

Elevation Angle from Reflectance Monotonicity: Photometric Stereo for General Isotropic Reflectances

Boxin Shi*, Ping Tan[†], Yasuyuki Matsushita[‡], and Katsushi Ikeuchi*

*The University of Tokyo,

[†]National University of Singapore, [‡]Microsoft Research Asia

Abstract. This paper exploits the monotonicity of general isotropic reflectances for estimating elevation angles of surface normal given the azimuth angles. With an assumption that the reflectance includes at least one lobe that is a monotonic function of the angle between the surface normal and half-vector (bisector of lighting and viewing directions), we prove that elevation angles can be uniquely determined when the surface is observed under varying directional lights densely and uniformly distributed over the hemisphere. We evaluate our method by experiments using synthetic and real data to show its wide applicability, even when the assumption does not strictly hold. By combining an existing method for azimuth angle estimation, our method derives complete surface normal estimates for general isotropic reflectances.

1 Introduction

Photometric stereo estimates a pixel-wise surface normal direction from a set of images taken under varying illumination and a fixed viewpoint [1]. However, the estimation suffers from inaccuracy when the surface has non-Lambertian reflectances. Although more flexible parametric reflectance models (*e.g.*, the Ward model [2]) can be integrated to describe the non-Lambertian phenomenon, there exists a trade-off between the computational complexity and generality of the materials that can be dealt with. In this paper, we exploit the reflectance monotonicity for estimating elevation angles of surface normal given the azimuth angles (*e.g.*, from the method in [3]) to fully determine the surface normal for a broad class of reflectances.

As described by Chandraker and Ramamoorthi [4], an isotropic Bidirectional Reflectance Distribution Function (BRDF) consists of a sum of *lobes*, and the reflectance of each lobe monotonically decreases as the surface normal deviates from the lobe’s projection direction, along which the reflectance function is ‘concentrated’. Following their work, we assume that a surface reflectance contains a single dominant lobe projected on the half-vector (the bisector of lighting and viewing direction). Under this assumption, we perform a pixel-wise 1-D search for the elevation angle in the range $[0, \pi/2]$, given an azimuth angle for each pixel. We prove that the reflectance monotonicity is maintained only when the

correct elevation angle is found, if the scene is observed under dense directional lights uniformly distributed over the whole hemisphere.

Based on the evaluations using MERL BRDF database [5] and the conclusion in [4], we show that many materials such as acrylic, phenolic, metallic-paint, and some shiny plastics can be well approximated by the 1-lobe half-vector BRDF model. While some other materials, such as fabric, require two or more lobes for precise representation, in practice, our algorithm still shows robustness against the deviations and performs accurate estimation when a dominant lobe projected on half-vector exists. We assess the applicability of our method to 2-lobe BRDFs as well and verify the effectiveness of the proposed method on various real data.

2 Related Works

Conventional photometric stereo algorithm [1] assumes Lambert’s reflectance and can recover surface normal directions from as few as three images. With additional images, non-Lambertian phenomenon such as shadow can be handled [6]. With more images, various robust techniques can be applied to statistically handle non-Lambertian outliers. There are approaches based on RANSAC [7], median filtering [8] and rank minimization [9]. Spatial information has also been used for robustly solving the problem by using expectation maximization [10]. Some photometric stereo methods explicitly model surface reflectances with parametric BRDFs. Georgiades [11] use the Torrance-Sparrow model [12] for specular fitting, and Goldman *et al.* [13] build their methods on the Ward model [2].

For surfaces with general reflectance, various properties have been exploited to estimate surface normal, such as radiance similarity [14] and attached shadow codes [15]. Higo *et al.*’s method [16] uses monotonicity and isotropy of general diffuse reflectances. Alldrin and Kriegman [3] show that the azimuth angle of normal can be reliably estimated for isotropic materials using reflectance symmetry, if lights locate on a view-centered ring. Later, their method is extended to solve for both shape and reflectance by further restricting the BRDF to be bivariate [17]. Based on the symmetry property, a comprehensive theory is developed in [18], and a surface reconstruction method using a special lighting rig is introduced in [19]. By focusing on the low-frequency reflectance, the biquadratic model [20] can also be applied to solve the problem here.

We exploit reflectance monotonicity to compute the per-pixel elevation angle of surface normal given estimated azimuth angles using [3]. Unlike [17]’s approach that involves complex optimization for iteratively estimating shape and reflectance, our method avoids the complex optimization by taking the advantages of the monotonicity from the 1-lobe BRDFs.

3 Elevation Angle Estimation

Assume that we have already known the azimuth angles of surface normal. For example, we can use the method in [3] to obtain the azimuth angles first, while our method is not limited to a particular azimuth angle estimation method.




Fig. 1. Coordinate system and key variables

Given the azimuth angles, our method performs 1-D search for determining the elevation angles ranging from 0 to $\pi/2$ for visible normals.

3.1 1-lobe BRDF

We use bold letters to denote unit 3-D vectors. \mathbf{l} and \mathbf{v} represent the lighting and viewing directions, respectively, and \mathbf{h} is their unit bisector. For photometric stereo, we use the view-centered coordinate system with $\mathbf{v} = (0, 0, 1)^T$. As illustrated in Fig. 1, we use spherical coordinates (θ, ϕ) and (θ_l, ϕ_l) to represent surface normal \mathbf{n} and lighting direction \mathbf{l} respectively.

