

Uncalibrated Photometric Stereo for Unknown Isotropic Reflectances

Feng Lu^{1*} Yasuyuki Matsushita² Imari Sato³ Takahiro Okabe¹ Yoichi Sato¹

¹The University of Tokyo, Japan ²Microsoft Research Asia, China

³National Institute of Informatics, Japan

{lufeng,takahiro,yasato}@iis.u-tokyo.ac.jp, yasumat@microsoft.com, imarik@nii.ac.jp

Abstract

We propose an uncalibrated photometric stereo method that works with general and unknown isotropic reflectances. Our method uses a pixel intensity profile, which is a sequence of radiance intensities recorded at a pixel across multi-illuminance images. We show that for general isotropic materials, the geodesic distance between intensity profiles is linearly related to the angular difference of their surface normals, and that the intensity distribution of an intensity profile conveys information about the reflectance properties, when the intensity profile is obtained under uniformly distributed directional lightings. Based on these observations, we show that surface normals can be estimated up to a convex/concave ambiguity. A solution method based on matrix decomposition with missing data is developed for a reliable estimation. Quantitative and qualitative evaluations of our method are performed using both synthetic and real-world scenes.

1. Introduction

Photometric stereo recovers the surface normals of a scene from a set of images recorded under varying lighting conditions. The original method of Woodham [32] assumes Lambertian reflectance and known directional lightings. To make the approach more practical, there have been a variety of recent studies on relaxing these assumptions. Most focus on relaxing either of the two assumptions while retaining the other.

There are previous approaches that estimate surface normals under less restricted conditions by making use of intensity profiles [16, 19, 26, 25]. An intensity profile is an ordered sequence of measured intensities at a pixel under varying illumination. The ordered measurement can offer information that cannot be found in the bare set of measure-

ments. However, the previous methods have certain limitations in recovering the surface normals. They either still require calibrated lightings or Lambertian reflectance [25] or only cluster similar surface orientations [19], or require additional assumptions on occluding boundaries, certain reflectance models [26], and calibrated reference objects [16].

In this paper, we further exploit the observation on intensity profiles and develop a photometric stereo method that works with general and unknown isotropic bidirectional reflectance distribution functions (BRDFs) and unknown lighting directions. We show that the proposed method can reliably recover surface normals up to a binary convex/concave ambiguity when the scene has a uniform (up to albedo differences) reflectance.

Our method has the following advantages over the previous methods. First, our method solves a more general problem compared with the prior ones. We assume uncalibrated lights and unknown and general isotropic BRDFs. Our solution method is highly deterministic without assuming occluding boundaries or certain reflectance models, and it recovers surface normals up to only a binary convex/concave ambiguity. Second, we show the relation between the surface reflectance property and the intensity distribution of the observed intensity profile. In particular, we calculate the skewness of the intensity distribution and demonstrate that the skewness can be used to infer an important linear coefficient for recovering the surface normals for unknown reflectance, without using reference objects or other priors. Finally, we develop a robust solution technique that only selects and uses reliable measurements for highly correlated surface normals, which makes a step toward practical photometric stereo.

1.1. Previous work

A variety of previous studies have been conducted to relax the constraints of the Lambertian model and known lighting directions in photometric stereo. Non-Lambertian reflectances have been handled either by 1) regarding non-Lambertian components as outliers, or 2) using more general reflectance models. The former class of method-

*Part of this work was done while the first author was visiting Microsoft Research Asia as a research intern. This work was also supported in part by Grant-in-Aid for Scientific Research on Innovative Areas “Shitsukan” from MEXT, Japan.

s finds non-Lambertian observations in a robust estimation framework using color-cues [4], median filter [22], rank-minimization [33], or Markov random field with hidden variables [34]. The latter class of methods studies reflectance properties, such as bilateral symmetry [1], reflective symmetry about the halfway vector [18], isotropy and monotonicity [17, 28], and other reflectance symmetries [30]. There are methods that use sophisticated parametric models, such as ones that use the Ward model [14, 9], Lambertian+specular models [23] and bi-variate BRDF representation [3]. These methods assume that the lighting directions are known.

There are uncalibrated photometric stereo techniques wherein the lighting directions are unknown. Most of the existing methods attempt to resolve the Generalized Bas-Relief (GBR) ambiguity [6] with the Lambertian model. Various surface properties are used for resolving the GBR ambiguity, such as diffuse maxima [12], specularity [11, 10], low-dimensional space [5], minimum entropy [2], interreflections [8], color profiles [27], reflectance symmetry [30], and certain configuration of the light sources [35]. These methods rely on the assumption that the diffuse reflectance component follows the Lambertian model.

