

Symmetric-light Photometric Stereo

Kazuma Minami Hiroaki Santo Fumio Okura Yasuyuki Matsushita
Graduate School of Information Science and Technology, Osaka University

{kazuma.minami,santo.hiroaki,okura,yasumat}@ist.osaka-u.ac.jp

Abstract

This paper presents symmetric-light photometric stereo for surface normal estimation, in which directional lights are distributed symmetrically with respect to the optic center. Unlike previous studies of ring-light settings that required the information of ring radius, we show that even without the knowledge of the exact light source locations or their distances from the optic center, the symmetric configuration provides us sufficient information for recovering unique surface normals without ambiguity. Specifically, under the symmetric lights, measurements of a pair of scene points having distinct surface normals but the same albedo yield a system of constrained quadratic equations about the surface normal, which has a unique solution. Experiments demonstrate that the proposed method alleviates the need for geometric light source calibration while maintaining the accuracy of calibrated photometric stereo.

1. Introduction

Photometric stereo estimates surface normals from a set of measurements captured under varying light directions [23, 21]. Because of its capability of estimating the surface normal at each pixel, photometric stereo is useful for recovering a detailed shape of real-world scenes. Conventional photometric stereo methods [8, 17], including recent learning-based methods [14, 12, 7], assume *known* light directions obtained by geometric light source calibration [26, 15] for determining surface normals without ambiguity. To eliminate the need for the light source calibration, uncalibrated photometric stereo [10, 5] has been studied. However, their surface normal estimates generally suffer from ambiguities [2] due to the lack of knowledge about the light source directions. More recently, a ring-light configuration [27, 20, 3], in which point light sources are located on a circle centered at the camera, has been introduced by exploiting viable constraints in light configurations. While the ring-light configuration is useful for better constraining the problem, their estimates still suffer from ambiguities [27].

In this paper, we show that the subset of ring-light con-




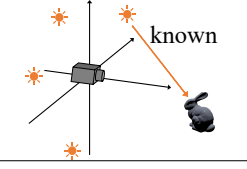
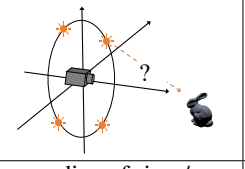
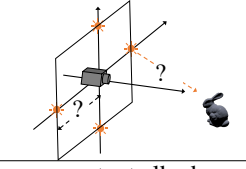
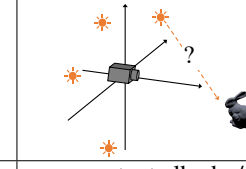
Figure 1. Symmetric- and ring-light configurations. The planes on which point light sources are placed are parallel to the image plane and aligned to the optic center.

figuration, namely, a *symmetric-light configuration* (Fig. 1), can sufficiently constrain the problem by exploiting the symmetry of the lights, yielding unique surface normal estimates without ambiguity. In the symmetric-light configuration, two sets of two point light sources are symmetrically placed w.r.t. the x - and y - axes of the camera in the same interval, respectively. We assume that these four point light sources are placed infinitely far away from a scene, resulting in symmetric directional lightings. The differences of the symmetric measurements produce independent constraints over x -, y -, and z - components of surface normals. By taking two scene points having distinct surface normals but the same albedo, our method creates a system of constrained quadratic equations about the surface normals that has a unique solution. The experiments demonstrate that the proposed method achieves comparable accuracy to the calibrated photometric stereo even without the need for directional light source calibration.

The chief contributions of this paper are twofold:

1. We introduce symmetric-light photometric stereo, in which the light symmetry is exploited for constraining the problem, for determining surface normals.
2. By being a subset of ring-light configuration, the symmetric-light configuration is simple and easy to set up and avoids explicit geometric light source calibration, while it achieves comparable accuracy to fully

Table 1. Summary of (a) calibrated, (b) ring-light, (c) symmetric-light, and (d) uncalibrated photometric stereo.

	(a) calibrated [23]	(b) ring-light [27, 20]	(c) symmetric-light (ours)	(d) uncalibrated [10, 18]
Light sources				
Unique solution with	-	radius of ring / integrability	constant albedo (two pixels)	constant albedo / integrability

calibrated photometric stereo.

2. Related works

We briefly review previous photometric stereo methods that do not require explicit geometric light source calibration. This direction has been studied as uncalibrated photometric stereo [10] (Table 1 (d)), which assumes unknown light directions and aims at simultaneously estimating both surface normals and light directions. Hayakawa [10] introduced a method based on matrix factorization and showed there existed a linear ambiguity in the surface normal estimates. Later, it is shown that surface integrability [11] can reduce the linear ambiguity to a three-parameter linear transform, which is known as the generalized bas-relief (GBR) ambiguity [2, 25]. To resolve the GBR ambiguity, various external assumptions, such as albedo priors [18], specular spike [6], mutual reflections [4], a certain type of parametric reflectance models [9], symmetry of reflectance [22, 24], have been exploited.

It has also been studied to constrain the problem of uncalibrated photometric stereo by putting assumptions on light source distributions. Zhou and Tan [27] have proposed a ring-light configuration (Table 1 (b)), in which the point light sources are located on a circle centered on the camera. They show that the constraint from the ring lights can reduce the ambiguity of uncalibrated photometric stereo up to mirror transformation, rotation, and scaling ambiguities. To resolve these ambiguities, they use (1) known rotation and radius of ring lights with equal intensities or intervals, (2) the assumption of surface integrability and scene points with the same albedo, or (3) multi-view information. Shiradkar *et al.* [20] also show that ring lights that have an equal interval and known radius but unknown rotation still exhibit scaling and rotation ambiguities, which may be resolved by multi-view information. Chandraker *et al.* [3] use denser ring lights realized by rotating a point light source around the camera. From the differences of the measurements when slightly rotating the light source, their method can estimate the surface normals of a scene with unknown but isotropic reflectances.

