

# Uncalibrated Photometric Stereo Based on Elevation Angle Recovery from BRDF Symmetry of Isotropic Materials

Feng Lu  
The University of Tokyo  
Japan  
lufeng@ut-vision.org

Imari Sato  
National Institute of Informatics  
Japan  
imarik@nii.ac.jp

Yoichi Sato  
The University of Tokyo  
Japan  
ysato@iis.u-tokyo.ac.jp

## Abstract

This paper addresses the problem of uncalibrated photometric stereo with isotropic reflectances. Existing methods face difficulty in solving for the elevation angles of surface normals when the light sources only cover the visible hemisphere. Here, we introduce the notion of “constrained half-vector symmetry” for general isotropic BRDFs and show its capability of elevation angle recovery. This sort of symmetry can be observed in a 1D BRDF slice from a subset of surface normals with the same azimuth angle, and we use it to devise an efficient modeling and solution method to constrain and recover the elevation angles of surface normals accurately. To enable our method to work in an uncalibrated manner, we further solve for light sources in the case of general isotropic BRDFs. By combining this method with the existing ones for azimuth angle estimation, we can get state-of-the-art results for uncalibrated photometric stereo with general isotropic reflectances.

## 1. Introduction

Photometric stereo is an important research topic for 3D scene recovery in computer vision. Compared with other techniques, e.g., multi-view stereo, it is better at handling the essential physical process by which light is modulated by the surface. In particular, it can deal with the interaction between the surface normal, illumination, and surface reflectance simultaneously, and thus, it has much broader prospects than other methods do and shows a potential to influence many other fields. For instance, although there are promising methods for 3D object recognition [5], they are typically multi view-based and can be further enhanced by considering the photometric cue. In this paper, we exploit surface reflectance properties to recover surface normals, especially their elevation angles, in the most challenging case of unknown light sources and surface reflectances.

The problem is difficult because real-world scenes show

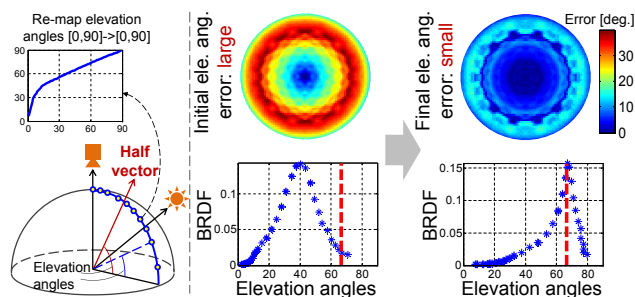


Figure 1. Elevation angle re-mapping using BRDF symmetry. Previous methods [17, 10, 3] can solve for the azimuth angles of surface normals accurately, whereas our method provides good estimates of the elevation angles. To this end, we propose “constrained half-vector symmetry”: BRDF values in the 1D slice (bottom-right figure) should have a symmetric distribution about the half vector (red dashed line). This symmetry is observed on a set of surface normals with varying elevation angles but the same azimuth angle with a source (bottom-left figure); it helps re-map the elevation angles (top-left figure) to reduce the final error.

a variety of different surface reflectances, resulting in diverse and complex bidirectional reflectance distribution functions (BRDFs). As a result, most of the previous uncalibrated (i.e., when light sources are unknown) methods assume simple surface reflectances, e.g., pure Lambertian reflectance, or that the diffuse component of the BRDF follows the Lambertian law. In doing so, the surface normals can be routinely solved up to the Generalized Bas-Relief (GBR) ambiguity [2]. The GBR ambiguity can be further resolved by using various properties observed in real-world BRDFs, including symmetries [20, 25], diffuse maxima [4] and pixel color profiles [18]. Since the GBR ambiguity only has a few spatial-invariant parameters, and it only works for the Lambertian BRDF, resolving it may not fully exploit the surface reflectance properties for scene recovery.

Several recent methods are capable of recovering the azimuth angles of surface normals without assuming a Lambertian BRDF [17, 10, 3]. However, they still face a challenge as far as elevation angle recovery goes. Although


[17, 10] provide estimates for the elevation angles, they must assume that the light sources cover the whole sphere uniformly or else their accuracies drop significantly. In practice, real scenes are usually hard to light from the back uniformly; it requires a special device to hold the object without blocking the lighting from the back. Moreover, lighting from the back always casts shadows that produce outliers. Therefore, it is expected that the lighting is only in the visible hemisphere, i.e., the camera-side hemisphere, like most conventional methods for Lambertian BRDF assume.

In this paper, we propose to leverage reflectance properties in order to meet this remaining challenge, under the conditions that the roughly uniform light sources only cover the visible hemisphere and the reflectance is homogeneous over the surface. In particular, we define a 1D BRDF slice based on the surface normals sharing the same azimuth angle with a light source, from which we can observe the constrained half-vector symmetry in general isotropic BRDFs. As shown in Fig. 1, as the elevation angles' error becomes smaller, the computed BRDF data in the 1D slice is more symmetrically distributed on the two sides of the half vector (red dashed line), which is the bisector of the lighting direction and viewing direction. This feature constrains the elevation angles and it enables us to devise an efficient method to solve for them. In addition, our method solves for light sources, and thus, it can work in an uncalibrated manner.