Handling both non-Lambertian reflectances and uncalibrated light sources is far more challenging and, as a result, has been less studied. Silver [29] and Hertzmann and Seitz [16] use reference objects that have unknown but the same reflectance as a target object for estimating its surface normals. Georgiades [13] uses the Torrance-Sparrow model and proposes to optimize over a large set of variables including the model parameters, surface normals, and illumination. Chandraker *et al.* [7] recover surface iso-contours from the differential images by restricting the positions of the light sources to a circle around the camera axis. They need additional information such as an initial normal to determine surface normals. Sato *et al.* [26] propose a method that uses intensity profiles for estimating surface normals, but the method is limited to the Lambertian or Torrance-Sparrow reflectance models, while our method can deal with general isotropic reflectances. Okabe *et al.* [24] use the similarity of attached shadow codes to deal with general BRDFs. Both of these methods assume that the surface has visible occluding contours, which provide knowledge of surface normals perpendicular to the viewing direction for resolving the ambiguity. Unlike these approaches, our method does not require such assumptions about reference object, initial normals, or visible occluding contours.

2. What does intensity profile tell us?

An intensity profile is a sequence of radiance intensities recorded at a pixel across multi-luminance images. It has been used in various problems due to its following properties when assuming no cast shadow or interreflection.




Figure 1. Illustration of intensity profiles. Surface points A, B, and C have the same reflectance, but D is different. A, C, and D have the same surface normal, while B has a different normal.

Orientation-consistency: Intensity profiles become exactly the same, if and only if they correspond to the same surface normal orientation and material (A and C in Fig. 1). Using this simple observation, surface normals can be determined by looking up a pre-stored table indexed with surface brightness values [29], or match the intensity profiles to those from a reference object [16]. Methods based on this observation require the use of reference objects of known shapes and the same material as the target.

Geometry-extrema: For many materials, intensity profiles reach the extremas synchronously, if and only if they correspond to the same surface normal (A and D in Fig. 1). This fact is used for clustering surface orientations [19] without determining the orientations.

Similarity: Similarity between intensity profiles is strongly related with the difference between surface normals for the same material (A and B in Fig. 1). Sato *et al.* [26] analyze this relationship and exploit it to recover surface normals in the cases of Lambertian and Torrance-Sparrow reflectance and evenly distributed light sources with an assumption of having occluding boundaries.

This paper makes a further observation about intensity profiles and introduces the notion of *conditional linearity*. Different from [26], we do not restrict our analysis to certain reflectance models. Instead, we take into account more general isotropic reflectances in the MERL BRDF database [21].

Conditional linearity: For most real-world isotropic reflectances, we observe a strong *linear* relation between the distance among intensity profiles seen under evenly distributed lightings and the angular difference of surface normals, up to a certain normal angular difference. We also observe that the linear coefficient is material-dependent and closely related to the intensity distribution of the observed intensity profile. These observations allow us to develop an uncalibrated photometric stereo method that works with general and unknown isotropic reflectances.

2.1. Geodesic distance of intensity profiles and normal angular difference

Let us assume evenly distributed light directions and a scene with a uniform material; we show later that these can be relaxed to some extent. Let $\{\mathbf{n}_p, \mathbf{n}_q\}$ be a surface nor-




Figure 2. Geodesic distance calculation. Using geodesic distance (right) preserves a linear relationship over a greater range of angular differences in comparison with using Euclidean distance (left).

mal pair, and $\{\mathbf{I}_p, \mathbf{I}_q\}$ be the corresponding pixel intensity profiles in a normalized form as below:

$$\mathbf{I}_p = [I_p^1, \dots, I_p^L]^T = [\tilde{I}_p^1, \dots, \tilde{I}_p^L]^T / \sqrt{\sum_l (\tilde{I}_p^l)^2}, \quad (1)$$

where \tilde{I}_p^l is the recorded intensity at the p -th pixel for a scene point ($p = 1, \dots, P$), under the l -th lighting direction ($l = 1, \dots, L$), and I_p^l is the normalized intensity.

The previous methods have shown that the similarity of intensity profiles and surface normals are strongly correlated [19, 26]. The similarity can be straightforwardly defined using the Euclidean distance of two intensity profiles as $\|\mathbf{I}_p - \mathbf{I}_q\|_2$. It indeed correlates with the angular difference $\cos^{-1}(\mathbf{n}_p^T \mathbf{n}_q)$ of surface normals \mathbf{n}_p and \mathbf{n}_q at scene points p and q ($\mathbf{n}_p, \mathbf{n}_q \in \mathbb{R}^{3 \times 1}$); however, the linear relationship holds only in a limited range as depicted in Fig. 2 (left). To extend the range, our method uses a geodesic distance $d_G(\mathbf{I}_p, \mathbf{I}_q)$ instead of $\|\mathbf{I}_p - \mathbf{I}_q\|_2$ as it is used in [26] to measure the similarity of more diverse normals.

