some surface points. These recovered normals play a key role in estimating the global reflectance parameters. In the second step, with the reduced number of unknowns, a robust iterative optimization process is applied to estimating the surface normals and diffuse albedo for all surface points.

In this unified framework, our method can also be applied to dealing with surfaces assembled with several different materials. The material segmentation, parameter estimation, and normal computation can be automatically achieved in our approach.

2. Related Work

Most related work on non-Lambertian surface reconstruction uses various bi-directional reflection distribution function (BRDF) models. The pioneer work includes recovering the surface normals using a Torrance-Sparrow model [14]. Tagare and deFigueiredo [21] estimated both the shape and reflectance map for diffuse surfaces which consist of foreshadow lobes, normal lobes, and backscatter lobes. In our method, we estimate the shape and reflectance of a surface containing both specular and diffuse components.

Coleman and Jain [5] used four-source configurations to detect highlight. In [20], the surface roughness was recovered using the similar four-light setup. In [2], Barsky and Petrou demonstrated that color can facilitate the detection of specular pixels with the assumption of no specular reflection in at least three images for each pixel.

Using a reference object, Hertzmann [11] proposed an example-based surface reconstruction method with arbitrary BRDFs. It employs the orientation consistency to detect correspondence between the target object and the reference one. Recently, Wu and Tang [26] introduced a photometric stereo method using dense image input to estimate surface normals using graph cuts. Complex surfaces with discontinuities, highlight, and transparency can be recovered. In [9], spatially-variant BRDFs were assumed and the object surface was modeled by a linear combination of two fundamental materials. An iterative optimization on BRDF parameters and surface normals is applied. This method assumes that the specular lobe is narrow so that the highlight pixels can be removed by intensity threshold.

In order to robustly reconstruct surface, the specular components, i.e., the highlight, are usually required to be removed such that traditional photometric stereo method on Lambertian surface can be applied. Most highlight removal methods are based on dichromatic reflection model [19]. Klinker *et al.* [12] separated specular and diffuse components using the RGB color space. Sato *et al.* [18] proposed a method based on the color difference between the diffuse and specular components to estimate reflectance. Mallick *et al.* [13] introduced a data-dependent SUV color space obtained by rotating the RGB color space such that R channel is aligned with the color of the incident light. Then the

other two channels contain only diffuse color component. Since these approaches only consider colors, they do not work well on textured objects where the diffuse color of a surface and the color of a light are similar.

The highlight separation method proposed in [17] works on inhomogeneous surfaces using additional light sources. This method is limited by the problem of color similarity between the light source and object surface. In [23, 22], highlight components were removed using single image information. Most of these methods assume that the diffuse color does not change too much inside and outside the highlight region or the specular color is different from that on surface.

Most of the above image-based highlight removal methods make assumptions on colors, shape, or brightness. In this paper, we consider the situation that the specular surface becomes rough and the specular lobes have wide and severely blurred boundary. This situation brings difficulty for previous methods to determine the specular components in the reflected light. Consequently, the highlight removal methods easily fail on these objects especially when the surface has complex textures.

Georghiadis [7] showed that the specularities of surfaces can be used in uncalibrated photometric stereo. In [8], he proposed to recover surface normals and reflectance using iterative optimization over a very large set of variables, including reflectance parameters, light directions, and surface normals. Because there exists a large search space, a good initialization of all unknowns are generally required. Besides, the efficiency and optimality are also issues for that optimization process. In [4] and [1], shadow clue was used to constrain the depth construction of object surface. In our method, the shadow information is explored in a novel way - that is, using the cast shadow boundary normal to assist recovery of surface normals and BRDF parameters.

3. Our System

Similar to the configuration of calibrated photometric stereo methods in previous work, our input images $\mathbf{I} = \{I_1, I_2, \dots, I_n\}$ are taken by a static camera where the target object is illuminated by varying directional light source. We put light source far from the object such that parallel light can be assumed. Input images are assumed to be taken under orthographic projection. The photometric calibration is carried out such that the pixel intensities represent the radiance. The incident radiance direction in each image is measured using a chrome sphere.

For simplicity's sake, we first assume that the object is built of a single material. Our method can be naturally extended to the objects assembled with different materials, which will be described in Section 6.

