A Microfacet-based Reflectance Model for Photometric Stereo with Highly Specular Surfaces

Lixiong Chen *[1], Yinqiang Zheng[1], Boxin Shi[2], Art Subpa-Asa[3], and Imari Sato[1]

[1]National Institute of Informatics, Tokyo, Japan
[2]Artificial Intelligence Research Center, National Institute of AIST, Tokyo, Japan
[3]Tokyo Institute of Technology, Tokyo, Japan

Abstract

A precise, stable and invertible model for surface reflectance is the key to the success of photometric stereo with real world materials. Recent developments in the field have enabled shape recovery techniques for surfaces of various types, but an effective solution to directly estimating the surface normal in the presence of highly specular reflectance remains elusive. In this paper, we derive an analytical isotropic microfacet-based reflectance model, based on which a physically interpretable approximate is tailored for highly specular surfaces. With this approximate, we identify the equivalence between the surface recovery problem and the ellipsoid of revolution fitting problem, where the latter can be described as a system of polynomials. Additionally, we devise a fast, non-iterative and globally optimal solver for this problem. Experimental results on both synthetic and real images validate our model and demonstrate that our solution can stably deliver superior performance in its targeted application domain.

1. Introduction

The object appearance is a compound of illumination, surface reflectance and surface shape. Under directional light $\vec{l}$, the appearance I of a surface point with normal $\vec{n}$ observed from direction $\vec{v}$ is described as a product of shading signal $\vec{n}^\intercal\vec{l}$ and the reflectance signal specified by the Bidirectional Reflectance Distribution Function (BRDF) $\rho(\cdot)$:

$$I = \rho(\vec{v}, \vec{l}, \vec{n})\vec{n}^\intercal\vec{l}. \quad (1)$$

Photometric stereo aims at inversely recovering $\vec{n}$ given multiple observations with different lighting directions.

*Part of this work was finished when the author was a Master student of McGill University, Montreal, Canada, and was visiting National Institute of Informatics (NII), Japan, as an intern funded in part by the NII MOU internship program.

(a) chrome steel (b) metallic paint (c) two layer gold

Figure 1: Examples of highly specular reflectance targeted in this paper are pervasive in the real world, but effectively estimating the shape of the specular surfaces is a challenging task.

Therefore, finding a proper tradeoff between the expressiveness and complexity of Equation 1 is critical. For example, the Lambertian model [43] allows straightforward normal recovery, but a pure Lambertian surface rarely exists in real world objects. In recent years, the photometric stereo has seen dramatic development in shape recovery techniques for surfaces made of a great variety of materials [9, 15, 16, 22, 37, 44]. After all, sufficient know-how is already at hand to help extract shape cues from low-frequency reflectance. However, how to properly handle highly specular surfaces, like the ones illustrated in Figure 1, remains to be a hard nut. This is because when signals representing specular reflectance become predominant, knowledge about inferring shape solely from specularities for a typical photometric stereo setup is still limited.

On the contrary, how to render visually realistic specular/general surfaces has been extensively studied in the computer graphics community. For example, the microfacet reflectance model and its variants [4, 6, 10, 12, 25, 41] offer significant insight into the formation of specularity given the detailed knowledge of the surface geometry. Unfortunately, it is not straightforward to apply these models for the inverse problem that seeks to recover the surface geometry, because analysis is absent to warrant model invert-

ibility and estimation accuracy.

In this paper, we carefully investigate the state-of-the-art microfacet theory involving the ellipsoid microfacet normal distribution function (a.k.a. ellipsoid NDF), and have succeed in deriving an analytical model serviceable for photometric stereo. We further introduce a physical interpretable approximate that brings appealing algebraic properties for specular surface normal estimation, without sacrificing its expressiveness. Essentially, with this approximate, we identify that the calibrated photometric stereo problem boils down to an ellipsoid of revolution fitting problem, for which we devise a fast, non-iterative and globally optimal solver targeting a system of polynomials. We use empirical results on both synthetic and real images to justify our theory and model, and demonstrate that our solution outperforms its peers when applied to highly specular surfaces. Moreover, based on our model, we also discuss the possibility to devise a generalized solution to get the best of both worlds.

To sum up, our contributions mainly lie in

1. Deriving an analytical form based on the microfacet theory and the ellipsoid NDF for photometric stereo;
2. Developing a physically interpretable approximation for highly specular reflectance, which equates the problem of normal estimation with an ellipsoid of revolution fitting problem;
3. Designing a fast, non-iterative and globally optimal solver to stably obtain the normal of specular surfaces.

The remaining of this paper is organized as follows: Section 2 discusses the related works, Section 3 presents our analytical microfacet reflectance model, and its reduction for highly specular reflectance is derived in Section 4. Section 5 explains our normal recovery algorithm, then in Section 6 we discuss our experiment results obtained from both synthetic images and real images. Section 7 concludes this paper.

2. Related Works

Reflectance models play important roles in realistic rendering (forward problem) and shape inference using photometric stereo (inverse problem).

Reflectance models for rendering always seek to capture the finest real-world subtleties. Early models are either empirically based, where appearance is directly compared with observations [30, 23], or microfacet-based but in a constrained parametric form [36, 35, 7, 42, 28]. Recent works, such as the Ashikhmin-Shirley model for anisotropic materials [4] and the GGX distribution [8] for general rough surfaces, acknowledge the dominant effect of the NDF and leave it as a design choice.