The geodesic distance corresponds to the shortest path between two nodes in a graph, and is computed by adding up small Euclidean distances of neighboring points along the path [31]. In our case, we first compute and keep the Euclidean distances for nearby $\{\mathbf{I}_p, \mathbf{I}_q\}$

$$d(\mathbf{I}_p, \mathbf{I}_q) = \begin{cases} \|\mathbf{I}_p - \mathbf{I}_q\|_2 & \text{if } \|\mathbf{I}_p - \mathbf{I}_q\|_2 < \varepsilon_p \\ +\infty & \text{otherwise,} \end{cases} \quad (2)$$

where ε_p is a threshold at the point p . The geodesic distance $d_G(\mathbf{I}_p, \mathbf{I}_q)$ can be computed using these $d(\mathbf{I}_p, \mathbf{I}_q)$ as

$$d_G(\mathbf{I}_p, \mathbf{I}_q) = D_{sp}(\{d(\mathbf{I}_p, \mathbf{I}_q)\}), \quad (3)$$

where function $D_{sp}(\cdot)$ calculates the shortest path on the graph using Dijkstra’s algorithm. In this way, the geodesic distance $d_G(\mathbf{I}_p, \mathbf{I}_q)$ comprises a set of small Euclidean distances that are within the linear range. Therefore, the linearity is well preserved in the geodesic distance over a greater range of angular differences, as shown in Fig. 2 (right).

We examine this linearity for all 100 materials in the MERL BRDF database [21] by plotting the $\cos^{-1}(\mathbf{n}_p^T \mathbf{n}_q)$ and $d_G(\mathbf{I}_p, \mathbf{I}_q)$ values. Fig. 3 shows four typical plots, from which we can make the following observations. First, the




Figure 3. Examples of the linear relationship. Four typical shapes (dark regions) are shown for the synthetic surfaces under uniform lightings. Solid lines indicate linear fitting results within limited regions, while dotted lines are the results for non-uniform lights.

use of geodesic distance generally shows the linear relationship with the angular difference of normals in a large range. Second, most materials, even those having complex reflectances, show obvious linear relationships. Such a linear relationship generally holds in the range of 0 to 70° of the angular difference, but it does not span the entire range for many materials. Finally, the linear coefficient, or slope, varies with the material (solid lines in Fig. 3). The slope is insensitive to random noise in lighting directions as shown by the dotted lines in Fig. 3, which are the line fittings to the plots produced with fluctuating light directions by 7° standard deviations. A more thorough analysis on non-uniform light distributions is given in Sec. 4.1.

Interestingly, the first case in Fig. 3 is actually the problem solved in [26], which corresponds to an ideal subset of the more general reflectances handled by our method. Based on the above observations, we define a partial linear conversion from the geodesic distance $d_G(\mathbf{I}_p, \mathbf{I}_q)$ to the angular difference of surface normals $\cos^{-1}(\mathbf{n}_p^T \mathbf{n}_q)$ in a certain range bounded by a threshold δ as

$$\cos^{-1}(\mathbf{n}_p^T \mathbf{n}_q) = \alpha_m d_G(\mathbf{I}_p, \mathbf{I}_q) \text{ if } \cos^{-1}(\mathbf{n}_p^T \mathbf{n}_q) < \delta, \quad (4)$$

with the material-dependent slope α_m . Please note that inferring such linear coefficient α_m is important for determining the surface normals without assuming the occluding boundaries used in [26] or other priors.

2.2. Reflectance property and linear coefficient

The linear coefficient α_m described in the previous section is material dependent, *i.e.*, it is related to the surface reflectance property of a material. To characterize such a reflectance property, we show that the *intensity distribution* observed in an intensity profile conveys information about the reflectance property for the material. Fig. 4 show four original intensity profiles, where each row corresponds to the




Figure 4. Intensity profiles w.r.t. material and surface normal. The top row shows intensity profiles captured at two surface normals for a specular material. The bottom row shows intensity profiles captured at two surface normals for a diffuse material. The intensity values are plotted in 1D to show their distributions.




Figure 5. Average skewness and α_m^{-1} values of 100 materials. The correlation coefficient is 0.98. Error bars for different surface normals and the line fitting result are shown.

same material but different surface normals. These figures indicate that an intensity profile’s shape depends on both material (reflectance property) and surface normal. However, its intensity distribution, which does not rely on the intensity order as shown by the 1D plots in Fig. 4, appears stable against surface normal changes for the same material. Based on these observations, we compute the skewness of the intensity distribution for characterizing it with an aim of deriving the linear coefficient α_m via the skewness.

The skewness γ of an intensity distribution, which is *irrelevant* to the intensity/lighting order, is calculated as

$$\gamma(\mathbf{I}) = L^{\frac{1}{2}} \frac{\sum_l (I^l)^3}{\left(\sum_l (I^l)^2\right)^{\frac{3}{2}}}, \quad (5)$$

where \mathbf{I} is an intensity profile, and I^l is its l -th element that corresponds to the l -th lighting direction. Indeed, the skewness of the intensity distributions has high correlation with the inverse of the linear coefficient α_m^{-1} . To examine this, we plot the skewness γ and inverse slope α_m^{-1} using all 100 materials using synthetic scenes as shown in Fig. 5. It shows a linear relation with a correlation coefficient of 0.98. Fig. 5 also shows error bars that demonstrate the stability of skewness values computed across diverse surface normals.