As summarized in Table 1, our method is in between calibrated and uncalibrated photometric stereo. The

symmetric-light setting (Table 1 (c)) can be regarded as a subset of ring lights; thus the setup is simpler than the ring-light setting. By exploiting the new constraints yielded from the light symmetry, we show that the proposed method can uniquely determine the surface normals without ambiguity, even without the knowledge of the radius, assumption of surface integrability, and multi-view information that is needed in previous approaches.

3. Proposed method

From four images captured under symmetric lights, our photometric stereo uses a pixel pair having the same albedo to estimate the surface normals for these pixels. In what follows, we describe our formulation and the solution method starting from the image formation model. We here assume the appropriate pixel pairs are given, while we later introduce a practical method to select reasonable pixel pairs. Figure 2 shows the overview of the proposed symmetric-light photometric stereo.

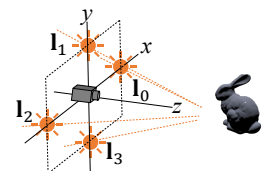
3.1. Image formation model

Suppose that a Lambertian surface is illuminated by directional lights and observed by a fixed camera. The measurement at a pixel $m^{(s)} \in \mathbb{R}_+$ under s -th light source can be described by

$$m^{(s)} = \rho \mathbf{n}^\top \mathbf{l}_s, \quad (1)$$

where $\mathbf{n} \in \mathcal{S}^2 \subset \mathbb{R}^3$ and $\rho \in \mathbb{R}_+$ are the surface normal and albedo, and $\mathbf{l}_s \in \mathcal{S}^2$ is the s -th light direction.

The proposed method uses four light sources that are symmetrically placed and have a uniform intensity. The light directions are written as



$$\begin{aligned} \mathbf{l}_0 &= [l_x, 0, l_z]^\top, \\ \mathbf{l}_1 &= [0, l_y, l_z]^\top, \\ \mathbf{l}_2 &= [-l_x, 0, l_z]^\top, \\ \mathbf{l}_3 &= [0, -l_y, l_z]^\top, \end{aligned} \quad (2)$$

where $l_x, l_y, l_z > 0$ are unknown. The light sources can rotate around the z -axis; however, for the simplicity of notation, we explain the method under this setting. The measurements under these four light directions are described by

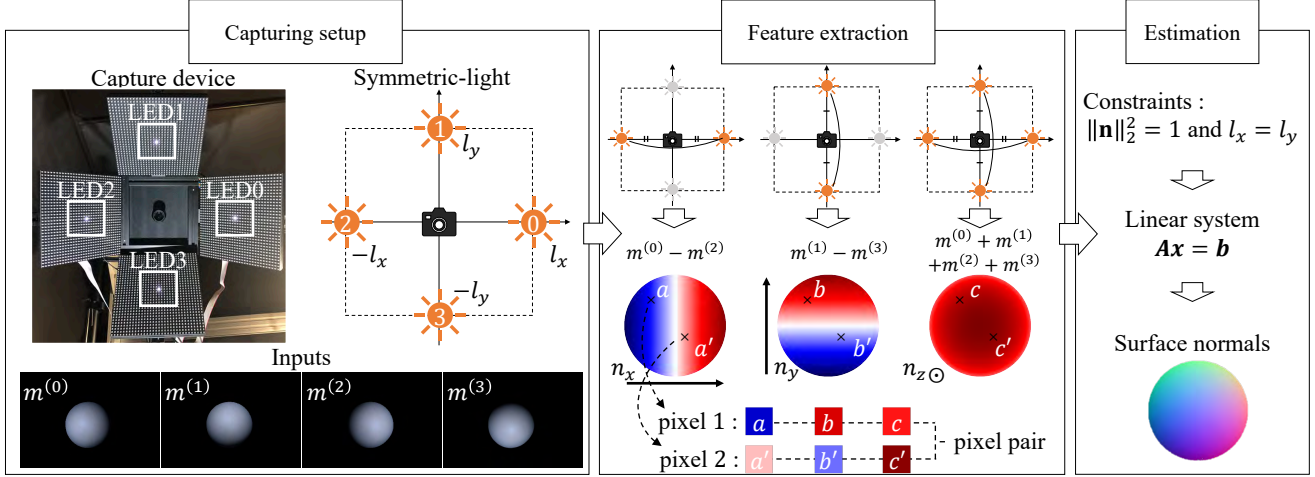


Figure 2. Overview of the proposed method. We capture four images under the symmetric-light configuration and extract features using the symmetry of the measurements. We construct a linear system from the features and constraints and solve it to estimate the surface normals.

Eq. (1) using the unit surface normal $\mathbf{n} = [n_x, n_y, n_z]^\top$ as

$$\begin{cases} m^{(0)} = \rho \mathbf{n}^\top \mathbf{l}_0 = \rho n_x l_x + n_z l_z, \\ m^{(1)} = \rho \mathbf{n}^\top \mathbf{l}_1 = \rho n_y l_y + n_z l_z, \\ m^{(2)} = \rho \mathbf{n}^\top \mathbf{l}_2 = -\rho n_x l_x + n_z l_z, \\ m^{(3)} = \rho \mathbf{n}^\top \mathbf{l}_3 = -\rho n_y l_y + n_z l_z. \end{cases} \quad (3)$$

3.2. Solution method

From the symmetry of the light directions, we obtain the features a , b , and c by linearly combining the measurements in Eq. (3) as:

$$\begin{aligned} \frac{m^{(0)} - m^{(2)}}{2} &= \rho n_x l_x \equiv a, \\ \frac{m^{(1)} - m^{(3)}}{2} &= \rho n_y l_y \equiv b, \\ \frac{m^{(0)} + m^{(1)} + m^{(2)} + m^{(3)}}{4} &= \rho n_z l_z \equiv c. \end{aligned} \quad (4)$$

These features represent the relative magnitudes for each component of the surface normals, n_x , n_y , and n_z , as illustrated in Fig. 2.