The contributions of this paper include: 1) a derivation of constrained half-vector symmetry (Sec 3); 2) an algorithm to determine the elevation angles of surface normals by using this symmetry (Sec 4); 3) light source estimation in the case of general isotropic reflectances (Sec 5). In experiments, our methods shows state-of-the-art accuracy for the challenging problem of uncalibrated photometric stereo with general isotropic reflectance.

## 2. Related Works

Whereas conventional photometric stereo methods assume known light directions and Lambertian reflectance [23], many of the recent studies relax these assumptions. When the positions of the light sources are unknown, the surface normals can be solved for up to the GBR ambiguity [2] by assuming Lambertian reflectance. The GBR ambiguity can be further resolved by using additional scene properties such as diffuse maxima [4], reflectance symmetry [20, 25], color profiles [18], and special light source positions [26]. However, for many general reflectances that do not obey the Lambertian law, the problem becomes harder, and known light sources are still needed. Some methods find and remove non-Lambertian components from observations by using techniques including median filter [13], rank-minimization [24] and sparse regression [9], while other methods obtain the surface geometry directly by leveraging various surface reflectance properties [1, 8, 19, 20].



(a) Local coord. sys. (b) View-centered coord. sys.  
 Figure 2. (a) Standard BRDF parameterization; (b) in a view-centered coordinate system, we examine the surface normals having the same azimuth angle as the light source. We observe the *constrained half-vector symmetry* in the resulting 1D BRDF slice.

In these methods, dense light sources with known positions usually provide crucial cues for obtaining a solution.

There are methods that handle both uncalibrated light sources and general isotropic reflectances. Some of them put additional objects with known shapes and reflectances in the scene as references for calibration [7, 15], and use certain reflectance models in their formulation [6]. Other methods include those using light sources placed in a ring [3] or exploiting the similarity of pixel intensity profiles when the light sources cover the whole sphere [17, 14, 10]. They work well in recovering azimuth angles of surface normals, but face difficulty [17, 14, 10] or require additional inputs [3] to determine the elevation angles of the surface normals. In particular, if the light sources only cover the visible hemisphere on the camera side, the methods described in [17, 14, 10] produce large errors in the elevation angles of the surface normals, while the azimuth angles and iso-depth contours remain accurate. Compared with these methods, our method accurately solves for the elevation angles without requiring light sources from the back, and it achieves state-of-the-art results. Note that a recent method developed by Shi *et al.* [19] refines the elevation angles, but it requires the positions of the light sources as input.

## 3. Constrained Half-Vector Symmetry

In this section, we review the definition and parameterizations of BRDFs and some of the important symmetries. Then, we propose *constrained half-vector symmetry*, which can be observed in a 1D BRDF slice, to optimize the solution of the elevation angles of the surface normals in photometric stereo problems.

### 3.1. BRDF and Its Parameterization

A BRDF measures the ratio of the reflected radiance from a surface patch. It is a function  $f(\omega_{in}, \omega_{out})$  of incoming and outgoing light directions in a local coordinate system. Many recent studies have used halfway/difference parameterization [16] to represent BRDFs. By introducing the half vector, which is the bisector of  $\omega_{in}$  and  $\omega_{out}$ , a BRDF can be parameterized as  $f(\theta_h, \phi_h, \theta_d, \phi_d)$ , as illustrated in Fig. 2 (a). Thus, a pixel value captured at a surface

patch can be expressed as

$$I = f(\theta_h, \phi_h, \theta_d, \phi_d)(\mathbf{n}^T \mathbf{s}), \quad (1)$$

where  $\mathbf{n}$  is the surface normal and  $\mathbf{s}$  is the point light source.

The above 4D BRDF parameterization domain can be reduced by using some of the symmetries observed in many real-world materials. For instance, the widely observed ‘isotropy’ reduces the BRDF to a 3D function  $f(\theta_h, \theta_d, \phi_d)$  which is invariant with respect to  $\phi_h$  [11]. This is the most widely adopted assumption in the related research, and nearly all materials in the MERL BRDF database conform to it [12]. Moreover, the ‘reciprocity’ and ‘bilateral’ symmetries further fold the domain of  $\phi_d$  from  $[0, \pi]$  to  $[0, \pi/2]$ . Finally, by considering the ‘half-vector’ symmetry, which assumes an invariant BRDF against a rotation  $\phi_d$  of both  $\omega_{in}$  and  $\omega_{out}$  around the half vector, the BRDF function can be further reduced to a bivariate one, denoted as  $f(\theta_h, \theta_d)$ .

### 3.2. Constrained Half-Vector Symmetry and 1D BRDF Slice

Let us examine a special case of isotropic reflectance wherein a set of surface normals  $\{\mathbf{n}\}$  with different elevation angles share the same azimuth angle with a light source  $\mathbf{s}$  in the view-centered coordinate system, as shown in Fig. 2 (b). This results in a fixed  $\theta_d$  and also  $\phi_d = 0$  or  $\pi/2^1$  in the conventional BRDF parameterization with local coordinates in Fig. 2 (a). Consequently, this simplifies the isotropic BRDF function to either  $f_{\theta_d, \phi_d=0}(\theta_h)$  or  $f_{\theta_d, \phi_d=\pi/2}(\theta_h)$ . If one further assumes reciprocity, meaning that the roles of  $\omega_{in}$  and  $\omega_{out}$  can be interchanged, the two representations can be unified to be  $f_{\theta_d, \phi_d=0}(\theta_h)$ , which represents a 1D slice from the original BRDF, as shown in Fig. 2 (b). Since  $f_{\theta_d, \phi_d=0}(\theta_h)$  is symmetric about the half-vector (i.e.,  $\theta_h = 0$ ), it leads to the following observation shown in Fig. 2 (b).