Therefore, it is efficient to estimate α_m for unknown materials from the skewness of the intensity distribution. In this way, Eq. (4) becomes deterministic in our method.

Fig. 5 also shows that the skewness increases as the materials vary from matte to shiny ones. This is because specular components generate large pixel intensities only under a limited light directions, which increases the skewness of the intensity distribution. Note that a few outliers exist in Fig. 5; they correspond to materials that have both significant diffuse and very narrow specular lobes. An example of such materials is one in Fig. 3 (bottom-right), and these materials show larger errors in our experiment.

3. Surface normal recovery

We describe our proposed method for recovering surface normals based on the discussion in Sec. 2. From the observed intensity profiles $\{\mathbf{I}_p\}$, we compute the geodesic distance $d_G(\mathbf{I}_p, \mathbf{I}_q)$ using Eq. (2) and Eq. (3). We then convert $d_G(\mathbf{I}_p, \mathbf{I}_q)$ into a normal angular difference $\cos^{-1}(\mathbf{n}_p^T \mathbf{n}_q)$ using Eq. (4). As discussed in the previous section, since the conversion is only valid when $\cos^{-1}(\mathbf{n}_p^T \mathbf{n}_q) < \delta$ holds, we rewrite Eq. (4) as the following:

$$\cos^{-1}(\mathbf{n}_p^T \mathbf{n}_q) = \begin{cases} \alpha_m d_G(\mathbf{I}_p, \mathbf{I}_q) & \text{if } \alpha_m d_G(\mathbf{I}_p, \mathbf{I}_q) \leq \delta \\ \text{Undefined} & \text{otherwise,} \end{cases} \quad (6)$$

where α_m is obtained by using the skewness (Sec. 2.2).

In our implementation, the threshold ε_p in Eq. (2) is empirically set to the 10-th shortest distance from \mathbf{I}_p to any other \mathbf{I}_q , and threshold δ in Eq. (6) is set to $\pi/4$ to ensure good linear regions in Fig. 3 for different materials.

3.1. Formulation

We wish to recover surface normals of scene points that correspond to P pixels in the observed image from a set of images taken under varying unknown lightings. Let the surface normal matrix be $\mathbf{N} = [\mathbf{n}_1, \mathbf{n}_2, \dots, \mathbf{n}_P] \in \mathbb{R}^{3 \times P}$ that we solve for. We define the observation matrix \mathbf{O} as

$$\mathbf{O} = \{o_{p,q} = \mathbf{n}_p^T \mathbf{n}_q\} \in \mathbb{R}^{P \times P}, \quad (7)$$

whose elements $o_{p,q} = \mathbf{n}_p^T \mathbf{n}_q$ are readily obtained from Eq. (6). In particular, the diagonal elements $o_{p,p}$ in \mathbf{O} are all ones, which ensures the unit normal length constraint for $\{\mathbf{n}_p\}$. Notice that Eq. (6) has undefined cases for some p and q ; therefore, only a portion of \mathbf{O} ’s elements have well-defined values. In other words, the observation matrix \mathbf{O} has missing elements. With a sparse error matrix \mathbf{E} that accounts for the errors due to the missing entries, the relationship between the observation matrix \mathbf{O} and surface normal \mathbf{N} can be written as

$$\mathbf{N}^T \mathbf{N} = \mathbf{O} + \mathbf{E}. \quad (8)$$

We wish to solve for surface normal N by using the incomplete matrix O and unknown but sparse error matrix E .

3.2. Matrix decomposition with missing data

Solving Eq. (8) for N involves recovering and decomposing the incomplete observation matrix O . We use a matrix A for $A = N^T N$, where we know that $\text{rank}(A) = 3$ since $\text{rank}(N) = 3$. Let Ω be a set of indices where $o_{p,q}$ are well-defined in O , and let its complement set be Ω^c . By restricting the error matrix E only to account for the missing entries Ω^c , the original problem of Eq. (8) can be written as

$$\underset{A}{\text{argmin}} \|A - O - E\|_F^2 \quad \text{s.t.} \quad \text{rank}(A) = 3, k_\Omega(E) = 0, \quad (9)$$

where $k_\Omega(\cdot)$ is an operator that only keeps the entries in Ω unchanged and sets others zero. We solve the problem of Eq. (9) by alternatingly estimating A and E . The optimization begins by initializing $E = 0$ and setting the missing entries of O to zeros so that $k_{\Omega^c}(O) = 0$.