Our method uses a pair of scene points (pixels) having distinct surface normals but the same albedo. Under the assumption of directional lights, all the scene points share the same light directions. We compute the ratios of features in Eq. (4) from a pixel pair with the same albedo $\rho = \rho'$:

$$\begin{aligned} \left(\frac{a}{a'}, \frac{b}{b'}, \frac{c}{c'} \right) &= \left(\frac{\rho n_x l_x}{\rho' n'_x l_x}, \frac{\rho n_y l_y}{\rho' n'_y l_y}, \frac{\rho n_z l_z}{\rho' n'_z l_z} \right) \\ &= \left(\frac{n_x}{n'_x}, \frac{n_y}{n'_y}, \frac{n_z}{n'_z} \right), \end{aligned} \quad (5)$$

where prime indicates the quantities of the other pixel in the pair. Additionally, by assuming that light sources are placed

on a circle, *i.e.*, $l_x = l_y$, we obtain the following constraints:

$$\left(\frac{b}{a}, \frac{b'}{a'} \right) = \left(\frac{n_y l_y}{n_x l_x}, \frac{n'_y l_y}{n'_x l_x} \right) = \left(\frac{n_y}{n_x}, \frac{n'_y}{n'_x} \right). \quad (6)$$

By squaring both sides of Eqs. (5) and (6) and plugging in the unit norm constraints, $\|\mathbf{n}\|_2^2 = 1$ and $\|\mathbf{n}'\|_2^2 = 1$, we have a system of quadratic equations as

$$\underbrace{\begin{bmatrix} a'^2 & 0 & 0 & -a^2 & 0 & 0 \\ 0 & b'^2 & 0 & 0 & -b^2 & 0 \\ 0 & 0 & c'^2 & 0 & 0 & -c^2 \\ b^2 & -a^2 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & b'^2 & -a'^2 & 0 \\ 1 & 1 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 1 & 1 \end{bmatrix}}_{\mathbf{A}} \underbrace{\begin{bmatrix} n_x^2 \\ n_y^2 \\ n_z^2 \\ n_x'^2 \\ n_y'^2 \\ n_z'^2 \end{bmatrix}}_{\mathbf{x}} = \underbrace{\begin{bmatrix} 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 1 \\ 1 \end{bmatrix}}_{\mathbf{b}}. \quad (7)$$

Our method treats the squares of unknowns $[n_x^2, n_y^2, \dots, n_z'^2]^\top$ as a new variable $\mathbf{x} (\geq \mathbf{0})$. The least squares approximate solution \mathbf{x}^* can be derived by minimizing the objective function:

$$\mathbf{x}^* = \underset{\mathbf{x}}{\operatorname{argmin}} \|\mathbf{A}\mathbf{x} - \mathbf{b}\|_2^2 \quad \text{s.t.} \quad \mathbf{x} \geq \mathbf{0}. \quad (8)$$

We use a method of non-negative least squares [13] for obtaining the approximate solution \mathbf{x}^* .

Because $l_x, l_y, l_z, \rho, c, c' > 0$ for any pixel in Eq. (4), the following conditions should also be satisfied:

$$\begin{cases} n_x \geq 0 & (\text{if } a \geq 0) \\ n_x < 0 & (\text{otherwise}) \end{cases}, \begin{cases} n_y \geq 0 & (\text{if } b \geq 0) \\ n_y < 0 & (\text{otherwise}) \end{cases}, n_z \geq 0 \quad (9)$$

Using the constraints in Eq. (9), we can uniquely determine the surface normals $\mathbf{n} = [n_x, n_y, n_z]^\top$ and $\mathbf{n}' = [n'_x, n'_y, n'_z]^\top$ from the estimated \mathbf{x} .

3.3. Solvability condition

For ensuring that the objective Eq. (8) has a unique solution, the coefficient matrix \mathbf{A} in Eq. (7) should be full rank. This implies that the two pixels should have distinct surface normals in terms of the elevation angles.

Let θ and ϕ as the elevation and azimuth angles of a surface normal, respectively, and suppose that the two pixels have the same elevation angle, *i.e.*, $\theta = \theta'$. With the angular expression of surface normals $\mathbf{n} = [n_x, n_y, n_z]^\top = [\sin \theta \cos \phi, \sin \theta \sin \phi, \cos \theta]^\top$, Eqs. (5) and (6) are written as

$$\begin{cases} \frac{a}{a'} = \frac{\cos \phi}{\cos \phi'} \\ \frac{b}{b'} = \frac{\sin \phi}{\sin \phi'} \\ \frac{c}{c'} = 1 \end{cases}, \quad \begin{cases} \frac{b}{a} = \frac{\sin \phi}{\cos \phi} \\ \frac{b'}{a'} = \frac{\sin \phi'}{\cos \phi'} \end{cases}. \quad (10)$$

Equations in (Eq. (10)) are independent of the elevation angles θ and θ' , and it indicates that we cannot determine the unique solution under the condition of $\theta = \theta'$. Indeed, this degenerate condition introduces rank deficiency in the coefficient matrix \mathbf{A} in Eq. (7).

To sum up, for uniquely determining surface normals, the proposed method requires a pixel pair that has (1) the same albedo, *i.e.*, $\rho = \rho'$ and (2) distinct surface normals in terms of the elevation angles.

4. Pixel-pair selection

Our method uses a pair of scene points with the same albedo but distinct elevation angles of their surface normals. For stably selecting such pixel pairs, we develop a simple yet effective method based on albedo and surface normal clustering.

Albedo clustering To select pixels having the same albedo, we use a chromaticity-based clustering [18] to approximate the albedo clustering. From the four RGB measurements under our symmetric-light setting for each pixel, we compute the mean RGB measurements $(\bar{m}^{(R)}, \bar{m}^{(G)}, \bar{m}^{(B)})$. Using the mean RGB measurements, we compute the chromaticity as

$$\left(\frac{\bar{m}^{(R)}}{\bar{m}^{(R)} + \bar{m}^{(G)} + \bar{m}^{(B)}}, \frac{\bar{m}^{(G)}}{\bar{m}^{(R)} + \bar{m}^{(G)} + \bar{m}^{(B)}} \right).$$

Once the albedos are clustered, our method selects a pixel pair that belongs to the same cluster.