**Observation 1.** *By assuming isotropy and reciprocity, if the elevation angles of the surface normals and the light source are correctly measured, the BRDF data computed from the observations in Fig. 2 (b) should be distributed symmetrically about the half vector.*

As explained above, this property is a direct result of assuming isotropy and reciprocity. Meanwhile, it can also be derived from the half-vector symmetry by fixing  $\theta_d$  in a bivariate BRDF function<sup>2</sup>. For convenience, we will use the term “constrained half-vector symmetry” in the rest of this paper since the concept of the half vector is important.

<sup>1</sup> $\omega_{in}$  and  $\omega_{out}$  can be rotated by  $\pi/2$  to interchange their positions, while still ensuring that  $\mathbf{n}$ ,  $\omega_{in}$ , and  $\omega_{out}$  are coplanar.

<sup>2</sup>In most cases, assuming half-vector symmetry is stronger than assuming isotropy and reciprocity. Therefore, our symmetry is in fact fit for more general reflectances than conventional half-vector symmetry.

## 4. Elevation Angle Re-mapping

In this section, we describe the recovery of surface normals of test scenes. In particular, we propose an accurate elevation angle recovery by using the constrained half-vector symmetry. Note that we will use known light positions to describe the method. However, this assumption will be discarded in Sec 5.

### 4.1. Symmetric Data Acquisition from Images

Scene images are captured by a fixed camera with roughly uniform light sources in the visible hemisphere. We employ the previously proposed method [17] for obtaining the initial estimates of the surface normals from these images. Pixel profiles are obtained in order to measure the surface normals’ similarities. Then the surface normals are recovered from these similarity measurements. Because the light sources are only in the visible hemisphere, the measurements are asymmetric along the viewing direction. Preliminary experiments indicate the average error of the elevation angle is more than twice that of the azimuth angle.

We use the initialized surface normals to extract symmetric pairs from 1D BRDF slices for estimating the elevation angles. For simplicity, let us describe this procedure for a given light source  $\mathbf{s}$ . We will do the same for each of the other light sources. In the view-centered coordinate system, we find a set of captured pixels whose corresponding surface normals  $\{\mathbf{n}^j\}$  share the same azimuth angle with  $\mathbf{s}$ . In practice, a threshold of the azimuth angle difference can be used, e.g.,  $1^\circ$ , to ensure that the number of selected  $\{\mathbf{n}^j\}$  is larger than 10. Then, each pixel value is converted into a BRDF value by dividing it by  $\mathbf{n}^j \cdot \mathbf{s}$ , using Eq. (1). Finally, we obtain a 1D slice  $\rho$  containing all BRDF values observed at  $\{\mathbf{n}^j\}$ .

The next step is to extract symmetric pairs. Given the viewing direction  $[0, 0, -1]^T$  and the light source  $\mathbf{s}$ , we first compute the half vector as their bisector and record its elevation angle as  $h$ . We also find the surface normal in  $\{\mathbf{n}^j\}$  corresponding to the largest BRDF value in  $\rho$  and record its elevation angle as  $\delta$ . It is reasonable to assume that the largest BRDF is observed at  $h$ , and thus,  $\delta$  should be close to  $h$ . Then, we extract the elevation angles of  $\{\mathbf{n}^j\}$  in a pairwise manner that satisfy constrained half-vector symmetry. We traverse all elements of  $\rho$ , and for each of them, indexed by the elevation angle  $\alpha^j$ , we find in  $\rho$  another element  $\beta^j$  such that 1)  $\alpha^j$  and  $\beta^j$  correspond to the same BRDF value; 2)  $\alpha^j$  and  $\beta^j$  are located on two sides of  $\delta$ . Note that  $\alpha^j$ ,  $\beta^j$  and  $\delta$  are the initial elevation angles or their current estimates during the iteration (see Sec 4.3).

In practice,  $\beta^j$  does not usually have exactly the same BRDF value as  $\alpha^j$ . Therefore, instead of finding the exact  $\beta^j$ , we find two elevation angles  $\beta^j$  and  $\gamma^j$ , whose BRDF values are closest to that of  $\alpha^j$ . By assuming the BRDF values change smoothly, we can compute the weights  $\mu^j$  and

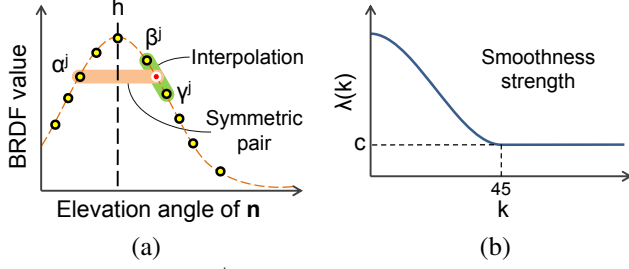


Figure 3. (a) For each  $\alpha^j$ , we find its symmetric partner in the 1D slice by using interpolation; (b) smoothness level in Eq. (4).