At the k -th iteration, we update A_{k+1} by

$$\begin{aligned} USV^T &\leftarrow \text{SVD}(O + E_k), \\ A_{k+1} &\leftarrow U \left(S \begin{bmatrix} I_3 & 0 \\ 0 & 0 \end{bmatrix} \right) V^T = US_{(3)}V^T, \end{aligned} \quad (10)$$

where A_{k+1} is reconstructed using only the first three singular values of S to ensure $\text{rank}(A_{k+1}) = 3$. We then update E_{k+1} by

$$E_{k+1} \leftarrow k_{\Omega^c}(O - A_{k+1}). \quad (11)$$

Eq. (11) assigns values to E_{k+1} only for those entries in Ω^c . This ensures the second constraint of Eq. (9) to hold.

The iteration stops when it converges: $\|E_k\|_F - \|E_{k+1}\|_F \leq \xi \|O\|_F$, where ξ is a small value (set to 10^{-4}). We finally obtain the solution of the surface normals \hat{N} as

$$\hat{N} = S_{(3)}^{\frac{1}{2}} U^T = S_{(3)}^{\frac{1}{2}} V^T. \quad (12)$$

3.3. Concave/Convex ambiguity

Like any other uncalibrated approach, our solution contains ambiguity. In our case, any matrix $Q \in \mathbb{R}^{3 \times 3}$ that satisfies $Q^T Q = I$ can be multiplied with the solution for Eq. (8), without violating the equality:

$$N^T Q^T Q N = (QN)^T (QN) = O + E. \quad (13)$$

Therefore, $QN \in \mathbb{R}^{3 \times P}$ is also a solution to the problem. We state that such ambiguity, which corresponds to rotations and reflections, can be reduced by using the integrability constraint. Belhumeur *et al.* [6] show that if $\text{rank}(N) = 3$, the ambiguity due to a general 3×3 matrix




Figure 6. Surfaces used for synthesis. Notice that these are side view images of three spherical caps. In experiments, the capture direction is from the top.

can be reduced to the GBR ambiguity using the integrability constraint. This is also true in our case since $Q \in \mathbb{R}^{3 \times 3}$. Therefore, by enforcing integrability of N , our original rotation and reflection ambiguity can be reduced to intersect with the GBR ambiguity. As a result, the resulting ambiguity should take the form of a GBR transformation, and also satisfy $Q^T Q = I$ in our case. Then, it must be

$$\tilde{Q} = \frac{1}{\lambda} \begin{bmatrix} 1 & 0 & \mu \\ 0 & 1 & \nu \\ 0 & 0 & \lambda \end{bmatrix} = \frac{1}{\lambda} \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & \lambda \end{bmatrix}, \quad \lambda = \pm 1. \quad (14)$$

This is a binary ambiguity where $\lambda = \pm 1$ corresponds to physically valid convex/concave surfaces that are not distinguishable without light calibration. Therefore, by using integrability constraint, our method recovers surface normals up to only a binary convex/concave ambiguity.

4. Experimental results

We evaluate the proposed method using synthetic and real-world data. The experiments using synthetic data are for making a quantitative evaluation, and the real-world experiments are for making a qualitative assessment.

4.1. Synthetic data

We use three different synthetic surfaces, *i.e.*, a hemisphere, a spherical cap whose surface normals deviate from the viewing direction by $0^\circ \sim 75^\circ$, and another spherical cap with the smaller range of $0^\circ \sim 60^\circ$, as shown in Fig. 6. For each scene, images are synthesized using all 100 materials in the MERL BRDF database. We densely arrange 642 uniform light sources via icosahedron-division on the entire sphere. For each light direction, an image with a resolution of 80×80 is synthesized. It contains about 5000 valid pixels, at which we estimate the surface normals.

Accuracy with known α_m . We first show the results with known linear coefficient α_m values to factor out the effect of estimation of α_m . Table 1 shows the average errors of the 100 materials with the three surfaces. Our method performs better for normals that are less perpendicular to the viewing direction, because the intensity distribution is more stable for these normals. Table 1 also shows the results for some materials on which our method works best. The estimation accuracy is generally high, in particular, it is quite accurate for 50 out of 100 materials in the database.

	Normal range		
	hemisphere	0° ~ 75°	0° ~ 60°
All 100 materials	10.87°	7.25°	5.76°
Best 75 materials	7.79°	4.69°	3.04°
Best 50 materials	5.14°	3.27°	2.00°

Table 1. Recovery errors using the known α_m .

	Normal range		
	hemisphere	0° ~ 75°	0° ~ 60°
All 100 materials	10.86°	7.68°	6.73°
Best 75 materials	7.93°	5.39°	4.23°
Best 50 materials	5.76°	4.25°	2.93°

Table 2. Recovery errors using the estimated α_m .