Surface normal clustering To select a pixel pair having distinct surface normals in terms of their elevation angles, we cluster surface normals based on their elevation angles. While the elevation angles are unknown, let us suppose that two pixels have elevation angles θ and $\theta' \in [0, \frac{\pi}{2}]$. We use a metric r for assessing the difference of the elevation angles defined as

$$r = |\tan^2 \theta - \tan^2 \theta'|. \quad (11)$$

From Eq. (4), we have

$$(n_x, n_y, n_z) = \left(\frac{a}{\rho l_x}, \frac{b}{\rho l_y}, \frac{c}{\rho l_z} \right). \quad (12)$$

Substituting Eq. (12) into the metric r in Eq. (11), we have

$$\begin{aligned} r &= |\tan^2 \theta - \tan^2 \theta'| = \left| \frac{1 - \cos^2 \theta}{\cos^2 \theta} - \frac{1 - \cos^2 \theta'}{\cos^2 \theta'} \right| \\ &= \left| \frac{1 - n_z^2}{n_z^2} - \frac{1 - n_z'^2}{n_z'^2} \right| = \left| \frac{n_x^2 + n_y^2}{n_z^2} - \frac{n_x'^2 + n_y'^2}{n_z'^2} \right| \\ &= \frac{l_z^2}{l_x^2} \left| \frac{a^2 + b^2}{c^2} - \frac{a'^2 + b'^2}{c'^2} \right|, \end{aligned}$$

with the assumptions that $l_x = l_y$ and $\rho = \rho'$. Since the light directions l_x and l_z are constant across a scene, we can define a distance $\hat{r} = \left| \frac{a^2 + b^2}{c^2} - \frac{a'^2 + b'^2}{c'^2} \right|$ that is free from the light directions. Using the metric \hat{r} , we use K-means clustering to obtain clusters of similar surface normals in terms of their elevation angles. Once the clusters are obtained, a pixel pair is selected by taking pixels from different clusters.

For robustifying the pixel selection, we avoid selecting pixels from clusters that have large intra-cluster variances of the distance \hat{r} . Namely, we avoid clusters with top 50 percentile intra-cluster variances and randomly sample from the remaining clusters. In our implementation, we sample 100 pixel pairs and take the median of those estimates by vector median filter [1].

5. Experiment

We evaluate the proposed method using synthetic and real-world data. For comparison, we use a calibrated photometric stereo method by Woodham [23], denoted as ‘‘Calibrated.’’ The estimation accuracy by the calibrated method can be considered as the upper bound of the accuracy, as our method only uses less information about the light directions. For the real-world data, we use the DiLiGenT dataset [19], which is a benchmark dataset for calibrated photometric stereo, as well as the data captured by our setup. For the albedo and surface normal clustering, the numbers of clusters are set to 20 and 50, respectively. Since the proposed method is designed for grayscale measurements, the color measurements are first converted to grayscale in our implementation.