$\nu^j$  such that  $\rho(\alpha^j) = \mu^j \rho(\beta^j) + \nu^j \rho(\gamma^j)$  and use them to produce the symmetry partner of  $\alpha^j$  with  $\mu^j \beta^j + \nu^j \gamma^j$ . As a result,  $(\alpha^j, \mu^j \beta^j + \nu^j \gamma^j)$  is considered to be a symmetric pair. This situation is illustrated in Fig. 3 (a).

Ideally, the data collected on the light source  $s$  should satisfy constrained half-vector symmetry in the following ways: 1)  $\delta = h$  and 2)  $\frac{\alpha^j + (\mu^j \beta^j + \nu^j \gamma^j)}{2} = h$ , which means  $h$  should be the bisector of the elevation angles of any symmetric pair, as shown in Fig. 3 (a). On the other hand, these data are computed from inaccurate estimates of the elevation angles and thus may break the symmetry, as shown in Fig. 1. Therefore, we can use symmetry as a criterion to correct the elevation angles of the surface normals.

## 4.2. Formulation with 1D BRDF Slices

By using the symmetric data described in the previous section, we propose to refine the elevation angles of the surface normals so that they satisfy constrained half-vector symmetry for all 1D BRDF slices. We model the refinement as an elevation angle re-mapping process. By assuming a global mapping and correct azimuth angles, which are supported by results from [17, 10], we establish a one-to-one mapping  $\varepsilon \mapsto \hat{\varepsilon} : \hat{\varepsilon} = m(\varepsilon)$  with boundary conditions  $m(0) = 0$  and  $m(\pi/2) = \pi/2$  to refine any elevation angle  $\varepsilon$ . The mapping can be simple or highly non-linear, depending on the accuracy of the initial elevation angles.

We formulate the problem by using constraints derived from the constrained half-vector symmetry. In particular, for the  $i^{th}$  light source  $s_i$ , the collected elevation angle pairs should satisfy 1)  $m(\delta_i) = h_i$  and 2)  $\frac{m(\alpha_i^j) + [\mu_i^j m(\beta_i^j) + \nu_i^j m(\gamma_i^j)]}{2} = h_i$  after the re-mapping, where the subscript  $i$  indicates data from a 1D BRDF slice lit by a source  $s_i$ . This directly leads to the formulation,

$$E_{data}(m(\cdot)) = \sum_i \left[ \underbrace{\kappa [m(\delta_i) - h_i]^2}_{\text{align extreme to } h_i} + \underbrace{\sum_j [m(\alpha_i^j) + [\mu_i^j m(\beta_i^j) + \nu_i^j m(\gamma_i^j)] - 2h_i]^2}_{\text{make every elevation angle pair symmetric about } h_i} \right], \quad (2)$$

where  $\kappa$  balances the influence of different terms. In practice, we use  $\kappa = 10$ , due to the importance of information directly from the half vector. Note that Eq. (2) uses an outer summation to take into account equations from all light sources  $\{s_i\}$ .

In addition, a smoothness term can be added to Eq. (2) to suppress the influence of outliers. To this end, we force the mapping curve to change gradually by minimizing the sum of the squares of the second derivative of  $m(\cdot)$  as

$$E_{smooth}(m(\cdot)) = \sum_{k=0+1}^{90-1} \lambda(k) [m''(k)]^2, \quad (3)$$

where  $k$  represents sampling angles to define  $m(\cdot)$  in a discrete form (see the x-axis values in the top-left plot of Fig. 1), and  $\lambda(k)$  determines the level of elevation-angle-dependent smoothness. In this discrete formulation, we use  $m''(k) = m(k-1) - 2m(k) + m(k+1)$  to compute the second derivative. Note that, here,  $m(\cdot)$  is defined at elevation angles from  $0^\circ$  to  $90^\circ$  in steps of  $1^\circ$ , while a formulation with a higher sampling density is straightforward.

The balance weight  $\lambda(\cdot)$  is specially designed to be

$$\lambda(k) = \begin{cases} c \times (2 \times \cos(4k) + 3), & k < 45 \\ c, & \text{otherwise} \end{cases}, \quad (4)$$

where  $c$  is a constant value ( $= 100$  in our experiments). The curve of the smoothness level is in Fig. 3 (b). By using this spatially varying smoothness, small elevation angles (which tend to suffer from noise because their surface points are always dark without lighting from the back) are smoothed out to avoid effects due to the low signal-noise-ratio.

The final formulation of the problem takes into consideration both terms as follows:

$$\begin{aligned} m(\cdot) &= \underset{m(\cdot)}{\operatorname{argmin}} (E_{data} + E_{smooth}), \\ \text{s.t. } & m(0) = 0, m(\pi/2) = \pi/2, \\ & m([0, \pi/2]) \subseteq [0, \pi/2]. \end{aligned} \quad (5)$$