We model the reflectance property of an object's glossy

surface using Ward BRDF model [24]:

$$f(\mathbf{i}, \mathbf{o}) = \frac{\rho_d}{\pi} + \frac{\rho_s}{4\pi\alpha^2\sqrt{\cos\theta_i\cos\theta_o}} \exp\frac{-\tan^2\beta}{\alpha^2}, \quad (1)$$

where \mathbf{i} and \mathbf{o} denote the incident and outgoing light directions. θ_i and θ_o are polar angles of \mathbf{i} and \mathbf{o} respectively. α is the roughness coefficient, determining the size of the specular lobe on object surface. Larger α value corresponds to wider specular lobe with severely blurred boundary. ρ_d and ρ_s measure the diffuse and specular reflectance respectively. β is the angle between the surface normal and a halfway vector [24]. Note that other parametric reflectance models can also be employed similarly in our method.

Ward model is a combination of two terms. $\frac{\rho_d}{\pi}$ is the diffuse term. Other parameters form the specular term which is highly non-linear with respect to the unknowns. Directly optimizing all the unknowns in (1) usually results in a slow and unstable process. In our method, we separate the unknowns into two classes and estimate them separately:

- The global parameters ρ_s and α have fixed values for all the surface points and they can be estimated in the first pass. For computation efficiency, a few selected samples are sufficient to robustly estimate these parameters. This step is described in Section 4.
- The estimated global parameters are taken back to (1) to simplify the specular term. We then apply iterative optimization to computing the surface normals and other local parameters respectively for each pixel. This step is described in Section 5.

4. Estimation of Global Parameters

We first show that if the normal of one surface point is known, the global parameters can be computed. In order to use the normal information, we first study the BRDF measurement on object surface using single directional light source. The amount of outgoing light, L_o , is given by:

$$L_o = BRDF(\theta_i, \phi_i, \theta_o, \phi_o)L_i \cos\theta_i, \quad (2)$$

where L_i is the incoming light. Substituting (1) into (2), and generally assuming the unit intensity of the incoming light, we obtain

$$L_o = \frac{\rho_d}{\pi} \cos\theta_i + \frac{\rho_s}{4\pi\alpha^2} \sqrt{\frac{\cos\theta_i}{\cos\theta_o}} \exp\frac{-\tan^2\beta}{\alpha^2}. \quad (3)$$

Note that if the corresponding normal of one surface point q is known (the details of deriving this normal will be depicted in Section 4.2 and 4.3), the formulation of L_o can be greatly simplified since θ_o , θ_i , and β are determined. We then estimate other BRDF parameters by minimizing

$$g(q) = \sum_{0 \leq i < n} (I_i(q) - L_o(q, \rho_s, \alpha, \mathbf{n}, \rho_d))^2, \quad (4)$$

where I_i is one of the input images. In the rest of the paper, without causing ambiguity, we simplify the notation of $L_o(q, \rho_s, \alpha, \mathbf{n}, \rho_d)$ to $L_o(q)$. We propose the following iterative optimization algorithm to minimize g and estimate the BRDF parameters.

1. Initialize BRDF parameters.

α is initialized to a value between 0 and 0.4 while ρ_d and ρ_s are initialized to between 0 and 1.

2. Optimize ρ_d by fixing other parameter values.

In the diffuse term, ρ_d is the only parameter to be estimated for one pixel. It can be exactly computed by setting $\frac{\partial g}{\partial \rho_d} = 0$.

3. Optimize ρ_s and α by fixing ρ_d .

In the specular term, ρ_s and α are optimized together using the Levenberg-Marquet method.

4. Terminate.

We iterate steps 2 and 3 until the energy of the objective function does not decrease in successive two iterations.

Since energy g in steps 2 and 3 monotonically decreases, the whole process is guaranteed to converge to an optimal solution. This optimization process, in our experiments, converges rapidly given the small number of unknowns.

The above optimization process shows that the global BRDF parameters can be efficiently estimated from one or a set of known surface normals. In the following, we describe our method of estimating the surface normals using the shadow visual clue.

4.1. Shadow Separation

Since directional light is used in photometric stereo, the lighting in each image not only forms specular lobes on object surfaces but also produces cast shadow of the object on the background. This shadow is caused by the occlusion of incident light.