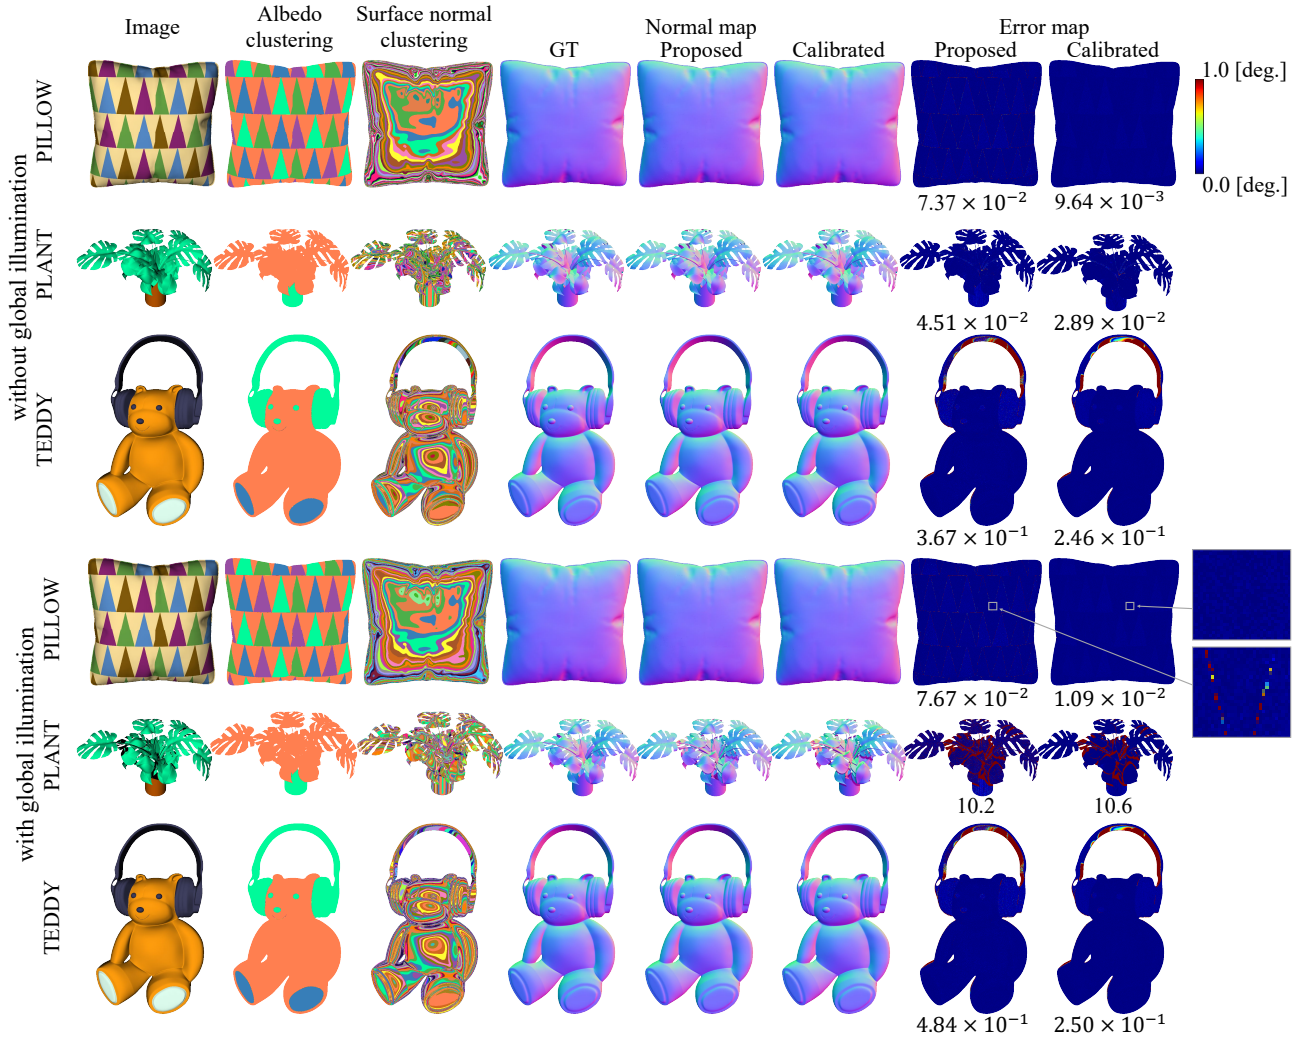


Figure 3. Estimation results for our synthetic dataset. For each row, we show one of the input images, results of albedo and surface normal clustering, the ground-truth and estimated normal maps for each method, and corresponding error maps. For each error map, we show the mean angular errors in degrees. For each scene, we have two rendering conditions: with and without global illumination effects.

5.1. Evaluation with synthetic data

We generate our synthetic dataset by Blender¹ using the scenes PILLOW, TEDDY, and PLANT² rendered with the Lambertian reflectance. The scenes are rendered with and without the global illumination effect, *i.e.*, cast shadow and interreflections, to assess the robustness of the methods. As for the light directions, we set the elevation angles to 80 degrees, *i.e.*, $l_x = l_y = 0.1736$, $l_z = 0.9848$.

Figure 3 shows the visual results for our synthetic dataset along with the mean angular errors of the estimated normal maps in degrees. The proposed method shows almost com-

¹Blender 2.83 LTS, <https://www.blender.org>, last accessed on Aug. 17, 2021

²CGTrader, <https://www.cgtrader.com>, last accessed on Aug. 17, 2021

parable accuracy with the calibrated photometric stereo, while for some objects with slightly larger errors. In the PILLOW scene, the proposed method has a relatively large error at the boundary of the surface texture due to the error of the albedo clustering. On the other hand, the PLANT scene shows the applicability of the proposed method for scenes that do not satisfy the surface integrability.

5.2. Evaluation with real-world data

Figure 4 shows the results for the DiLiGenT dataset. Since the DiLiGenT dataset uses 96 light directions roughly arranged in a grid manner, we choose four light directions³ that approximately satisfy the symmetric-light assumption,

³The results with different choices of the light directions are shown in the supplementary material.

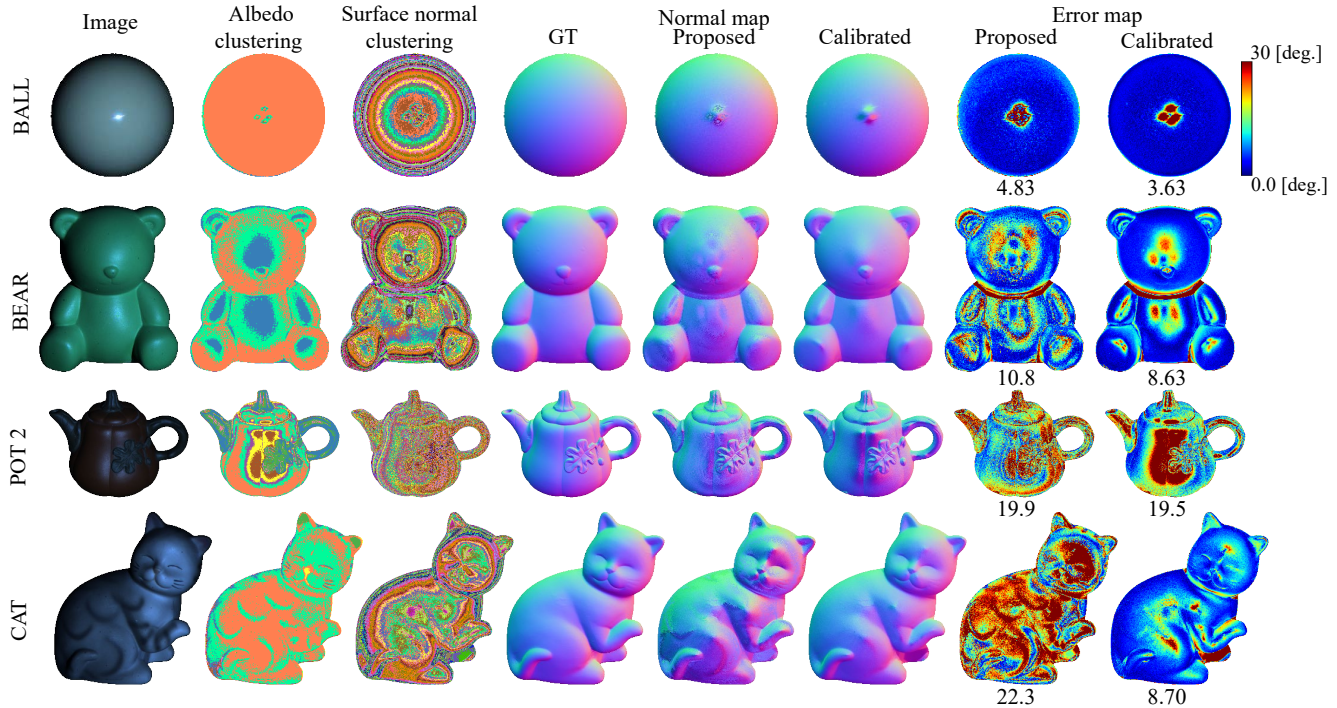


Figure 4. Estimation results for the DiLiGenT dataset. For each row, we show one of the input images, results of albedo and surface normal clustering, the ground-truth and estimated normal maps for each method, and corresponding error maps. For each error map, we show the mean angular errors in degrees.