## 4.3. Solution and Uniqueness

We solve the unknown mapping function  $m(\cdot)$  through matrix calculations. Under the condition that the elevation angles are sampled from  $0^\circ$  to  $90^\circ$  in steps of  $1^\circ$ , we can re-define  $m(\cdot)$  as a 1D vector  $\mathbf{m} \in \mathbb{R}^{91 \times 1}$ , while its  $k^{th}$  element equals  $m(k)$  which is the mapping result for elevation angle  $k$ . Accordingly,  $\mathbf{m}$  numerically defines a mapping of elevation angles, which allows us to re-write Eq. (5) as

$$\mathbf{m} = \underset{\mathbf{m}}{\operatorname{argmin}} \|\mathbf{A}\mathbf{m} - \mathbf{b}\|_2 \quad \text{s.t. } \mathbf{m} \geq 0, \quad (6)$$

where the matrices  $\mathbf{A}$  and  $\mathbf{B}$  are constructed as follows:

$$\mathbf{A} = \begin{pmatrix} \text{column:} & [\beta_i^j] & [\alpha_i^j] & [\delta_i] & [\gamma_i^j] \\ \vdots & \vdots & \vdots & \vdots & \vdots \\ \dots & \dots & \dots & \dots & \dots \\ \dots & \mu_i^j & 1 & \dots & \nu_i^j \\ \dots & \dots & \dots & \lambda(k) & -2 \cdot \lambda(k) & \lambda(k) & \dots \\ \dots & \dots & \dots & \dots & \dots & \dots & \dots \\ \text{Inf} & \dots & \dots & \dots & \dots & \dots & \dots \\ & & 0 & & & & \\ & & & & & & \text{Inf} \end{pmatrix}, \mathbf{b} = \begin{pmatrix} \vdots \\ \vdots \\ 2 \cdot h_i \\ 2 \cdot h_i \\ \vdots \\ 0 \\ \vdots \\ \vdots \\ 0 \\ 90 \cdot \text{Inf} \end{pmatrix}, \quad (7)$$

where  $[\beta_i^j]$  takes the rounded value of  $\beta_i^j$  and indicates the column index. Note that the same row from the matrix  $\mathbf{A}$  and vector  $\mathbf{b}$  exactly constitute a sub-formula in either Eq. (2) or Eq. (3), except for the last two rows representing the boundary conditions for 0 and  $\pi/2$ , where ‘Inf’ indicates a constant large value ( $= 10^5$  in our experiments). Therefore, the solution of Eq. (6) numerically approximates the solution of Eq. (5). The solution can be obtained by using least squares solvers with non-negative constraints.

**Iteration** In Sec 4.1, we compute the BRDF values by using the initial elevation angles, and thus, those values are inaccurate. After the elevation angles are updated using our solution, we re-compute the BRDF values and solve for the elevation angles again. Note that we do not update the choice of  $\{\mathbf{n}^j\}$  since the azimuth angles are assumed correct. In practice, 3 iterations suffice to get a final solution.

**Solution uniqueness** It is important to know whether the solution to Eq. (6) is unique or not in order to determine the surface normals. For this, we can prove that

**Proposition 1.** *Without considering the smoothness term, the problem in Eq. (5) (Eq. (6)) has a unique solution under the following conditions: 1) the surface reflectance is not Lambertian; 2) there are at least two light sources with different elevation angles.*

We prove the second condition in the Appendix and give an intuitive explanation of it here. Let us suppose that all light sources have the same elevation angle, which produces the same half vector. Our symmetry cannot determine the elevation angles of the surface normals with only one half vector, because the distribution of elevation angles on one side of the half vector can be freely changed, and the other side can be symmetrically changed after that. Only another half vector can remove this additional degree of freedom.

As for the first condition, it is obvious because Lambertian reflectance has an identical BRDF value everywhere, which makes it too ambiguous to use to measure the symmetry. Many diffuse materials like latex and rubber cause this problem, which make the output of our method unstable and non-convergent during the iteration. Meanwhile, we can easily identify such materials by checking the convergence of the output; we can simply keep the original results produced by methods like [17], since Lambertian-like materials are easy for the conventional methods to handle.

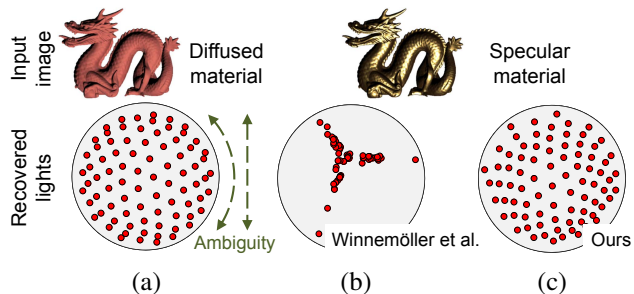


Figure 4. Recovered 3D light source positions in the visible hemisphere. (a) and (b): results using raw pixel profiles as in [21]; (c): results of our method.

## 5. Robust Light Source Estimation

The previous sections use light sources to compute the half vectors. If they are not given beforehand, we have to estimate them from the input images. By assuming that the light sources cover the visible hemisphere, Winnemöller *et al.* [22] show that 3D light source positions can be recovered by using similarities among images taken under different lighting conditions. In particular, a pixel profile can be constructed so as to comprise all pixels in the same image with a fixed order (e.g., raster scan). Let us denote a profile consisting of pixel intensities  $\{I_i^p\}$  from the image captured under the  $i^{th}$  light source as  $\chi_i = [I_i^1, \dots, I_i^p, \dots]^T$ . As demonstrated in [22], by varying the light source position moderately, the differences among  $\{\chi_i\}$  can be strongly correlated to the differences in the light sources positions. Following this observation, one can apply dimensionality reduction (e.g., ISOMAP [21]) to embed all pixel intensity profiles  $\{\chi_i\}$  into 3D space and map them to the visible hemisphere to recover the unknown light sources.