In order to capture the shadow of the object under different lightings, we place a small Lambertian board on the back of the object. To reduce the possible interreflection, the board is painted with mid-grey color to absorb most of the incident light while the shadow can still be faithfully detected. We show one example in Figure 2 (a). To verify that using the board does not significantly influence the light, we carried out several experiments using and without using the grey board and found that the differences are constantly very small. One comparison is given in Figure 2 where the root mean square (RMS) difference of pixel colors between (a) and (b), taking account of image noise, is only 0.0024.

For automatically separating the shadow and the foreground object in the input images, we observe that the

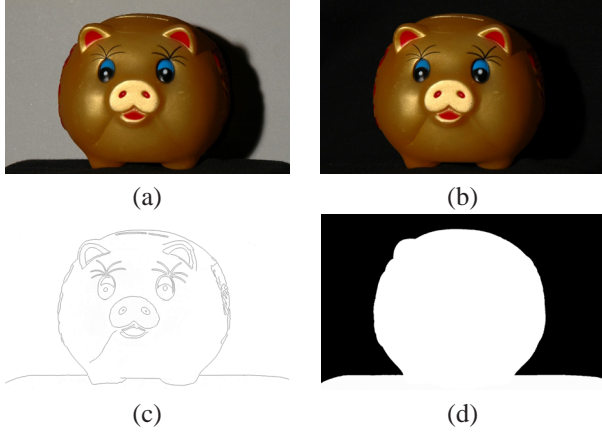


Figure 2. Mid-grey backboard configuration. (a) A mid-grey board is placed behind the object to capture the cast shadow. The reflection from the board is low. (b) Same scene captured with black background. The object color difference between (a) and (b) is very small. (c) and (d) show the automatic foreground extraction. (c) The edge detected using input image information. Most shadows boundary pixels are removed. (d) Using the level set method, the foreground map is formed (with white pixels) by curve evolving from the background.

shadow region varies among the images due to the change of lighting directions whereas the object in the images has fixed position. So the object boundary and textures constantly appear in multiple images. We detect the consistent object edges in multiple images using the method proposed in [27] where a two-component GMM is constructed for each pixel to robustly reject outliers. The output of this method includes edge pixels detected on the object as shown in Figure 2(c). To extract the object in all images, a closure of outermost edge points is computed using the flood fill or the geometric level set [16] methods. One example of the extracted object is shown in Figure 2(d).

When the object has been extracted, only shadow remains in the images. Due to the use of directional light, the shadow penumbra is narrow. We, thus, can detect shadow boundary using another edge detection process.

4.2. Normal Estimation Using Shadow Boundary

Once the shadow boundary pixels are detected, we compute image gradients on these pixels. As shown in Figure 3(a), each shadow boundary pixel p in back plane Φ can be mapped to at least one object surface point p^* . We study the following two situations in order to compute the surface normal \mathbf{n}_{p^*} at p^* . An illustration is shown in Figure 3.

Case 1: The shadow boundary pixel, e.g., p in Figure 3(a), maps to a single surface point p^* along the light ray and the first order derivative exists at p^* . Ideally, the surface normal \mathbf{n}_{p^*} at p^* is projected to the shadow boundary normal \mathbf{n}_p . However, in discrete image space, there may exist

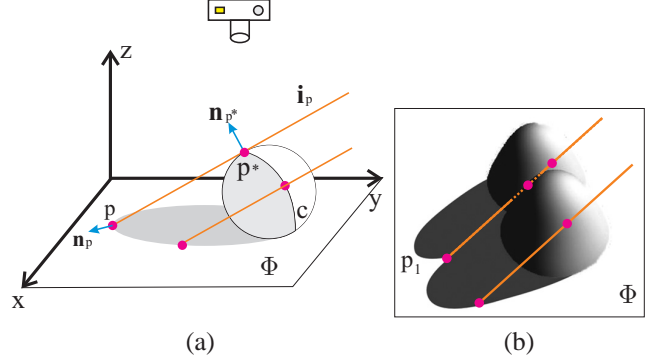


Figure 3. (a) Shadow formation. \mathbf{i}_p indicates the light direction. The object casts shadow on back plane Φ . \mathbf{n}_p is the 2-D shadow boundary normal of p in Φ . p^* is the corresponding surface point and \mathbf{n}_{p^*} is the surface normal. (b) Shadow silhouette. p_1 is at the junction of two hills' cast shadow. The gradient computed on p_1 does not map to any surface normal.

quantization errors by computing shadow boundary gradient. In order to reduce them, we adopt a robust estimation algorithm, which will be depicted in Section 4.3.