The semi-parametric model proposed in [4] suggests that BRDFs can be well represented by the summation of several lobes in the form of $\rho(\mathbf{n}^T \mathbf{k})$, where $\rho(\cdot)$ is a monotonically increasing function, and \mathbf{k} is referred as a projection direction. We further assume a dominant lobe (weighting is larger than other lobes) of $\rho(\mathbf{n}^T \mathbf{h})$ exists. We analyze the monotonicity of this lobe, and use this property for determining elevation angles.

3.2 BRDF Profile

At a scene point, each hypothesized elevation angle θ' uniquely determines a normal direction \mathbf{n}' , which in turn leads to multiple hypothesized BRDF values by dividing the observed scene radiance with the foreshortening term $\mathbf{n}'^T \mathbf{l}$. For simplicity, we refer to these BRDF values across varying $\mathbf{n}'^T \mathbf{h}$ as a *BRDF profile*. Given the observed scene radiance i at a pixel, the hypothesized BRDF profile ρ' is computed as

$$\rho'(\mathbf{n}'^T \mathbf{h}) = \frac{i}{\mathbf{n}'^T \mathbf{l}} = \rho(\mathbf{n}^T \mathbf{h}) \frac{\mathbf{n}^T \mathbf{l}}{\mathbf{n}'^T \mathbf{l}}. \quad (1)$$

We prove that ρ' is only monotonic w.r.t. $\mathbf{n}'^T \mathbf{h}$ when \mathbf{n}' is the correct surface normal direction (*i.e.*, θ' is the correct elevation angle θ), except for some degenerate lighting configurations. Based on this, we find the correct elevation angle θ using the monotonicity of the BRDF profile ρ' .

3.3 Monotonicity of the BRDF Profile

Consider two different lighting directions \mathbf{l}_1 and \mathbf{l}_2 , and their associated half-vectors \mathbf{h}_1 and \mathbf{h}_2 . Without loss of generality, we assume a surface normal with the azimuth angle $\phi = 0$ and elevation angle θ , *i.e.*, $\mathbf{n} = (\cos \theta, 0, \sin \theta)^T$. The hypothesized elevation angle θ' leads to $\mathbf{n}' = (\cos \theta', 0, \sin \theta')^T$. For notation simplicity, we define $x_1 = \mathbf{n}^T \mathbf{h}_1, x_2 = \mathbf{n}^T \mathbf{h}_2, x'_1 = \mathbf{n}'^T \mathbf{h}_1, x'_2 = \mathbf{n}'^T \mathbf{h}_2$ and $y_1 = \rho(x_1), y_2 = \rho(x_2), y'_1 = \rho'(x'_1), y'_2 = \rho'(x'_2)$.

If ρ is a monotonically increasing function¹, the following condition holds: for $x_1 < x_2$, we have $y_1 < y_2$. There are two cases when ρ' becomes non-monotonic:

1. The ordering of x is swapped, but y is not swapped: $x'_1 > x'_2$ and $y'_1 < y'_2$;
2. The ordering of y is swapped, but x is not swapped: $x'_1 < x'_2$ and $y'_1 > y'_2$.

In the following, we first discuss the conditions for reordering of x and y respectively (which we call x -swap and y -swap for short hereafter), and then analyze under what condition an incorrect estimate of the elevation angle θ' will break the monotonicity of ρ' .

Note here we focus on the case of two observations under lighting directions \mathbf{l}_1 and \mathbf{l}_2 . In a photometric stereo setting, we often have far more than two input observations. The discussion applies to any pair of observations: ρ' becomes non-monotonic, if any observation pair breaks its monotonicity.

Necessary and Sufficient Condition for x -swap. Suppose that the ordering of x is changed by the hypothesized θ' , *i.e.*, $(x_1 - x_2)(x'_1 - x'_2) = (\mathbf{n}^T \mathbf{h}_1 - \mathbf{n}^T \mathbf{h}_2)(\mathbf{n}'^T \mathbf{h}_1 - \mathbf{n}'^T \mathbf{h}_2) < 0$. From the definition of \mathbf{h} and $\mathbf{v} = (0, 0, 1)^T$, we have

$$\mathbf{h} = \frac{\mathbf{l} + \mathbf{v}}{\|\mathbf{l} + \mathbf{v}\|} = \frac{(l_x, l_y, l_z + 1)^T}{\sqrt{l_x^2 + l_y^2 + (l_z + 1)^2}} = \frac{(l_x, l_y, l_z + 1)^T}{\sqrt{2 + 2l_z}}. \quad (2)$$

After some simple derivations, we have

$$\mathbf{n}^T \mathbf{h}_1 = \frac{1}{\sqrt{2 + 2l_{1z}}} ((l_{1z} + 1) \sin \theta + l_{1x} \cos \theta). \quad (3)$$

$\mathbf{n}^T \mathbf{h}_2$ can be computed in a similar way, and their difference becomes

$$\mathbf{n}^T \mathbf{h}_1 - \mathbf{n}^T \mathbf{h}_2 = A \sin \theta + B \cos \theta = \sqrt{A^2 + B^2} \sin(\theta + \alpha), \quad (4)$$

where $A = \frac{l_{1z}+1}{\sqrt{2+2l_{1z}}} - \frac{l_{2z}+1}{\sqrt{2+2l_{2z}}}$, $B = \frac{l_{1x}}{\sqrt{2+2l_{1z}}} - \frac{l_{2x}}{\sqrt{2+2l_{2z}}}$, and $\alpha = \arctan \frac{B}{A}$, $\alpha \in [-\pi/2, \pi/2]$. We obtain a similar equation for \mathbf{n}' by replacing θ in Eq. (4) with θ' .