We use this procedure to estimate light sources solely from images taken under different lighting conditions; our technique includes two major extensions to that of [22]:

- (1) [22] works well for simple Lambertian reflectance but not for glossy or specular ones, while our method works for general isotropic reflectances.
- (2) We resolve their remaining rotation/flip ambiguity by aligning the light sources to the surface normals.

**General Reflectances** An example of light sources recovered by [22] is shown in Fig. 4 (a), where the captured surface has Lambertian reflectance. However, as shown in Fig. 4 (b), in the case of general isotropic reflectances, e.g., specular ones, light source positions cannot be correctly recovered by using raw pixel intensity profiles  $\{\chi_i\}$ . This is because the pixel values of specular materials change greatly in a highly nonlinear manner and thus they cannot be used to measure differences between images.

To overcome this problem, note that although different materials produce very different pixel values, the locations of the bright and dark regions simply convey light

source positions, as in the case of the two “dragon” surfaces in Fig. 4. Therefore, we propose to modify the pixel intensity profiles so that they can handle more general isotropic reflectances. In particular, we replace  $\chi_i = [I_i^1, \dots, I_i^p, \dots]^T$  with  $\hat{\chi}_i = [o(I_i^1), \dots, o(I_i^p), \dots]^T$ , where  $o(I_p)$  is the ascending order of  $I_p$  of all pixel values in  $\chi_i$ . By using  $\{\hat{\chi}_i\}$ , we can recover accurate light source positions as shown in Fig. 4 (c) even with highly specular materials. An explanation of the dimensionality reduction using ISOMAP can be found in [21, 22].

**Rotation/Flip Ambiguity** The solution is invariant against rotations/flips of all light sources around/about the viewing direction, as shown in Fig. 4 (a). Winnemöller *et al.* [22] leave this ambiguity for the user to deal with manually, whereas we propose an automatic method that aligns the light sources to the recovered surface normals.

We first model the rotation and flip ambiguity as

$$Q = \begin{pmatrix} \eta \cdot \cos \phi_v & -\eta \cdot \sin \phi_v & 0 \\ \sin \theta_v & \cos \theta_v & 0 \\ 0 & 0 & 1 \end{pmatrix}, \quad (8)$$

where  $\theta_v$  controls the rotation around the viewing direction, and  $\eta = \pm 1$  flips the axis. The ground truth solution can be expressed as  $Qs_i$ , assuming the current solution is  $s_i$ .

A simple observation enables us to solve  $Q$ . That is, any light source  $s_i$  always shares the same azimuth angle with the surface normal  $\tilde{n}_i$ , which corresponds to the brightest pixel  $\tilde{I}_i$  that  $s_i$  can produce. Thus,

$$Q = \operatorname{argmin}_Q \sum_i \left[ \tilde{I}_i \cdot \arccos \frac{c_{2,1}(Qs_i)^T \cdot c_{2,1}(\tilde{n}_i)}{|c_{2,1}(Qs_i)| \cdot |c_{2,1}(\tilde{n}_i)|} \right], \quad (9)$$

where  $c_{2,1}(\cdot)$  keeps the first two rows and one column of the input vector, and the pixel value  $\tilde{I}_i$  balances the influence of the current term since bright pixels are reliable. Eq. (9) can easily be solved since the only unknown is the angle  $\theta_v$  and  $\eta = \pm 1$ . In this way,  $Q$  yields correct light sources  $\{Qs_i\}$ .

## 6. Experimental Evaluation

### 6.1. Synthetic Datasets




Figure 5. Positions of light sources in the hemisphere.

We quantitatively evaluate our method by using synthetic scenes. To render the scenes, we use 100 different BRDFs in the MERL BRDF database [12] with varying lighting positions. Unlike in [17, 10], where the lighting positions cover the whole sphere, ours only cover the visible hemisphere.

Table 1. Average normal/light errors of 100 materials. Our method automatically recognizes and skips Lambertian-like materials.

Materials	Initial err.	Final err.	Light err.
Processed (81)	18.66°	<b>6.21°</b>	7.41°
Skipped (19)	4.76°	4.76°	7.50°
Average (100)	15.98°	<b>5.89°</b>	7.43°




Figure 6. Examples of automatic material type classification.

All 82 light source positions are illustrated in Fig. 5. For each source, we produce scene images using 1) a standard hemispherical surface and 2) several 3D models, e.g., Bunny, Armadillo, Rabbit, and Happy Buddha. These images are used as inputs to our method. To initialize the surface normals, we use the method in [17].

**Accuracy for All 100 Materials** First, we test our method using hemispherical surfaces. Table 1 shows the average results for all 100 materials. The error is 5.89° for surface normal recovery, which is pretty accurate considering that both the materials and the light sources are unknown. Moreover, our method “skips” 19 Lambertian-like materials in Table 1; i.e., their initial surface normals are already accurate (4.76° in average) and are not updated as described in Sec 4.3. For the other materials, our method significantly increases the average accuracy, from 18.66° to 6.21°. The big difference in accuracy indicates the effectiveness of our method. The results for light source recovery are insensitive to the type of material and have an average error of 7.43°.