Figure 5. Our capturing setup. We use four symmetric LEDs for our method, and 100 calibrated LEDs for obtaining the ground truth of the surface normals.



as shown in the left figure. Since we focus on the Lambertian-like scenes, we use BALL, BEAR, POT2, and CAT scenes out of ten scenes in the dataset.

Although the light configuration does not strictly satisfy our

symmetric-light assumption (discussed in Sec. 6 later), for the BALL, BEAR, and POT2 scenes, the proposed method still shows almost comparable accuracy to the calibrated method even though the proposed method does not use the information of the light source calibration. For the CAT scene, the proposed method has a larger error compared to the calibrated method, due to the failure of the albedo clustering. The detail of the failure mode will be discussed in Sec. 6.

Next, we show the results of scenes captured by our setup with the symmetric-light configuration in Fig. 5. We use a CCD camera (FLIR BFS-U3-28S5-C) with the image resolution of 1464×1936 pixels with 12 bit color depth and a lens with the focal length of 25 mm. To obtain the ground truth, we use 100 LEDs on the LED arrays, which are calibrated [16]. We use objects with simple (SPHERE and KITTEN) and complicated (LIZARD) shapes for scenes with uniform albedo. SHEEP and CHALK scenes contain spatially-varying albedos. Figure 6 shows the estimated results for our dataset. Overall, the proposed method achieves comparable results to the calibrated method. For the CHALK scene, the proposed method exhibits a larger error compared to the calibrated method in the concave regions since these regions are clustered as different albedo due to cast shadows. In contrast, the other regions with spatially-varying albedos are correctly clustered and surface

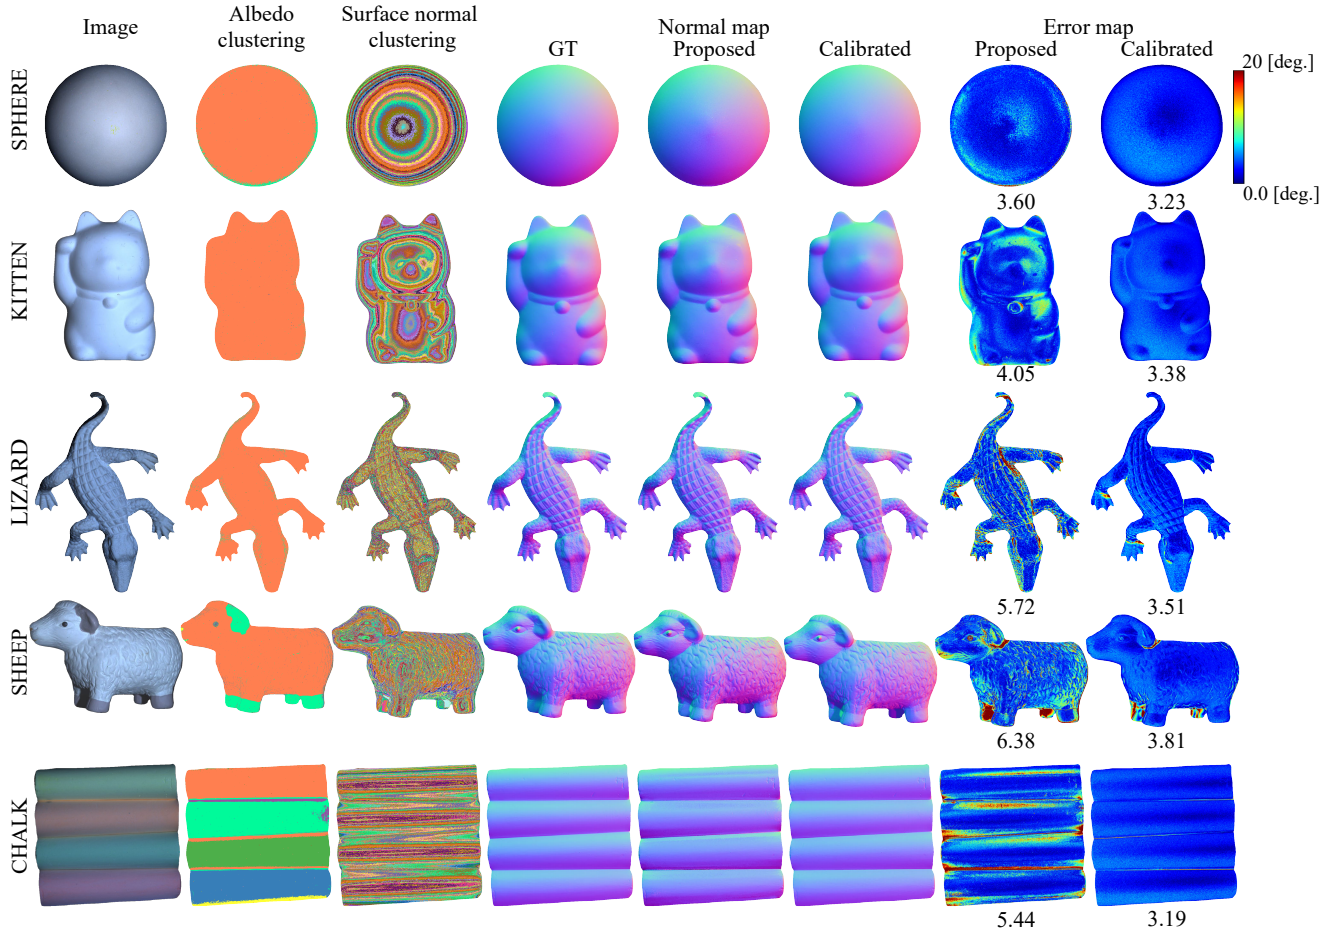


Figure 6. Estimation results for our dataset. For each row, we show one of the input images, results of albedo and surface normal clustering, the ground-truth and estimated surface normal maps for each method, and corresponding error maps. For each error map, we show the mean angular errors in degrees.

normals are accurately estimated.