Therefore the necessary and sufficient condition to change the ordering of x is $(\mathbf{n}^T \mathbf{h}_1 - \mathbf{n}^T \mathbf{h}_2)(\mathbf{n}'^T \mathbf{h}_1 - \mathbf{n}'^T \mathbf{h}_2) = \sin(\theta + \alpha) \sin(\theta' + \alpha) < 0$. In other words,

¹ We discuss only the monotonically increasing case in this paper. The similar analysis can also be applied to monotonically decreasing case.

$\sin(\theta + \alpha)$ and $\sin(\theta' + \alpha)$ should have different signs. This condition is true only when $\alpha < 0$ and $\theta < -\alpha < \theta'$ (or, $\theta' < -\alpha < \theta$), since θ and θ' are both within $[0, \pi/2]$, and α is within $[-\pi/2, \pi/2]$. Note that α is completely determined by the two lighting directions \mathbf{l}_1 and \mathbf{l}_2 , and independent of θ and θ' . Hence, a larger difference between θ and θ' gives higher possibility for an x -swap to happen.

Sufficient Condition for y -swap. Next, we consider when the ordering of y will be changed, *i.e.*, we discuss the case where $y_1 < y_2$ and $y'_1 > y'_2$. We assume that \mathbf{l}_1 is close to \mathbf{l}_2 such that $\rho(\mathbf{n}^T \mathbf{h}_1)$ is close to $\rho(\mathbf{n}^T \mathbf{h}_2)$. This is true when lighting is dense and neighboring lighting directions are sampled, and lighting directions do not cause the highlight reflection at \mathbf{n} . Note that by focusing on nearby lighting directions, we can only derive the *sufficient* condition for swapping the ordering of y , because there could be two parted lights causing y -swap to happen. Under these assumptions and according to Eq. (1), y' is mainly determined by $\mathbf{n}^T \mathbf{l} / \mathbf{n}^T \mathbf{l}'$. For \mathbf{l}_1 , we obtain

$$\frac{\mathbf{n}^T \mathbf{l}_1}{\mathbf{n}^T \mathbf{l}_1} = \frac{l_{1x} \cos \theta + l_{1z} \sin \theta}{l_{1x} \cos \theta' + l_{1z} \sin \theta'} = \frac{\sin(\theta + \beta_1)}{\sin(\theta' + \beta_1)}, \quad \tan \beta = \frac{l_{1x}}{l_{1z}}. \quad (5)$$

From the relationship $\frac{\sin(\theta + \beta_1)}{\sin(\theta' + \beta_1)} > \frac{\sin(\theta + \beta_2)}{\sin(\theta' + \beta_2)} > 0$, we can derive the sufficient condition for y -swap as

$$(\theta' - \theta) \mathbf{l}_{12y} < 0. \quad (6)$$

Here, \mathbf{l}_{12y} is y -component of the cross product of \mathbf{l}_1 and \mathbf{l}_2 . The derivation can be found in the Appendix. Eq. (6) indicates that the ordering of y depends on two factors: (1) the sign of $(\theta' - \theta)$, *i.e.*, the hypothesized θ' is larger or smaller than the true value, and (2) the relative positions of the two lights \mathbf{l}_1 and \mathbf{l}_2 .

Sufficient Condition for Unique Solution. For any hypothesized θ' , if it swaps the ordering of y (or, x) while keeps the ordering of x (or, y) unchanged, the function ρ' will become non-monotonic. Hence, if we can ensure an incorrect θ' will always break the monotonicity of ρ' , we can find the correct elevation angle θ by choosing the θ' that makes ρ' monotonic.

For a monotonically increasing ρ , we have $x_1 = \mathbf{n}^T \mathbf{h}_1 < x_2 = \mathbf{n}^T \mathbf{h}_2$, or equivalently $\theta_{h1} > \theta_{h2}$. We begin with the case where lights are densely distributed along the same longitude as the normal \mathbf{n} , *i.e.*, $\phi_l = \phi = 0$. Since we require $\theta_{h1} > \theta_{h2}$, \mathbf{h}_2 is closer to \mathbf{n} than \mathbf{h}_1 . Further, as illustrated in Fig. 2, \mathbf{h}_1 and \mathbf{h}_2 are restricted on the red dotted line when the lighting directions \mathbf{l}_1 and \mathbf{l}_2 move on the green dotted line. By observing the geometric relationship between \mathbf{h} and \mathbf{l} , we obtain the following results:

1. As shown in Fig. 2(a), if $\theta \leq \pi/4$, we always have $\theta_{l1} > \theta_{l2}$ (\mathbf{l}_1 is closer to \mathbf{v} than \mathbf{l}_2). This ensures $\mathbf{l}_{12y} > 0$. Then from Eq. (6), y -swap always happens if the hypothesized θ' becomes $\theta' < \theta$. In contrast, y -swap might not² happen when $\theta' > \theta$;

² Because we only have the sufficient condition for y -swap.