Note that the decision whether or not to skip a material is automatically determined via optimization as described at the end of Sec 4.3. To verify that the procedure is correct, we show examples of processed and skipped materials in Fig. 6. These examples, together with the accuracies in Table 1, prove the validity of our analysis in Sec 4.3.

In addition, we plot the errors for each material (Fig. 7) and show some representative examples of elevation angle mappings and recovered 1D symmetries. In particular, material A corresponds to a Lambertian-like material without obvious 1D half-vector symmetry, and thus our method does not update its results; materials B and C are processed correctly and their recovered 1D BRDF symmetries can be seen; material D has the worst accuracy due to its additional lobe that breaks our half-vector-based symmetry. Fig. 7 also compares the obtained elevation angle mapping curves of these four materials (blue) with the optimal curves (red).

**Comparison with Other Methods** We compare our results with those of previous methods. Experiments using the 100 hemispherical surfaces are conducted on the methods

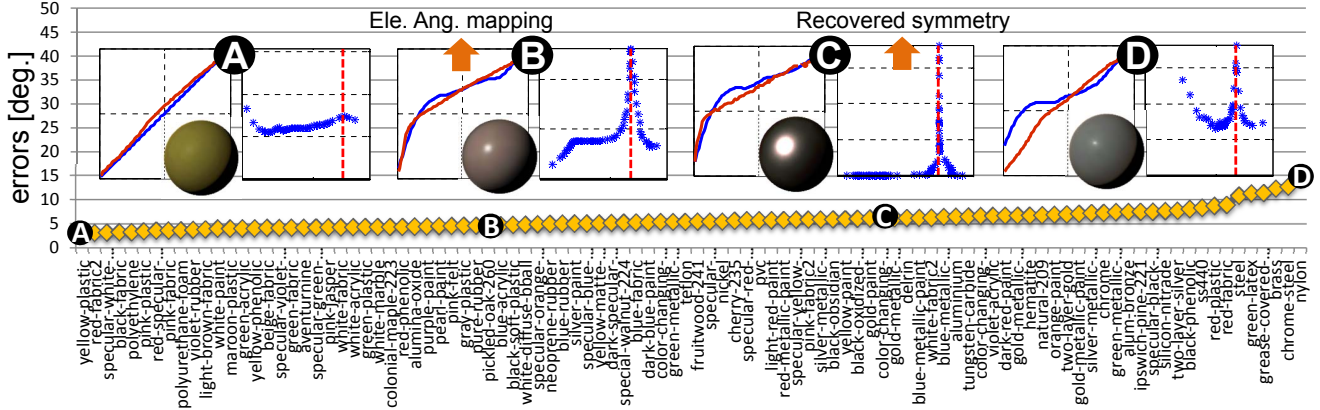


Figure 7. Detailed results for 100 materials. Computed elevation angle mapping curves (blue line: estimated; red line: ground truth) and recovered symmetry data (red dashed line for half vector) for four representative materials.

Table 2. Accuracies of ours and other uncalibrated methods.

	Ours	Lu <i>et al.</i> [10]	Sato <i>et al.</i> [17]	Wu <i>et al.</i> [24]
Normal err.	5.89°	14.95°	15.99°	20.00°
Light err.	7.43°	—	—	—

Table 3. Results for 3D synthetic models made with 10 materials.

	Bunny	Armadillo	Rabbit	Happy Buddha	Average
Ours	6.2°	6.0°	7.6°	6.3°	6.5°
[10]	12.2°	13.0°	12.8°	11.9°	12.5°
[17]	15.3°	16.0°	16.8°	15.5°	15.9°




Figure 8. Detailed comparison (large errors are cut at 50°).

of Lu *et al.* [10], Sato *et al.* [17] and Wu *et al.* [24]. Note that although we test “uncalibrated” methods, we provide the method in [24] with ground truth light sources to help it resolve the GBR ambiguity for some difficult materials.

Average results are listed in Table 2. Our method is significantly more accurate than the other methods. Note that although Lu *et al.* [10] and Sato *et al.* [17] report good accuracies when assuming uniform light sources on the whole sphere, their errors increase significantly with only light sources on the visible hemisphere. Also note that our method is the only one that can recover the light sources. Per-material errors are plotted in Fig. 8 for reference.

**Results for Complex 3D Models** We test our method by using 3D models. Their images are rendered by using ten typical materials in the MERL database under the light sources in Fig. 5. Examples of the rendered images and the recovery results are shown in Fig. 9, where we visually compare the

error maps produced by our method and by [10]. As expected, the errors from [10] are basically in the elevation angles, while our method can greatly improve on their accuracy and eventually achieves good estimates of the surface geometry. Exact numbers of recovery errors are compared with those of [10] and [17] in Table 3 (We do not include the results for the method in [24] since it fails to give reasonable outputs for some of the test materials.).

## 6.2. Real Scenes

Here, we show some qualitative results for real-world objects. The images are captured with light sources only on the visible hemisphere. There are typically 60 such sources. The recovered surface normal maps for different objects are shown in Fig. 10, together with the depth maps reconstructed from these surface normal maps. The correctness of the object shapes can be checked from the depth maps. Although these objects have various surface reflectances, our method still gives good estimates of their geometry.