Case 2: The shadow boundary pixel, for example, p_1 in Figure 3(b), is at the junction of two or more hills' cast shadow. The normal of the corresponding surface point cannot be determined since there exist at least two surface points projected to that pixel.

In order to estimate the surface normal using cast shadow, the pixels in case 2 should be rejected. We employ a feature detection method to detect the shadow boundary pixels at junctions. Specifically, we construct the Harris matrix $M(x, y)$ [10] for each shadow boundary pixel $I(x, y)$. The junctions are detected when the corner strength $S(x, y) = \frac{\det M(x, y)}{\text{tr} H(x, y)}$ is larger than a threshold.

After feature removal, we compute image gradients on the remaining cast shadow boundary pixels and denote them as $\{\mathbf{n}_1^i, \mathbf{n}_2^i, \dots, \mathbf{n}_{m_i}^i\}$ for each image I_i . The corresponding surface normal for \mathbf{n}_p^i is denoted by $\mathbf{n}_{p^*}^i$, which is computed under the following conditions.

- $\mathbf{n}_{p^*}^i$ is a surface normal; it must be perpendicular to light direction \mathbf{i}_p :

$$\mathbf{i}_p \cdot \mathbf{n}_{p^*}^i = 0. \quad (5)$$

- $\mathbf{n}_{p^*}^i$, \mathbf{n}_p^i , and \mathbf{i}_p are coplanar, we thus have

$$(\mathbf{i}_p \times \mathbf{n}_p^i) \cdot \mathbf{n}_{p^*}^i = 0. \quad (6)$$

- $\mathbf{n}_{p^*}^i$ is a unit vector

$$\|\mathbf{n}_{p^*}^i\| = 1. \quad (7)$$




Figure 4. Determining the surface points. On object surface, the point p^* , on which the normal \mathbf{n}_{p^*} is computed, is only known along the ray \mathbf{i}_p , as shown in green.

Combining (5), (6), and (7), we solve a linear equation system to compute the surface normals corresponding to the shadow boundary pixels.

Note that our feature removal may not reject all shadow pixels in case 2 in computing surface normals. In addition, errors possibly produced by inaccurate image gradient and shadow boundary computation may also contaminate the results. In the next section, we introduce a robust clustering method to further refine the parameter estimation.

4.3. Global Parameter Estimation and Refinement

We have recovered a set of possible surface normals corresponding to the cast shadow boundary pixels. In order to robustly estimate the BRDF parameters in (1), the 3-D positions of corresponding surface points should also be known such that (4) can be optimized.

In most cases, the attached shadow on object surfaces fades gradually and no clear boundary exists between the illuminated pixels and the attached shadow. One example is shown in Figure 4, in which we only know that p^* corresponding to normal \mathbf{n}_{p^*} is along the light ray \mathbf{i}_p . It is not obvious to determine its exact 3-D position. In our approach, a small number of surface candidates \mathbf{S}_p of p^* along \mathbf{i}_p are collected. All of them are in a transition region between fully illuminated and shadowed pixels. This region is shown in green in Figure 4.

Note that using shadow boundary, we have computed several surface normals in image i , each normal $\mathbf{n}_{p^*}^i$ has a set of position candidates \mathbf{S}_p^i . To simultaneously estimate the normal position and the global BRDF parameters, we introduce the following optimization process.

1. For each surface normal $\mathbf{n}_{p^*}^i$:
 - (a) We select one candidate normal position from \mathbf{S}_p^i and take it into the iterative optimization described in Section 4 to minimize (4). The residual error is recorded.
 - (b) If there exists candidate in \mathbf{S}_p^i not used in the above computation, repeat (a).
 - (c) We compare the recorded residual errors, and find the surface point in \mathbf{S}_p^i producing the small-

est one. Then we record the values of BRDF parameters $\rho_s(i, p)$ and $\alpha(i, p)$ produced using that surface point.

2. For all recorded values for BRDF parameters ρ_s and α , we employ the mean shift clustering [6] to robustly estimate the best set.