Fig. 2. Lights, half-vectors and normal on the same longitude. (a) $\theta \leq \pi/4$; (b) $\theta > \pi/4$.




Fig. 3. (a) α values for $\phi_l = 0$ and $\theta_{l1}, \theta_{l2} \in [0, \pi/2]$; (b) $\max(\alpha)$ values for $\phi_l \in (0, \pi/2]$.

2. As shown in Fig. 2(b), if $\theta > \pi/4$, the relative positions of \mathbf{l}_1 and \mathbf{l}_2 cannot be determined from $\theta_{h1} > \theta_{h2}$. There are two cases: (1) If \mathbf{h}_1 and \mathbf{h}_2 are both farther from \mathbf{v} than \mathbf{n} , we should have $\theta_{l2} > \theta_{l1}$ (\mathbf{l}_2 is closer to \mathbf{v} than \mathbf{l}_1). Hence, $\mathbf{l}_{12y} < 0$; (2) If \mathbf{h}_1 and \mathbf{h}_2 are both closer to \mathbf{v} than \mathbf{n} , we should have $\theta_{l1} > \theta_{l2}$ (\mathbf{l}_1 is closer to \mathbf{v} than \mathbf{l}_2 , which is a similar case as Fig. 2(a)). Hence, $\mathbf{l}_{12y} > 0$. When lights are densely distributed along the longitude, we can always find a pair of lights satisfying these two cases. As a result, Eq. (6) can hold to cause y -swap no matter what θ' is.

Now we know y -swap might not happen when $\theta \leq \pi/4$ and $\theta' > \theta$. In the next, we analyze when x -swap happens. If both swaps do not happen, an incorrect estimate of the elevation angle θ' will not break the monotonicity of ρ' , and we cannot tell whether that θ' is correct or not. Due to the complexity of the analytic form of α , we simulate and plot all α values for all combinations of \mathbf{l}_1 and \mathbf{l}_2 on the same longitude as \mathbf{n} in Fig. 3(a). It is interesting to note that α continuously changes in $[-\pi/2, \max(\alpha)]$, where $\max(\alpha) \approx -\pi/4$. In other words, $-\alpha$ covers the whole range of $[-\max(\alpha), \pi/2]$. Recall that the necessary and sufficient condition for x -swap to happen is that $\theta' < -\alpha < \theta$ (or, $\theta < -\alpha < \theta'$). Thus, if both θ' and θ are smaller than $-\max(\alpha)$, x -swap will never happen.

Combining the conditions that both x -swap and y -swap do not happen, we conclude $\theta' \in [\theta, -\max(\alpha)]$ is the degenerate interval for normals with $\phi_l = 0$, where the monotonicity of ρ' is not broken by any incorrect θ' .

Next, we consider lights along other longitudes, *i.e.* $\phi_l \in (0, \pi/2]$. The same analysis can be applied, and similarly, the monotonicity of ρ' is not broken by any incorrect θ' within the degenerate interval $[\theta, -\max(\alpha)_{\phi_l}]$. Here, $-\max(\alpha)_{\phi_l}$ depends on the azimuth angle of lighting directions at the longitude ϕ_l . We plot it as a function of ϕ_l for those longitudes in Fig. 3(b). We find that $\max(\alpha)$ increases from about $-\pi/4$ to 0 when ϕ_l approaches $\pi/2$. Since monotonicity of ρ' is lost if any pair of lights on a longitude can break its monotonicity, the final degenerate interval is the intersection of all degenerate intervals $[\theta, -\max(\alpha)_{\phi_l}]$ from each longitude. This makes the final intersection of the interval an empty set, because $-\max(\alpha)_{\phi_l}$ approaches 0. In other words, if we have lights on all longitudes with azimuth angles $\phi_l \in [0, \pi/2]$, we can uniquely determine the elevation angle θ for normals with $\phi = 0$.

To recover the elevation angle of arbitrary surface normals, we will need lights spanned all longitudes that cover the whole hemisphere. These lights form a dense and uniform distribution over the hemisphere. In practice, we capture images under random lighting directions that approximate the uniformity.