6. Discussions

This paper studies the conditions derived from the symmetric-light configuration for surface normal recovery. As a result, we present the symmetric-light photometric stereo, in which directional lights are symmetrically distributed w.r.t. the camera. Using a pixel pair that has the same albedo but different elevation angles of the surface normals, the proposed method can uniquely determine the surface normal without ambiguities. Unlike previous methods that use a ring-light, the symmetric-light alleviates the need for assuming a known radius of lights and surface integrability. Our experiments demonstrate that the proposed method works as good as calibrated photometric stereo even without the need for the geometric light source calibration.

For improving the surface normal estimation accuracy, there are a few venues that can be investigated. Here we

discuss the source of errors in albedo-clustering, surface normal clustering, and the symmetric-light configuration.

Albedo clustering For albedo clustering, our method uses chromaticities as an approximation of albedos, because the true albedo is inaccessible. To observe the error caused by this approximation, we compute true albedos from the ground truth surface normal, apply the same clustering method, and assess the surface normal estimation accuracy. Figure 7 shows the estimated results with and without true albedos. This result indicates that the proposed method can be improved by a better albedo approximation and albedo clustering.

Surface normal clustering We assess the effect of surface normal clustering to the surface normal estimation accuracy. Figure 8 shows the results of the synthetic scenes, PLANT and CUBE, with and without surface normal clus-

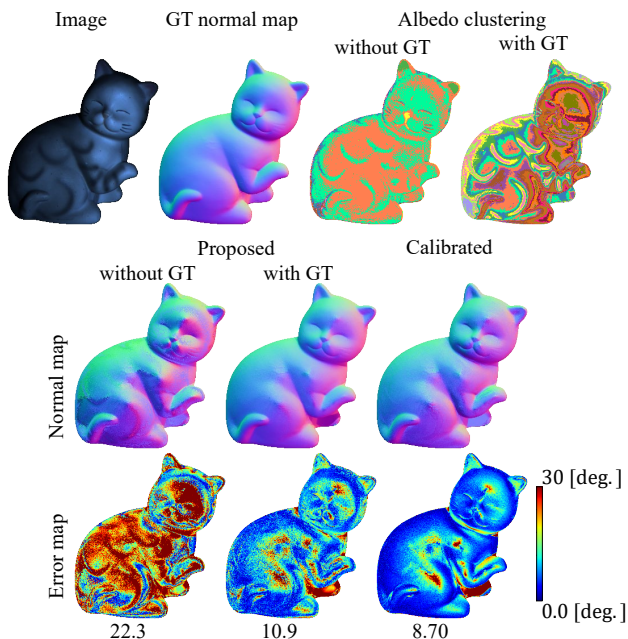


Figure 7. Comparison of surface normal estimation accuracies with and without true albedos (denoted as GT) for albedo clustering. The numbers below the error maps indicate the mean angular errors in degrees.

tering. Since PLANT scene has a complex shape and various surface normals, even without surface normal clustering, we could obtain a reasonable result by random sampling, while surface normal clustering still improves the accuracy. On the other hand, the CUBE scene has only three distinct surface normals, resulting in a larger error without surface normal clustering. It indicates that accurate surface normal clustering contributes to a better surface normal estimation accuracy.

Deviation of symmetric-light configuration We evaluate the effect of deviations of light directions from the symmetric-light assumption in the surface normal estimation. Figure 9 shows the results of the synthetic experiments with adding noise to the light directions. We use the same settings in Sec. 5.1 and render the scenes without the global illumination effects. The Gaussian noise with varying the noise level is added to the elevation and azimuth angles of the light directions, respectively. The results indicate that the accuracy remains comparable to the calibrated method even when the light directions are off from the symmetric-light assumption.

Limitations One of the limitations of the proposed method is a distant light assumption. To make photomet-

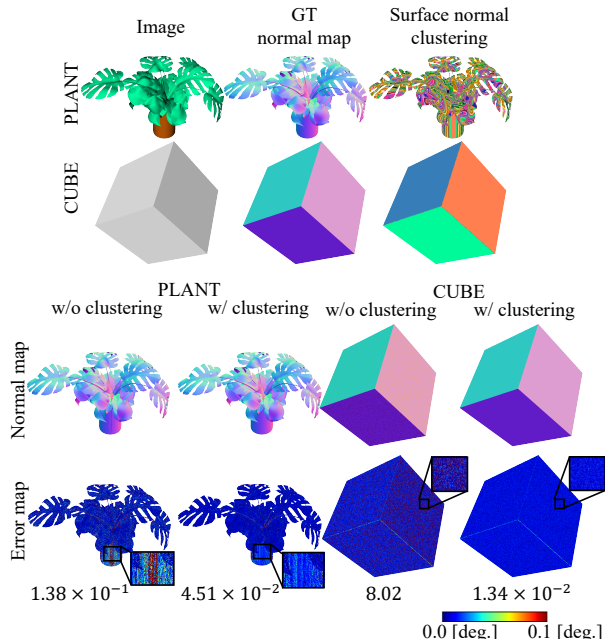


Figure 8. Estimation results for the PLANT and CUBE scenes by the proposed method with and without surface normal clustering. The numbers below the error maps indicate the mean angular errors in degrees. “w/o clustering” randomly selects the pixel pairs that have the same albedo.

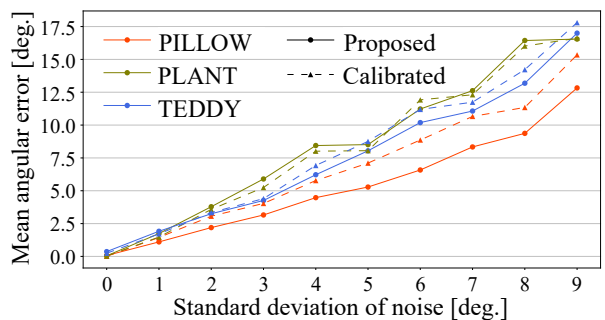


Figure 9. Estimation results for our synthetic datasets with adding noise to the light directions. The solid line is the proposed method, and the dotted line is the calibrated method. Each data point is an average of 30 trials.

ric stereo more practical, we typically want to put the light sources close to a target scene, resulting in a near light setting. One of our future directions is to extend the proposed method for the near light setting.