In regard to step 1, among all candidates for a surface point, we find the one producing the least residual error in the optimization and regard it as the estimated surface point. It is shown to be a reliable measurement in experiments. Its computation is also efficient.

Step 2 is based on an observation that there may exist noise and outliers. The correspondingly estimated parameter values using these outliers have small probability to be similar to each other. So we use the statistical clustering method to refine the parameter estimation.

5. Surface Reconstruction

After estimating the specular parameters, surface normals and diffuse albedo for all surface points can be reconstructed without highlight separation. The surface texture is the recovered diffuse albedo. We first rewrite (3) as

$$L_o(p) = \frac{\rho_d(p)}{\pi} \mathbf{i} \cdot \mathbf{n}(p) + \frac{\rho_s}{4\pi\alpha^2} \sqrt{\frac{\mathbf{i} \cdot \mathbf{n}(p)}{\mathbf{o} \cdot \mathbf{n}(p)}} \exp\left\{\frac{1}{\alpha^2} \left\{1 - \frac{1}{[\mathbf{h}(p) \cdot \mathbf{n}(p)]^2}\right\}\right\},$$

where \mathbf{i} and \mathbf{o} denote the incident and outgoing light directions respectively. \mathbf{h} and \mathbf{n} represent a half-way vector and the surface normal. We minimize energy function $g(p)$ where

$$g(p) = \sum_{0 \leq i < n} (I_i(p) - L_o(p))^2 \quad (8)$$

in a way similar to that for (4) to reconstruct surface normals. Comparing to previous methods using only Lambertian components [25, 5, 20, 13], our formulation has the same number of unknowns. We employ an iterative optimization to estimate them.

1. For each pixel q on an object surface:
 - (a) The normal \mathbf{n}_q of q is initialized as the linear interpolation of the recovered normals from the previous step.
 - (b) **Optimize $\rho_d(p)$.** $\rho_d(p)$ can be exactly computed by setting $\frac{\partial g(p)}{\partial \rho_d(p)} = 0$.
 - (c) **Optimize normal \mathbf{n}_q .** The normal vector is estimated using a gradient based optimization method. Given the small number of unknowns, it converges rapidly.
 - (d) Repeat (b) and (c) until convergence.

The energy in steps (b) and (c) monotonically decreases, so the convergence is guaranteed. In all our experiments, the computation terminates in 10 iterations.




Figure 5. Mouse. (a) and (b) show two out of six input images containing large region of highlight, (c) The recovered surface normal map encoded using RGB colors. (d) Estimated surface textures. (e) Rendered mouse under a novel lighting direction. (f) The camera captured image using the same lighting direction, which is similar to our rendered result (e).

6. Reconstructing Multiple-Material Surface

Our method can also be naturally applied to surface reconstruction of multiple materials. With the similar configuration, we can obtain object silhouettes projected to the backboard. Similarly applying the method described in Section 4, a set of surface normals are computed. Then we employ the method described in Section 4.3 to estimate several sets of optimal specular parameter values using these normals. We place these values in the parameter space. Assuming there are M different materials where $M > 1$, rather than using mean shift clustering to find a unique parameter configuration, we employ mean shift segmentation to form M partitions in the parameter space. Then the centers of all partitions with highest density are regarded as the locations of estimated specular parameter values.

After specular estimation, we use these parameter values to recover the surface normal for each pixel. Using the optimization described in Section 5, we are able to compute a set of optimal $\mathbf{n}^*(j)$ and $\rho_d^*(j)$, where $0 \leq j < M$, in terms of minimizing residual error $\xi_{p,j}$ which is given by

$$\xi_{p,j} = \sum_{0 \leq i < n} (I_i(p) - L_o(p, \rho_s(j), \alpha(j), \mathbf{n}^*(j), \rho_d^*(j)))^2.$$

There should exist only one true surface normal for each pixel p . However, simply comparing $\xi_{p,j}$ for all j 's and



Figure 6. Surface reconstruction result by our method for the example shown in Figure 5.