4 Solution Method

We perform a 1-D search for θ within $[0, \pi/2]$ at each pixel. For each hypothesized value θ' , we assess the reflectance monotonicity. Specifically, we compute the hypothesized BRDF values y' and evaluate its monotonicity w.r.t. to x' . To measure the reflectance monotonicity, we calculate the derivatives (discrete differentiation) of y' and sum up the absolute values of negative derivatives, denoted as $\delta(\theta')$:

$$\delta(\theta') = \sum_{x'} \max\left(-\frac{dy'}{dx'}, 0\right), \quad (7)$$

The cost δ penalizes monotonically decreasing sequence, *i.e.*, a larger δ value indicates a less monotonically increasing y' . We show a typical cost function in Fig. 4 for a normal with $\theta = \pi/6$ and $\phi = 0$. Fig. 4(a) is the cost computed where all lights locate on the same longitude as the normal. As we have proved in Sec. 3.3, there is a degenerate interval for $\theta' \in [\theta, -\max(\alpha)]$. Indeed, our cost function is almost a constant value in $[\pi/6, \pi/4]$, and we cannot tell which value within this interval is the correct elevation angle. Fig. 4(b) shows the cost computed with lights distributed on the whole hemisphere. It has a clear global minimum to estimate the elevation angle.

The complete normal estimation algorithm is summarized as Algorithm 1. There are a few implementation details: (1) We discard shadows by simple thresholding. For scene radiance values normalized to 1, we use a threshold $i_0 = 10^{-6}$ for synthetic data, and in real data the threshold is manually chosen through cross validation (typically set as 0.02 in our experiments); (2) When calculating y' , we set it as a big value ($i_\infty = 10^{10}$), if $\mathbf{n}'^T \mathbf{l} \leq 0$; (3) When evaluating the monotonicity, we apply a monotonic mapping as $y' \leftarrow y'^\gamma$ (γ is empirically determined as 5 in all experiments) for a data normalization purpose.

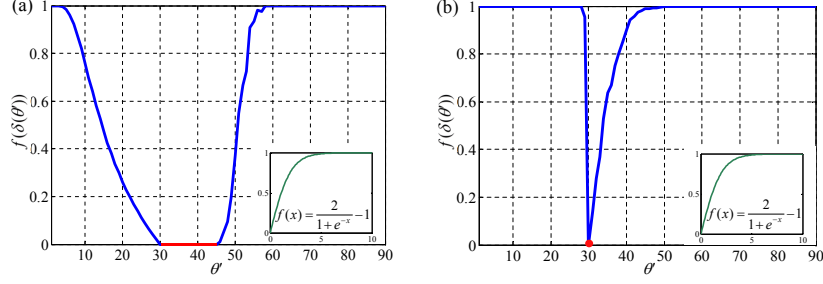


Fig. 4. Cost function for $\theta' \in [0^\circ, 90^\circ]$ (normalized by $f(\cdot)$ shown in the right bottom for visualization purpose). (a) Using lights on the same longitude as normal; (b) Using lights covering the whole hemisphere.

Algorithm 1 Photometric Stereo for General Isotropic Reflectances

INPUT: Scene radiance values i , lighting directions \mathbf{l} , threshold i_0, i_∞ .

Compute the azimuth angle of normal ϕ using the method in [3];

for each pixel **do**

for $\theta' \in [0, \pi/2]$ **do**

 Let $\mathbf{n}' = (\cos \theta' \cos \phi, \cos \theta' \sin \phi, \sin \theta')^T$;

 Calculate y' based on Eq. (1);

 If any \mathbf{l} causes $y \leq i_0$ (in shadow), set its corresponding $y' = i_0$;

 If any $\mathbf{n}'^T \mathbf{l} \leq 0$, set its corresponding $y' = i_\infty$;

 Order $y' = i / (\mathbf{n}'^T \mathbf{l})$ w.r.t. $x' = \mathbf{n}'^T \mathbf{h}$;

 Evaluate the cost values using Eq. (7);

end for

$\theta = \underset{\theta'}{\operatorname{argmin}} \delta(\theta')$;

$\mathbf{n} = (\cos \theta \cos \phi, \cos \theta \sin \phi, \sin \theta)^T$;

end for

OUTPUT: Normals for all pixels.

5 Experiments

We evaluate the accuracy of elevation angle estimation using all the 100 materials in the MERL BRDF database [5], also with varying number of lights. Synthetic experiments using 2-lobe BRDFs are performed to assess the robustness of our method against the deviation from the 1-lobe BRDF assumption. Finally, we show our normal estimates on real data.

5.1 Synthetic Data

Performance on Measured BRDFs. We sample 1620 normal directions from the visible hemisphere by uniformly choosing 36 longitudes and 45 altitudes. We generate their observations under 337 lighting directions uniformly sampled on the hemisphere. For the sampling, we use the vertices defined by an icosahedron