## 7. Conclusion

We propose constrained half-vector symmetry, which is the fact that the observed BRDF values in a 1D slice have a symmetric distribution about the half vector. By using such symmetric data, we develop an efficient and accurate modeling and solution method for elevation angle recovery. Our method also recovers light sources in the case of general isotropic BRDFs, and this helps in the elevation angle estimation. By combining our method with existing techniques




Figure 10. Exemplary results of surface normal maps and depth maps recovered from real-world images by using our method.




Figure 9. Representative results for synthetic 3D surfaces. Four out of the ten test materials are shown in the first row.

for estimating the azimuth angles, we obtain satisfactory solutions for the challenging problem of uncalibrated photometric stereo with general isotropic reflectances.

## Appendix A: Prove of Proposition 1

*Proof.* Without assuming the smoothness term and boundary conditions, we first consider the case with only one light source. The matrix  $\mathbf{A}$  in Eq. (6) is composed of 1) one row of  $[\dots, 2\kappa, \dots]$  and 2) multiple rows of  $[\dots, \mu_i^j, \dots, 1, \dots, \nu_i^j, \dots]$ , and the vector  $\mathbf{b}$  contains a unique half vector value. In Sec 4.1, weights of  $\mu_i^j$  and  $\nu_i^j$  are used only when the surface point with the exact elevation angle is not captured. In the ideal case where we can capture data at every sampling position of elevation angle, the rows become  $[\dots, 1, \dots, 1, \dots]$  due to the existence of any symmetric pairs. Because there is only one light source, every symmetric pair is unique, and has no overlap with any other pair. Therefore, the maximum number of unique rows (pairs) is  $(N - 1)/2 + 1$ , where  $N$  is the number of elements in  $\mathbf{m}$ , which is 91 in Sec 4.2 and Sec 4.3, but can be made much larger for better accuracy.

This shows that a system with only one light source has at most  $(N - 1)/2 + 1$  independent equations but  $N$  unknowns, and thus it is underdetermined and has multiple solutions. Now let us go back to the non-ideal case, where the data at a certain sampling position of elevation angle is not captured in practice; here, the corresponding row becomes  $[\dots, \mu_i^j, \dots, 1, \dots, \nu_i^j, \dots]$ . However, as described in Sec. 4.1, it approximates the ideal case numerically and does not increase the number of effective functions.

Furthermore, let us consider multiple light sources. If an additional light source exists but has the same elevation angle as the first one, its half vector will also be the same as the first one and thus their resulting symmetric BRDF pairs will be unique. Therefore, the two light sources at most produce the same  $(N - 1)/2 + 1$  equations and the problem is still underdetermined. Only if the second light source has a different elevation angle, it can produce at most another  $(N - 1)/2 + 1$  independent equations, so that the system can be determined.

□



## 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] 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. 1, 2
- [3] 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. 1, 2
- [4] P. Favaro and T. Papadhimetri. 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. 1, 2
- [5] Y. Gao, M. Wang, D. Tao, R. Ji, and Q. Dai. 3-d object retrieval and recognition with hypergraph analysis. *IEEE Trans. on Image Processing*, 21(9):4290–4303, 2012. 1
- [6] 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
- [7] 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. 2
- [8] 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
- [9] S. Ikehata, D. Wipf, Y. Matsushita, and K. Aizawa. Robust photometric stereo using sparse regression. In *Proc. of IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, pages 318–325, 2012. 2
- [10] F. Lu, Y. Matsushita, I. Sato, T. Okabe, and Y. Sato. Uncalibrated photometric stereo for unknown isotropic reflectances. In *Proc. of IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2013. 1, 2, 4, 6, 7
- [11] S. R. Marschner. *Inverse Rendering for Computer Graphics*. PhD thesis, Cornell University, 1998. 3
- [12] W. Matusik, H. Pfister, M. Brand, and L. McMillan. A data-driven reflectance model. In *Proc. of ACM SIGGRAPH*, pages 27–31, 2003. 3, 6
- [13] 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
- [14] 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
- [15] P. Ren, J. Wang, J. Snyder, X. Tong, and B. Guo. Pocket reflectometry. In *Proc. of ACM SIGGRAPH*, pages 45:1–45:10, 2011. 2
- [16] S. M. Rusinkiewicz. A new change of variables for efficient brdf representation. In *Rendering techniques 98*, pages 11–22. Springer, 1998. 2
- [17] 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, 4, 5, 6, 7
- [18] 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. 1, 2
- [19] 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
- [20] 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. 1, 2
- [21] J. Tenenbaum, V. De Silva, and J. Langford. A global geometric framework for nonlinear dimensionality reduction. *Science*, 290(5500):2319–2323, 2000. 5, 6
- [22] H. Winnemöller, A. Mohan, J. Tumblin, and B. Gooch. Light waving: Estimating light positions from photographs alone. *Computer Graphics Forum*, 24:433–438, 2005. 5, 6
- [23] R. Woodham. photometric method for determining surface orientation from multiple images. *Optical Engineering*, 1(7):139–144, 1980. 2
- [24] 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, 7
- [25] Z. Wu and P. Tan. Calibrating photometric stereo by holistic reflectance symmetry analysis. In *Proc. of IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2013. 1, 2
- [26] Z. Zhou and P. Tan. Ring-light photometric stereo. In *Proc. of European Conf. on Computer Vision (ECCV)*, pages 265–279, 2010. 2