selecting the normal corresponding to the smallest $\xi_{p,j}$ is unreliable in this case. Considering the material continuity property of the object surface, we treat the selection of an optimal specular parameter set as a labeling problem in Markov Random Field where the parameter set of each pixel can be selected from M different assignments. Thus we minimize an energy defined as

$$U(f) = \sum_{p \in \mathcal{P}} \xi_{p,f(p)} + \alpha \sum_{\{p,q\} \in \mathcal{N}} \mathcal{H}(s_{f(p)}, s_{f(q)}), \quad (9)$$

where $f : \mathcal{P} \rightarrow \mathcal{S}$, \mathcal{P} and \mathcal{S} denote the pixel set and label set respectively. So $f(p)$ returns the label at p . s_j denotes the j th set of the specular parameters, where $s_j = [\rho_s(j), \alpha(j)]^T$ and $0 \leq j < M$. \mathcal{N} represents the neighborhood set and α is a weight. $\sum \xi_{p,f(p)}$ is the data term. $\sum \mathcal{H}(a, b)$ is the smoothness term given by

$$\mathcal{H}(a, b) = 1 - \exp\left[-\frac{1}{2}(a - b)^T \Sigma^{-1}(a - b)\right],$$

which encodes the labeling smoothness between adjacent pixels. Σ denotes the covariance matrix between vectors a and b . We minimize (9) using graph cuts [3]. The whole surface is thereby segmented into partitions with estimated optimal normal and albedo for each pixel.

In order to faithfully reconstruct normals for an object surface consisting of M different materials, M sets of global parameters must be recovered. It implies there exist surface points from each material projected to the cast shadow boundaries. This condition, in our experiments, can be generally satisfied when several input images under different lightings are used.

7. Results

We show experimental results in this section. Most of the surfaces contain wide specular lobes in input images. Since our algorithm does not rely on separation of specular components, we can faithfully recover surface textures even if the highlight color is similar to the texture color.

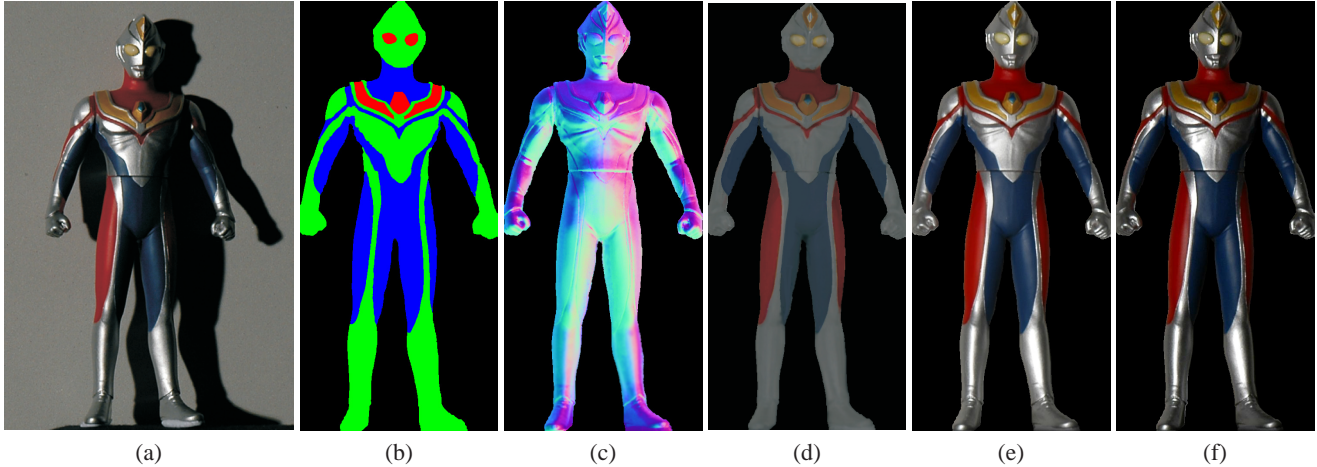


Figure 8. Toy. (a) The toy surface consists of multiple materials, and the reflection is very complex. (b) Recovered material composition map. (c) and (d) The computed surface normals and textures. (e) We re-render the toy under a novel lighting condition. (f) The ground truth image taken by the camera for comparison.