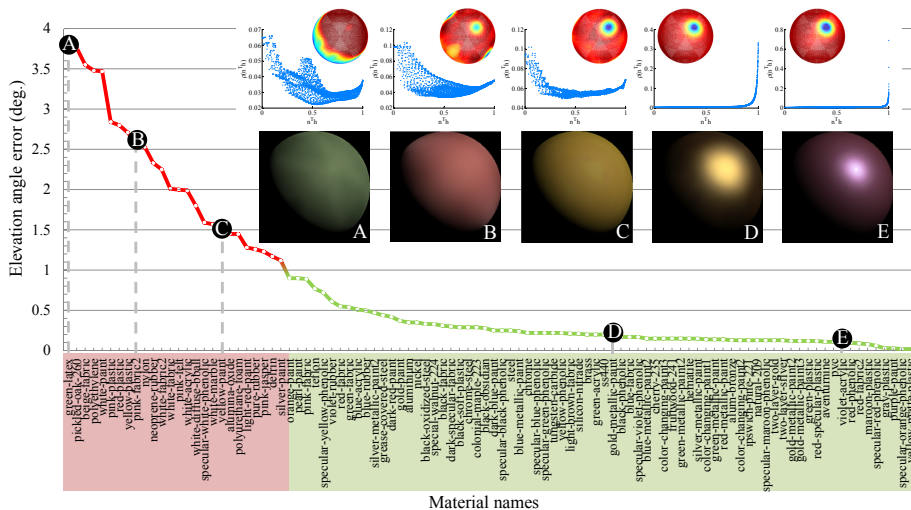


Fig. 5. Elevation angle errors (degree) on 100 materials (ranked by errors in a descending order). Rendered spheres of some representative materials and their BRDF values in the $\rho\text{-}\mathbf{n}^T\mathbf{h}$ space are shown near the curve. The spheres in the $\rho\text{-}\mathbf{n}^T\mathbf{h}$ plots show the 1-D projection RMS errors using all the directions on the hemisphere (dark blue means small, and red means large errors).

with a tessellation order of three. The average elevation angle errors of the 1620 samples for all materials in the MERL BRDF database are summarized in Fig. 5. The average error on all the 100 materials is 0.77° from the ground truth.

We observe that the error is relatively large for materials that cannot be well represented by a monotonic 1-lobe BRDF with half-vector as the projection direction, for example the material A and B in Fig. 5 (see the $\rho\mathbf{-n}^T\mathbf{h}$ plot above the rendered spheres). For those materials, their optimal 1-lobe BRDF representations have a much different projection direction from \mathbf{h} (see the spheres of RMS error distribution in the $\rho\mathbf{-n}^T\mathbf{h}$ plots). Therefore, if we force them to project on \mathbf{h} and apply our method, the results have relatively larger errors.

Performance with Varying Number of Lights. Next, we evaluate the system performance variation with different numbers of lighting directions (input images). We perform the same experiment using 25, 50, 100 and 200 random lighting directions. The errors and light distributions are shown in Fig. 6. Empirically, about 100 random lights provide average elevation angle error of around 1° . Thus in the following test, we fix the number of lights to 100.

Performance on 2-lobe BRDFs. To demonstrate the robustness of our method, we evaluate our method on synthetic data using 2-lobe BRDFs in Fig. 7. The weighting of the two lobes k_1 and k_2 are varied from 0 to 1. Fig. 7(a) shows




Fig. 6. Average elevation angle errors (degree) on 100 materials varying with the number of lighting directions. Next to the curve, the blue dots on spheres show the light distribution in each case.




Fig. 7. Elevation angle errors (degree) on 2-lobe BRDFs. (a) The Cook-Torrance model, where k_1 is the Lambertian and k_2 is the specular strength; (b) $k_1\rho_1(\mathbf{n}^T\mathbf{v}) + k_2\rho_2(\mathbf{n}^T(\mathbf{v} + 2\mathbf{l}))$; (c) $k_1\rho_1(\mathbf{n}^T\mathbf{h}) + k_2\rho_2(\mathbf{n}^T\mathbf{k})$, where \mathbf{k} is a random direction.

our result on the Cook-Torrance model [21] (roughness $m = 0.5$) with a Lambertian diffuse lobe and a specular lobe. Note the specular lobe of the Cook-Torrance model is not centered at \mathbf{h} . However, our method always gives accurate result for different relative strength of the two lobes. The error in estimated elevation angles is the largest (about 1°) when BRDF is completely dominated by the specular lobe.

As discussed in [4], some fabric materials contain lobes with projection direction \mathbf{v} or $(\mathbf{v} + 2\mathbf{l})$. Hence, we deliberately create such a BRDF as $k_1\rho_1(\mathbf{n}^T\mathbf{v}) + k_2\rho_2(\mathbf{n}^T(\mathbf{v} + 2\mathbf{l}))$ and evaluate our method on this BRDF. Note both lobes are not centered at \mathbf{h} . Here, $\rho_1(x) = \rho_2(x) = x$. In Fig. 7(b), we plot the elevation angle errors for different values of k_1, k_2 . The errors are smaller than 3° for most of the combinations.

At last, we create a 2-lobe BRDF $k_1\rho_1(\mathbf{n}^T\mathbf{h}) + k_2\rho_2(\mathbf{n}^T\mathbf{k})$ by combining a lobe centered at \mathbf{h} and another one with a random direction \mathbf{k} . \mathbf{k} is a randomly sampled direction on the sphere for each pixel under each lighting direction. The elevation angle errors are shown in Fig. 7(c). We can observe that for more than half of all cases, our method generates errors smaller than 5° . Our method can work reasonably well with errors smaller than 3° when the relative strength of the random lobe is below 0.3.




Fig. 8. Results on GOURD1 (102) (The number in the bracket indicates the number of images in the dataset.) [17]. (a) One input image; (b), (c) Our normal and Lambertian shading; (d), (e) Normal and shading from Lambertian photometric stereo. Note that (b) and (d) have an average angular difference of about 12° .