In addition, we conduct one more experiment using 162 uniform lightings and the hemispherical surface. The average errors are 11.04°, 7.76° and 5.44° for the 100, 75 and 50 materials, respectively. This shows that reducing the illumination number to 162 does not affect the accuracy much.

Accuracy with the estimated α_m . Next we estimate the linear coefficient α_m via the skewness of intensity distributions and perform the whole pipeline. As shown in Table 2, the errors do not change significantly compared with the case in which we know the exact α_m (Table 1). Fig. 7 shows the results of the hemispherical surface scene with 100 BRDFs. As the selected examples show, the accuracy clearly depends on the degree of linearity.

Comparison with other methods. Strictly speaking, it is not easy to find prior methods that can completely handle unknown reflectances and uncalibrated illuminations without additional assumptions to ours. For instance, [26] does not work when occluding boundaries with normals perpendicular to the viewing direction are unavailable. Therefore, we choose the following ones that at least separate the reflectance and illumination factors without knowing the light directions:

1. *SVD* [15]: we implement it as a baseline method that assumes Lambertian reflectance and no shadows.
2. *RPCA* [33]: a state-of-the-art method that robustly handles non-Lambertian components and takes into account shadows.

and provide the ground truth light directions for their disambiguation, so that they will give most ideal results.

Fig. 8 shows the results for the same dataset (hemisphere/100 materials). The SVD method fails for nearly 30 materials. The RPCA method gives good estimates for dozens of materials that contain dominant Lambertian components, but it fails for many others, because the original method requires a shadow mask to be specified. On the




Figure 8. Comparison of methods. Normal recovery errors for all 100 materials using the proposed, the SVD [15], and the RPCA [33] methods. Errors larger than 70° are cut off.

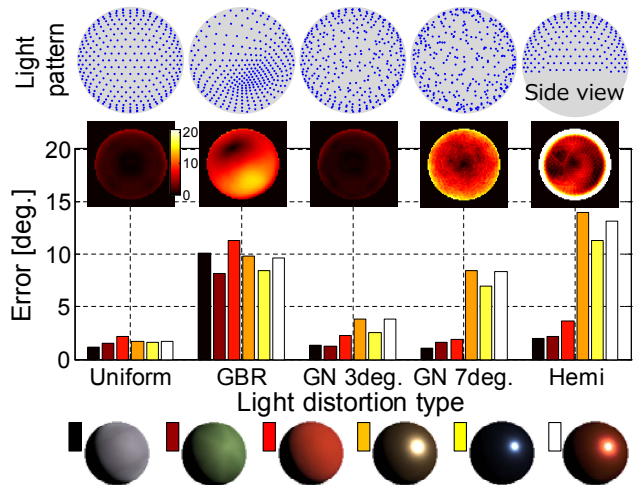


Figure 9. Non-uniform lights. Normal recovery errors under 1) uniform lights, 2) GBR transformed lights, 3) light sources distorted by Gaussian noise with standard deviations of 3° and 7°, and 4) lights from only the upper hemisphere. **Light patterns** and typical **normal error maps** are shown for each case.

other hand, our method works reliably for all 100 materials, even the extreme ones. Although its accuracy is not that high for some easy cases of Lambertian materials, notice that it works in a completely uncalibrated manner.

Non-uniform light sources. We investigate the effects of non-uniform lights beyond our uniform light assumption. We design four representative non-uniform light patterns, by applying the GBR transformation, adding Gaussian noise, and only using the upper lights. These lights and results are shown in Fig. 9, where each color of the error bars indicates one material. First, GBR distortion causes errors for all materials, and the error map has a similar distortion pattern. Second, diffuse surfaces are robust against random noise and hemi-lights, while specular surfaces are more affected. This is because specular materials have rapidly changing intensities, as shown in Fig. 4, and thus, the calculation is sensitive to light position biases. In particular, using only upper lights causes large errors especially near the occluding boundaries for specular materials.