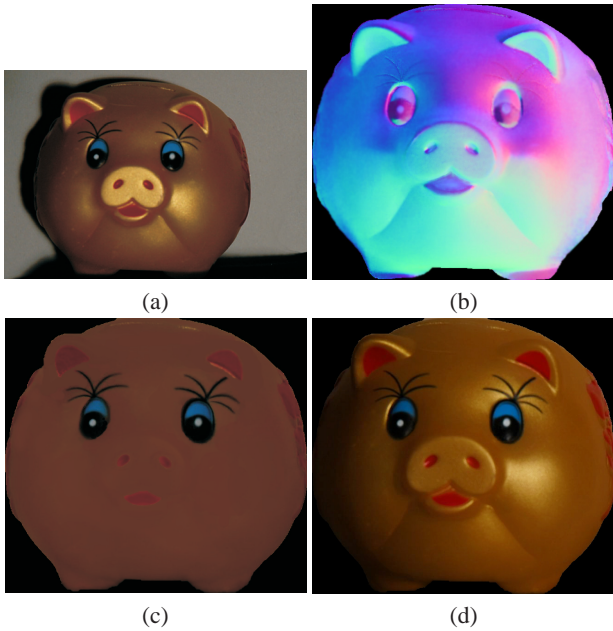


Figure 7. Saving box. (a) One of the input images, (b) and (c) Reconstructed normals and textures. (d) The rendered image under novel lighting.

Figure 5 shows a “mouse” example with large highlight area. (a) and (b) illustrate two out of the six input images. The object and shadow boundaries are automatically extracted from these images. The surface normal and texture shown in (c) and (d) are faithfully reconstructed. We evaluate the efficacy of our method by rendering the object under a novel lighting as shown in (e) and comparing it with the ground truth image (f). The RMS error is only 0.0326, showing that our recovered BRDF parameters and surface

normals are accurate.

In Figure 6, we show our surface reconstruction result using the “mouse” example. Since our method use shadow visual clue to estimate the global parameters, the complexity of general surface normal estimation is reduced. In comparison, the optimization in [8] takes hours to produce a surface result whereas our method uses only 50 minutes in computation.

Figure 7 shows another object in which the reflection is complex and highlight can be seen in different regions in the image. We only use 5 input images in this example. We faithfully recover surface normals and texture as shown in (b) and (c). The rendered image (d) under novel lighting is visually satisfying.

Figure 8 illustrates a very challenging example. It is a toy model with multiple materials on the surface. Each material has its own roughness and specular property. We use 10 input images and the reflections in the input images are complex. Using our method, different materials can be successfully recovered as shown in (b). The material shown in green is a kind of hard plastic which has strong and overlapping highlight in all images. The material in blue exhibits smaller highlight whereas the red material is near Lambertian. The recovered normals in (c) preserve sufficient fine details and structures. Comparing our rendered image with the one taking by a camera under the same lighting condition, the RMS error is only 0.0423.

8. Conclusion

To conclude, in this paper, we have proposed an efficient photometric stereo method to estimate the BRDF parameters and reconstruct surface normals. Our method does not require general specular component separation and outper-

forms direct specular fitting methods in terms of accuracy and convergence rate. Our method uses a new configuration to capture the cast shadow of an object on a back plane. Given the small number of input images, the shadow information is used to estimate surface normal samples and compute the global parameters. Once these parameters are estimated, the local reflectance and dense surface normals can be computed separately for all pixels. This largely reduces the computation complexity and improves the quality of the reconstructed object surfaces and texture maps. Our future work consists of exploring other BRDF models on multiple materials, and applying our optimization to anisotropic reflectance models.