Fig. 9. Results on APPLE (112) [17]. (a) One input image; (b), (c) Our normal and reconstructed surface; (d), (e) Normal and surface shown in paper [17]. Note here we use the color mapping of $((n_x + 1)/2, (n_y + 1)/2, n_z) \rightarrow (R, G, B)$ for a comparison purpose. The shapes look different partially due to we do not know exactly the same reconstruction method and rendering parameters used in [17].

5.2 Real Data

We show the estimated normals by our method on real data. We visualize the estimated normals by linearly encoding the x , y and z components of normals in RGB channels of an image (except for the APPLE data in Fig. 9 for a comparison purpose). First, in Fig. 8, our method is compared with the Lambertian photometric stereo [1] by showing the estimated normals and the Lambertian shadings calculated from the estimated normals with the same lighting direction as the input image. For such non-Lambertian materials, our method shows much more reasonable normal estimates. Next, we compare with the method in [17] by showing the surface reconstructions (according to the method in [22]) from the estimated normals in Fig. 9(b). Due to the lack of the ground truth, we cannot make a quantitative comparison, but qualitatively, our method shows similar results as [17]. In terms of the complexity, our method is much simpler and computationally inexpensive for deriving the elevation angle.

Finally, we show our estimated normals on other materials with various reflectances, such as plastic, metal, paint, *etc.*, as shown in Fig. 10. The datasets on the left side are from [17] and [23]; right part of the datasets are captured by ourselves with a Sony XCD-X710CR camera. By comparing the input images and the Lambertian shadings, we claim our estimated normals are of high ac-




Fig. 10. Real data results. Left part: GOURD2 (98) [17], HELMET_SIDE_RIGHT (119) (We use only 119 out of all the 253 images in the original dataset.) [23], KNEELING_KNIGHT (119) [23]; Right part: PEAR (77), GOD (57), DINOSAUR (118). We show one input image, the estimated normals and the Lambertian shadings for each dataset.

curacy. Some noisy points observed on the results are mainly caused by pixels with too dark intensities and hence low signal-to-noise ratio.

6 Conclusions

We show a method for estimating elevation angles of surface normal by exploiting reflectance monotonicity given the azimuth angles. We assume the BRDF contains a dominant monotonic lobe projected on the half-vector, and prove that the elevation angle can be uniquely determined under dense lights uniformly distributed on the hemisphere. In synthetic experiments, we first demonstrate the accuracy of our method on a broad category of materials. We further evaluate its robustness to deviations from our assumption about BRDFs. Various real-data experiments also show the effectiveness of the proposed method.

Limitations. Our method assumes known azimuth angles. Joint estimation of both azimuth and elevation angles makes the problem prohibitively difficult due to its highly non-linear nature of the problem. When the assumed azimuth angles are not accurate, the elevation angle estimation will be deteriorated accordingly. For example, in the experiment of Fig. 5, the average elevation angle errors over 100 materials increase to $\{1.16, 2.24, 3.32, 4.72\}$ degrees with additive azimuth angle errors normally distributed with mean 0 and standard deviation $\{0.5, 1, 1.5, 2\}$ degrees, respectively. To make the proposed method robust against

inaccuracy of azimuth angles is an interesting direction. Besides, while we have discussed only the sufficient condition, deriving a compact lighting configuration that uniquely determines the elevation angles is our future work.

Acknowledgement

The authors would like to thank Neil Alldrin for providing the code of [3] and Manmohan Chandraker for helpful discussions. Boxin Shi and Katsushi Ikeuchi was partially supported by the Digital Museum Project, Ministry of Education, Culture, Sports, Science & Technology in Japan. Ping Tan was supported by the Singapore ASTAR grant R-263-000-592-305 and MOE grant R-263-000-555-112.

References

1. Woodham, R.: Photometric method for determining surface orientation from multiple images. *Optical Engineering* **19** (1980) 139–144
2. Ward, G.: Measuring and modeling anisotropic reflection. *Computer Graphics* **26** (1992) 265–272
3. Alldrin, N., Kriegman, D.: Toward reconstructing surfaces with arbitrary isotropic reflectance: A stratified photometric stereo approach. In: *Proc. of International Conference on Computer Vision (ICCV)*. (2007)
4. Chandraker, M., Ramamoorthi, R.: What an image reveals about material reflectance. In: *Proc. of International Conference on Computer Vision (ICCV)*. (2011)
5. Matusik, W., Pfister, H., Brand, M., McMillan, L.: A data-driven reflectance model. In: *ACM Trans. on Graph. (Proc. of SIGGRAPH)*. (2003)
6. Barsky, S., Petrou, M.: The 4-source photometric stereo technique for three-dimensional surfaces in the presence of highlights and shadows. *IEEE Trans. Pattern Anal. Mach. Intell.* **25** (2003) 1239–1252
7. Mukaigawa, Y., Ishii, Y., Shakunaga, T.: Analysis of photometric factors based on photometric linearization. *Journal of the Optical Society of America* **24** (2007) 3326–3334
8. Miyazaki, D., Hara, K., Ikeuchi, K.: Median photometric stereo as applied to the segonko tumulus and museum objects. *International Journal of Computer Vision* **86** (2010) 229–242
9. Wu, L., Ganesh, A., Shi, B., Matsushita, Y., Wang, Y., Ma, Y.: Robust photometric stereo via low-rank matrix completion and recovery. In: *Proc. of Asian Conference on Computer Vision (ACCV)*. (2010)
10. Wu, T., Tang, C.: Photometric stereo via expectation maximization. *IEEE Trans. Pattern Anal. Mach. Intell.* **32** (2010) 546–560
11. Georgiades, A.: Incorporating the Torrance and Sparrow model of reflectance in uncalibrated photometric stereo. In: *Proc. of International Conference on Computer Vision (ICCV)*. (2003)
12. Torrance, K., Sparrow, E.: Theory for off-specular reflection from roughened surfaces. *Journal of the Optical Society of America* **57** (1967) 1105–1114
13. Goldman, D., Curless, B., Hertzmann, A., Seitz, S.: Shape and spatially-varying BRDFs from photometric stereo. *IEEE Trans. Pattern Anal. Mach. Intell.* **32** (2010) 1060–1071