Acknowledgments This work was supported by JSPS KAKENHI Grant Number JP19H01123.

References

- [1] Jaakko Astola, Petri Haavisto, and Yrjo Neuvo. Vector median filters. *Proceedings of the IEEE*, 78(4):678–689, 1990.
- [2] Peter N. Belhumeur, David J. Kriegman, and Alan L. Yuille. The bas-relief ambiguity. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1060–1066, 1997.
- [3] Manmohan Chandraker, Jiamin Bai, and Ravi Ramamoorthi. On differential photometric reconstruction for unknown, isotropic BRDFs. *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)*, 35(12):2941–2955, 2013.
- [4] Manmohan K. Chandraker, Fredrik Kahl, and David J. Kriegman. Reflections on the generalized bas-relief ambiguity. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, volume 1, pages 788–795, 2005.
- [5] Guanying Chen, Kai Han, Boxin Shi, Yasuyuki Matsushita, and Kwan-Yee K. Wong. Self-calibrating deep photometric stereo networks. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 8739–8747, 2019.
- [6] Ondrej Drbohlav and Mike Chaniler. Can two specular pixels calibrate photometric stereo? In *Proceedings of IEEE International Conference on Computer Vision (ICCV)*, volume 2, pages 1850–1857, 2005.
- [7] Kenji Enomoto, Michael Waechter, Kiriakos N. Kutulakos, and Yasuyuki Matsushita. Photometric stereo via discrete hypothesis-and-test search. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2308–2316, 2020.
- [8] Athinodoros S. Georghiades. Recovering 3-D shape and reflectance from a small number of photographs. In *Proceedings of Eurographics Workshop on Rendering (EGRW)*, pages 230–240, 2003.
- [9] Athinodoros S. Georghiades. Incorporating the Torrance and Sparrow model of reflectance in uncalibrated photometric stereo. In *Proceedings of IEEE International Conference on Computer Vision (ICCV)*, volume 2, pages 816–823, 2003.
- [10] Hideki Hayakawa. Photometric stereo under a light source with arbitrary motion. *Journal of the Optical Society of America (JOSA)*, 11(11):3079–3089, 1994.
- [11] Berthold K.P. Horn and Michael J Brooks. The variational approach to shape from shading. *Computer Vision, Graphics and Image Processing*, 33(2):174–208, 1986.
- [12] Satoshi Ikehata. CNN-PS: CNN-based photometric stereo for general non-convex surfaces. In *Proceedings of European Conference on Computer Vision (ECCV)*, pages 3–18, 2018.
- [13] Charles L. Lawson and Richard J. Hanson. *Solving least squares problems*. Series in Automatic Computation, 1995.
- [14] Hiroaki Santo, Masaki Samejima, Yusuke Sugano, Boxin Shi, and Yasuyuki Matsushita. Deep photometric stereo networks for determining surface normal and reflectances. *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)*, 2020.
- [15] Hiroaki Santo, Michael Waechter, Wen-Yan Lin, Yusuke Sugano, and Yasuyuki Matsushita. Light structure from pin motion: Geometric point light source calibration. *International Journal of Computer Vision (IJCV)*, 128(7):1889–1912, 2020.
- [16] Hui L. Shen and Yue Cheng. Calibrating light sources by using a planar mirror. *Journal of Electronic Imaging*, 20(1):1–6, 2011.
- [17] Boxin Shi, Kenji Inose, Yasuyuki Matsushita, Ping Tan, Sai-Kit Yeung, and Katsushi Ikeuchi. Photometric stereo using internet images. In *Proceedings of International Conference on 3D Vision (3DV)*, volume 1, pages 361–368, 2014.
- [18] Boxin Shi, Yasuyuki Matsushita, Yichen Wei, Chao Xu, and Ping Tan. Self-calibrating photometric stereo. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1118–1125, 2010.
- [19] Boxin Shi, Zhipeng Mo, Zhe Wu, Dinglong Duan, Ssai K. Yeung, and Ping Tan. A benchmark dataset and evaluation for non-lambertian and uncalibrated photometric stereo. *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)*, 41(2):271–284, 2019.
- [20] Rakesh Shiradkar, Ping Tan, and Sim H. Ong. Auto-calibrating photometric stereo using ring light constraints. *Machine Vision and Applications*, 25(3):801–809, Apr. 2014.
- [21] William M Silver. *Determining shape and reflectance using multiple images*. PhD thesis, Massachusetts Institute of Technology, 1980.
- [22] Ping Tan, Satya P. Mallick, Long Quan, David J. Kriegman, and Todd Zickler. Isotropy, reciprocity and the generalized bas-relief ambiguity. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1–8, 2007.
- [23] Robert J. Woodham. Photometric method for determining surface orientation from multiple images. *Optical Engineering*, 19(1):139–144, 1980.
- [24] Zhe Wu and Ping Tan. Calibrating photometric stereo by holistic reflectance symmetry analysis. In *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1498–1505, 2013.
- [25] Alan L. Yuille, Daniel Snow, Russell Epstein, and Peter N. Belhumeur. Determining generative models of objects under varying illumination: Shape and albedo from multiple images using SVD and integrability. *International Journal of Computer Vision (IJCV)*, 35(3):203–222, 1999.
- [26] Wei Zhou and Chandra Kambhampettu. Estimation of illuminant direction and intensity of multiple light sources. In *Proceedings of European Conference on Computer Vision (ECCV)*, pages 206–220, 2002.
- [27] Zhenglong Zhou and Ping Tan. Ring-light photometric stereo. In *Proceedings of European Conference on Computer Vision (ECCV)*, pages 265–279, 2010.