References

- [1] N. G. Alldrin and D. J. Kriegman. Toward reconstructing surfaces with arbitrary isotropic reflectance: a stratified photometric stereo approach. *ICCV*, 2007. 2
- [2] S. Barsky and M. Petrou. The 4-source photometric stereo technique for three-dimensional surfaces in the presence of highlights and shadows. *IEEE Trans. Pattern Anal. Mach. Intell.*, 25(10):1239–1252, 2003. 1, 2
- [3] Y. Boykov, O. Veksler, and R. Zabih. Fast approximate energy minimization via graph cuts. *IEEE Trans. Pattern Anal. Mach. Intell.*, 23(11):1222–1239, 2001. 6
- [4] M. K. Chandraker, S. Agarwal, and D. J. Kriegman. Shadowcuts: Photometric stereo with shadows. *CVPR*, 2007. 2
- [5] E. Coleman and R. Jain. Obtaining 3-dimensional shape of textured and specular surfaces using four-source photometry. *Comp. Vision, Graphics and Image Processing*, 1982. 2, 5
- [6] D. Comaniciu and P. Meer. Mean shift: A robust approach toward feature space analysis. *IEEE Trans. Pattern Anal. Mach. Intell.*, 24(5):603–619, 2002. 5
- [7] A. S. Georghiades. Incorporating the torrance and sparrow model of reflectance in uncalibrated photometric stereo. *iccv*, 02:816, 2003. 2
- [8] A. S. Georghiades. Recovering 3-d shape and reflectance from a small number of photographs. In *EGRW*, pages 230–240, 2003. 1, 2, 7
- [9] D. B. Goldman, B. Curless, A. Hertzmann, and S. M. Seitz. Shape and spatially-varying brdfs from photometric stereo. In *ICCV*, pages 341–348, 2005. 1, 2
- [10] C. Harris and M. Stephens. A combined corner and edge detector. In *Alvey Vision Conf.*, pages 147–151, 1988. 4
- [11] A. Hertzmann. Example-based photometric stereo: Shape reconstruction with general, varying brdfs. *IEEE Trans. Pattern Anal. Mach. Intell.*, 27(8):1254–1264, 2005. Member- Steven M. Seitz. 2
- [12] G. Klinker, S. Shafer, and T. Kanade. The measurement of highlights in color images. *International Journal of Computer Vision*, 2(1):7–32, 1988. 2
- [13] S. P. Mallick, T. E. Zickler, D. J. Kriegman, and P. N. Belhumeur. Beyond lambert: Reconstructing specular surfaces using color. In *CVPR*, pages 619–626, 2005. 1, 2, 5
- [14] S. Nayar, K. Ikeuchi, and T. Kanade. Determining shape and reflectance of lambertian, specular, and hybrid surfaces using extended sources. In *Workshop on Industrial Applications of Mach. Intell. and Vision*, pages 169 – 175, 1989. 2
- [15] S. 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. 1
- [16] S. Osher and J. A. Sethian. Fronts propagating with curvature-dependent speed: algorithms based on hamilton-jacobi formulations. *J. Comput. Phys.*, 79:12–49, 1988. 4
- [17] Y. Sato and K. Ikeuchi. Temporal-color space analysis of reflection. Technical Report CMU-CS-92-207, 1992. 2
- [18] Y. Sato, M. D. Wheeler, and K. Ikeuchi. Object shape and reflectance modeling from observation. In *SIGGRAPH*, pages 379–387, 1997. 2
- [19] S. A. Shafer. Using color to separate reflection components. *Color Research and Application*, 10(4):210–218, 1985. 2
- [20] F. Solomon and K. Ikeuchi. Extracting the shape and roughness of specular lobe objects using four light photometric stereo. *IEEE Trans. Pattern Anal. Mach. Intell.*, 18(4):449–454, 1996. 1, 2, 5
- [21] H. D. Tagare and R. J. P. deFigueiredo. Simultaneous estimation of shape and reflectance map from photometric stereo. *CVGIP: Image Underst.*, 55(3):275–286, 1992. 2
- [22] P. Tan, L. Quan, and S. Lin. Separation of highlight reflections on textured surfaces. In *CVPR*, pages 1855–1860, 2006. 2
- [23] R. T. Tan and K. Ikeuchi. Separating reflection components of textured surfaces using a single image. *IEEE Trans. Pattern Anal. Mach. Intell.*, 27(2):178–193, 2005. 2
- [24] G. J. Ward. Measuring and modeling anisotropic reflection. *SIGGRAPH Comput. Graph.*, 26(2):265–272, 1992. 3
- [25] R. J. Woodham. Photometric method for determining surface orientation from multiple images. *Optical Engineering*, 19(1):139–144, 1980. 1, 5
- [26] T.-P. Wu and C.-K. Tang. Dense photometric stereo using a mirror sphere and graph cut. In *CVPR*, pages 140–147, 2005. 2
- [27] W. Xu, M. Jenkin, and Y. Lesperance. A multi-channel algorithm for edge detection under varying lighting. In *CVPR*, pages 1885–1892, 2006. 4