14. Sato, I., Okabe, T., Yu, Q., Sato, Y.: Shape reconstruction based on similarity in radiance changes under varying illumination. In: Proc. of International Conference on Computer Vision (ICCV). (2007)
15. Okabe, T., Sato, I., Sato, Y.: Attached shadow coding: estimating surface normals from shadows under unknown reflectance and lighting conditions. In: Proc. of International Conference on Computer Vision (ICCV). (2009)
16. Higo, T., Matsushita, Y., Ikeuchi, K.: Consensus photometric stereo. In: Proc. of IEEE Conference on Computer Vision and Pattern Recognition (CVPR). (2010)
17. Alldrin, N., Zickler, T., Kriegman, D.: Photometric stereo with non-parametric and spatially-varying reflectance. In: Proc. of IEEE Conference on Computer Vision and Pattern Recognition (CVPR). (2008)
18. Tan, P., Quan, L., Zickler, T.: The geometry of reflectance symmetries. IEEE Trans. Pattern Anal. Mach. Intell. **33** (2011) 2506–2520
19. Chandraker, M., Bai, J., Ramamoorthi, R.: A theory of photometric reconstruction for unknown isotropic reflectances. In: Proc. of IEEE Conference on Computer Vision and Pattern Recognition (CVPR). (2011)
20. Shi, B., Tan, P., Matsushita, Y., Ikeuchi, K.: A biquadratic reflectance model for radiometric image analysis. In: Proc. of IEEE Conference on Computer Vision and Pattern Recognition (CVPR). (2012)
21. Cook, R., Torrance, K.: A reflectance model for computer graphics. ACM Trans. on Graph. **1** (1982) 7–24
22. Kovesi, P.: Shapelets correlated with surface normals produce surfaces. In: Proc. of International Conference on Computer Vision (ICCV). (2005)
23. Einarsson, P., et al.: Relighting human locomotion with flowed reflectance fields. In: Proc. of Eurographics Symposium on Rendering. (2006)

Appendix: Proof of Eq. (6)

Proof. If $\mathbf{n}^T \mathbf{l} > 0$ and $\mathbf{n}'^T \mathbf{l} > 0$, from Eq. (5), the following holds:

$$\frac{\sin(\theta + \beta_1)}{\sin(\theta' + \beta_1)} > \frac{\sin(\theta + \beta_2)}{\sin(\theta' + \beta_2)} > 0.$$

Therefore, we have

$$\sin(\theta + \beta_1) \sin(\theta' + \beta_2) - \sin(\theta' + \beta_1) \sin(\theta + \beta_2) > 0.$$

By applying the product-to-sum trigonometric identities, it is simplified as

$$\cos(\theta - \theta' + \beta_1 - \beta_2) - \cos(\theta' - \theta + \beta_1 - \beta_2) > 0.$$

This further becomes as following by applying the sum-to-product identities:

$$\sin(\theta' - \theta) \sin(\beta_1 - \beta_2) > 0.$$

Since $\theta, \theta' \in [0, \pi/2]$, we can use $(\theta' - \theta)$ to replace $\sin(\theta' - \theta)$ with inequality retained. From the definition of $\tan \beta$, we have the following relation

$$\begin{aligned} \sin(\beta_1 - \beta_2) &= \sin(\beta_1) \cos(\beta_2) - \cos(\beta_1) \sin(\beta_2) \\ &= \frac{l_{1x}}{\sqrt{l_{1x}^2 + l_{1z}^2}} \frac{l_{2z}}{\sqrt{l_{2x}^2 + l_{2z}^2}} - \frac{l_{1z}}{\sqrt{l_{1x}^2 + l_{1z}^2}} \frac{l_{2x}}{\sqrt{l_{2x}^2 + l_{2z}^2}}. \end{aligned}$$

This indicates the sign of $\sin(\beta_1 - \beta_2)$ is the same as $l_{1x}l_{2z} - l_{1z}l_{2x} = \mathbf{l}_2 \times \mathbf{l}_1 = -\mathbf{l}_{12y}$. Therefore, the inequality of Eq. (6) holds. \square