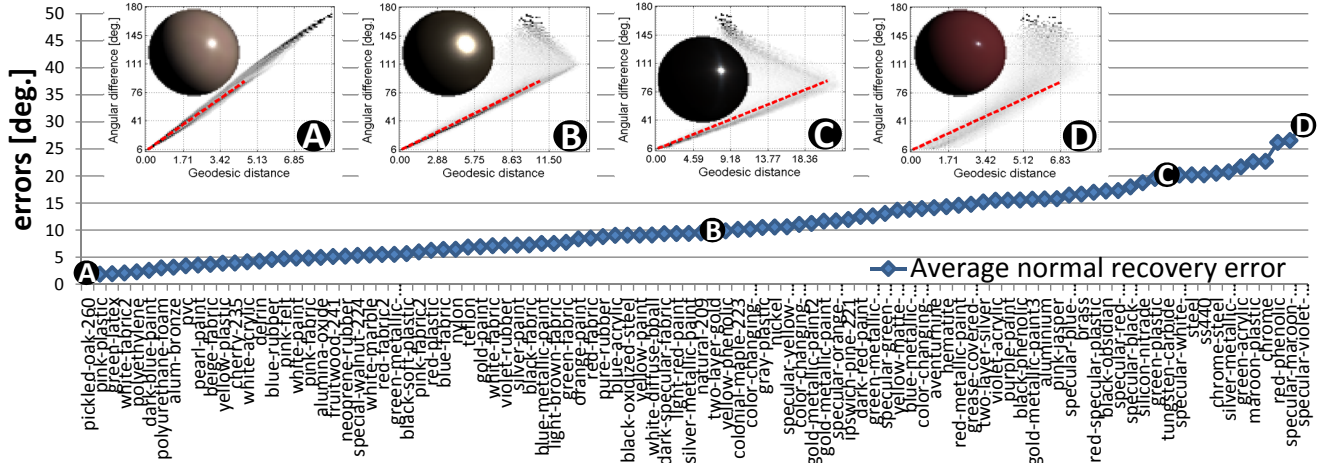


Figure 7. Average normal recovery errors for all 100 materials. The materials are listed by ranking their corresponding recovery errors.

4.2. Real-world data

Images for real-world objects are captured at a resolution of 640×480 , by rotating the light source around the objects to generate about 150 light directions. Some objects with large concavity suffer from severe cast shadows produced by sidelights and backlights. To avoid this, we use lights mainly from the upper hemisphere for these objects. The results are summarized in Fig. 10. As the results show, different colors/albedos on the same object are well handled, because of the normalization in Eq. (1) and the fact that difference in albedos does not essentially change the linear relation and its linear coefficient. We make a quantitative evaluation for a metallic hemisphere (the leftmost case) of which we know the exact shape, and the average angular error is 3.45° . A few scenes on the right do not have occluding boundaries, while we can still estimate their surface normals purely from the images of the surface patches. This shows the advantage of the proposed method over the prior methods [26, 24]. In all results, there exists a convex/concave ambiguity as discussed. We resolve the ambiguity by manually selecting one as our ambiguity is binary.

5. Conclusion

We present a photometric stereo technique that recovers surface normals with unknown real-world reflectances in an uncalibrated manner. We have shown that the information extracted from the pixel intensity profiles across images offers a strong cue for solving the problem. Our method solves the problem up to only a binary convex/concave ambiguity, and the effectiveness of our method has been shown by testing it on both synthetic and real-world data.

One limitation is that our method currently assumes a scene to have a uniform (up to albedo differences) reflectance. Relaxing this assumption will enhance the applicability of the proposed method. Since our method can

naturally distinguish different reflectances and recover surface normals accordingly, our next goal is to fully exploit such ability to deal with surfaces composed of more complex spatially-variant reflectances.

References

- [1] N. Alldrin and D. Kriegman. Toward reconstructing surfaces with arbitrary isotropic reflectance : A stratified photometric stereo approach. In *Proc. of Int'l Conf. on Computer Vision (ICCV)*, pages 1–8, 2007. 2
- [2] N. Alldrin, S. Mallick, and D. Kriegman. Resolving the generalized bas-relief ambiguity by entropy minimization. In *Proc. of IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2007. 2
- [3] N. Alldrin, T. Zickler, and D. Kriegman. Photometric stereo with non-parametric and spatially-varying reflectance. In *Proc. of IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 1–8, 2008. 2
- [4] S. Barsky and M. Petrou. The 4-source photometric stereo technique for three-dimensional surfaces in the presence of highlights and shadows. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 25(10):1239–1252, 2003. 2
- [5] R. Basri, D. Jacobs, and I. Kemelmacher. Photometric stereo with general, unknown lighting. *Int'l Journal of Computer Vision*, 72(3):239–257, 2007. 2
- [6] P. N. Belhumeur, D. J. Kriegman, and A. L. Yuille. The bas-relief ambiguity. *Int'l Journal of Computer Vision*, 35(1):33–44, 1999. 2, 5
- [7] M. Chandraker, J. Bai, and R. Ramamoorthi. A theory of differential photometric stereo for unknown isotropic brdfs. In *Proc. of IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 2505–2512, 2011. 2
- [8] M. Chandraker, F. Kahl, and D. Kriegman. Reflections on the generalized bas-relief ambiguity. In *Proc. of IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 788–795, 2005. 2
- [9] H. Chung and J. Jia. Efficient photometric stereo on glossy surfaces with wide specular lobes. In *Proc. of Int'l Conf. on Computer Vision (ICCV)*, pages 1–8, 2008. 2
- [10] O. Drbohlav and M. Chaniler. Can two specular pixels calibrate photometric stereo? In *Proc. of Int'l Conf. on Computer Vision (ICCV)*, pages 1850–1857, 2005. 2
- [11] O. Drbohlav and R. Šára. Specularities reduce ambiguity of uncalibrated photometric stereo. In *Proc. of European Conf. on Computer Vision (ECCV)*, pages 644–645, 2002. 2




Figure 10. Real-world results. Rows: captured images; color-coded normal maps; depth maps via [20]; and two views of rendered surfaces.

- [12] P. Favaro and T. Papadimitri. A closed-form solution to uncalibrated photometric stereo via diffuse maxima. In *Proc. of IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 821–828, 2012. 2
- [13] A. Georghiades. Incorporating the torrance and sparrow model of reflectance in uncalibrated photometric stereo. In *Proc. of Int'l Conf. on Computer Vision (ICCV)*, pages 816–823, 2003. 2
- [14] D. Goldman, B. Curless, A. Hertzmann, and S. Seitz. Shape and spatially-varying brdfs from photometric stereo. In *Proc. of Int'l Conf. on Computer Vision (ICCV)*, pages 341–348, 2005. 2
- [15] H. Hayakawa. Photometric stereo under a light source with arbitrary motion. *JOSA A*, 11(11):3079–3089, 1994. 6
- [16] A. Hertzmann and S. Seitz. Example-based photometric stereo: shape reconstruction with general, varying BRDFs. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 27(8):1254–1264, 2005. 1, 2
- [17] T. Higo, Y. Matsushita, and K. Ikeuchi. Consensus photometric stereo. In *Proc. of IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 1157–1164, 2010. 2
- [18] M. Holroyd, J. Lawrence, G. Humphreys, and T. Zickler. A photometric approach for estimating normals and tangents. *ACM Transactions on Graphics*, 27:133, 2008. 2
- [19] S. Koppal and S. Narasimhan. Clustering appearance for scene analysis. In *Proc. of IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 1323–1330, 2006. 1, 2, 3
- [20] P. Kovesi. Shapelets correlated with surface normals produce surfaces. In *Proc. of Int'l Conf. on Computer Vision (ICCV)*, pages 994–1001, 2005. 8
- [21] W. Matusik, H. Pfister, M. Brand, and L. McMillan. A data-driven reflectance model. In *Proc. of ACM SIGGRAPH*, pages 27–31, 2003. 2, 3
- [22] D. Miyazaki, K. Hara, and K. Ikeuchi. Median photometric stereo as applied to the segonko tumulus and museum objects. *Int'l Journal of Computer Vision*, 86(2):229–242, 2010. 2
- [23] 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, 1990. 2
- [24] T. Okabe, I. Sato, and Y. Sato. Attached shadow coding: Estimating surface normals from shadows under unknown reflectance and lighting conditions. In *Proc. of Int'l Conf. on Computer Vision (ICCV)*, pages 1693–1700, 2009. 2, 7
- [25] G. Oxholm, P. Bariya, and K. Nishino. The scale of geometric texture. In *Proc. of European Conf. on Computer Vision (ECCV)*, 2012. 1
- [26] I. Sato, T. Okabe, Q. Yu, and Y. Sato. Shape reconstruction based on similarity in radiance changes under varying illumination. In *Proc. of Int'l Conf. on Computer Vision (ICCV)*, pages 1–8, 2007. 1, 2, 3, 6, 7
- [27] B. Shi, Y. Matsushita, Y. Wei, C. Xu, and P. Tan. Self-calibrating photometric stereo. In *Proc. of IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 1118–1125, 2010. 2
- [28] B. Shi, P. Tan, Y. Matsushita, and K. Ikeuchi. Elevation angle from reflectance monotonicity: Photometric stereo for general isotropic reflectances. In *Proc. of European Conf. on Computer Vision (ECCV)*, 2012. 2
- [29] W. Silver. Determining shape and reflectance using multiple images. Master's thesis, M.I.T., Dept. of Electrical Engineering and Computer Science, 1980. 2
- [30] P. Tan, L. Quan, and T. Zickler. The geometry of reflectance symmetries. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 33(12):2506–2520, 2011. 2
- [31] J. Tenenbaum, V. De Silva, and J. Langford. A global geometric framework for nonlinear dimensionality reduction. *Science*, 290(5500):2319–2323, 2000. 3
- [32] R. Woodham. photometric method for determining surface orientation from multiple images. *Optical Engineering*, 1(7):139–144, 1980. 1
- [33] L. Wu, A. Ganesh, B. Shi, Y. Matsushita, Y. Wang, and Y. Ma. Robust photometric stereo via low-rank matrix completion and recovery. In *Proc. of Asian Conf. on Computer Vision (ACCV)*, pages 703–717, 2010. 2, 6
- [34] T. Wu, K. Tang, C. Tang, and T. Wong. Dense photometric stereo: A markov random field approach. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 28(11):1830–1846, 2006. 2
- [35] Z. Zhou and P. Tan. Ring-light photometric stereo. In *Proc. of European Conf. on Computer Vision (ECCV)*, pages 265–279, 2010. 2