A relevant line of research proposes to represent NDF directly with tabulated data [3]. This motivates the attempts to directly capture the real-world reflectance [25, 27], which in general requires dense sampling [14] and the ability to process large amount of data. One workaround is to leverage the prior knowledge of BRDF that leads to semi-parameterized fitting [11, 6]. Among the proposed models, we pay more attention to the ellipsoid NDF [40]. In case of isotropic reflectance, it is analogous to the microfacet BRDF with GGX/Trowbridge-Reitz distribution [41, 38]. It is worth noting that for rendering energy conservation is not warranted, but for finer appearance capturing dense light distribution is always desired [26], and the sampled surface is assumed to have regular geometry [12] (*e.g.* flat surface).

Reflectance models for photometric stereo seek to decompose surface geometry from the scene radiance, which usually contains highly complex surface reflectance. The Lambertian model is widely adopted for its simplicity [43]. To deal with more general materials, specular highlight could be discarded through outlier rejection [44, 22]. The general reflectance could also be modeled by parametric BRDF models [37, 15, 9], or a combined effect of several easy-to-model components [16].

Recent approaches adopt non-parametric formulation to handle a broader range of materials. Without explicit modeling of the BRDF, some general reflectance properties, such as isotropy [2], monotonicity [32], and their combination with visibility [19], are exploited to infer surface orientation. The BRDFs can also be explicitly represented as a bivariate function [1, 45], a constrained bivariate regression [21] or a sparse dictionary-based representation [20]. Overall, BRDF for photometric stereo is a delicate trade off between generality and complexity [33].

Benchmark evaluation [34] demonstrates that the state-of-the-art performance can be achieved with data containing less-specular observations [21, 33]. These approaches work well for a great diversity of real-world materials, but are challenged by specularity-dominant observations. In contrast, our method aims to attack this challenging problem with theoretical support.

3. A Microfacet BRDF with Ellipsoid Normal Distribution Function

The microfacet reflectance model postulates that the surface is made up by a large collection of tiny facets, and the surface radiance is essentially a composition of microfacet reflections, where the radiance intensity can be evaluated [10, 36, 5] as

$$I(\vec{v}) = \int_{\Omega_+} \max(\vec{m}\cdot\vec{l},0) D(\vec{m}) G(\vec{l},\vec{v}) \rho_{\vec{m}}(\vec{m},\vec{l},\vec{v})\, d\vec{m}. \tag{2}$$

As illustrated in Figure 2, Ω_+ denotes the the visible upper half sphere, $\rho_m(\vec{m},\vec{l},\vec{v})$ describes the reflectance of a

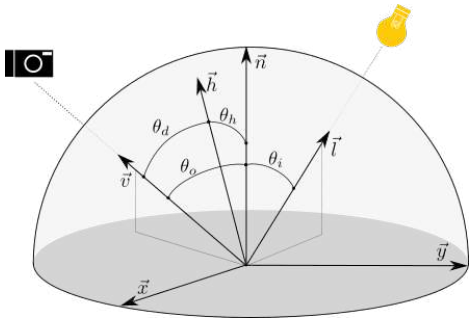

Figure 2: The coordinates in which BRDF is defined. By convention $\vec{n} = (0, 0, 1)$, and $\vec{v}$ and $\vec{l}$ are unit vectors that allow to orient arbitrarily above the positive half-sphere. Note that this is in contrast to the typical setup for photometric stereo, where $\vec{v} = (0, 0, 1)$

specific microfacet with normal $\vec{m}$ under directional light $\vec{l}$ while being perceived along $\vec{v}$, $D(\vec{m})$ is the microfacet Normal Distribution Function (NDF) counting the population of the microfacets of the same orientation, and $G(\vec{l}, \vec{v})$ is the masking-shadowing term ensuring power conservation. To model general reflectance, each microfacet can be effectively assumed to exhibit mirror reflection [41], namely, $\rho_m(\vec{m}, \vec{l}, \vec{v}) = F(\theta_d)\delta_{\vec{m}}(\vec{h})$ dictating that a microfacet contributes to the actual reflection only if its normal $\vec{m}$ and bisector $\vec{h} = \frac{\vec{l}+\vec{v}}{|\vec{l}+\vec{v}|}$ are perfectly aligned, and according to the Fresnel equations the amount of power it reflects is determined by the angle θ_d made by the normal and the incident light. Hence, Equation 2 can be rewritten as:

$$I(\vec{v}, \vec{l}) = G(\vec{l}, \vec{v})D(\vec{h})F(\theta_d). \tag{3}$$

Essentially, the microfacet model is to built upon the construction of a Gauss map that parameterizes the microfacet in Euclidean Space $\mathbf{R}^3$ with its normal $\vec{h}$, where the NDF evaluates its rate of the change over a unit sphere $\mathbf{S}^2$. In this regard, NDF is inherently the gaussian curvature of the surface that the Gauss map applies to. For example, a planar surface with zero Gaussian curvature leads to a Dirac delta NDF that only spikes along the normal of the plane. So, with identical setting given above, the NDF can also be implicitly defined as the inverse of the Gaussian curvature of the illusory surface covered by the microfacets. Moreover, recent study [40] demonstrates that the ellipsoidal microfacet arrangement and the general CGX NDF are equivalent. Whereas the success of the latter has been widely acknowledged in rendering, in the following we present the appealing algebraic properties the former manifests for shape analysis.

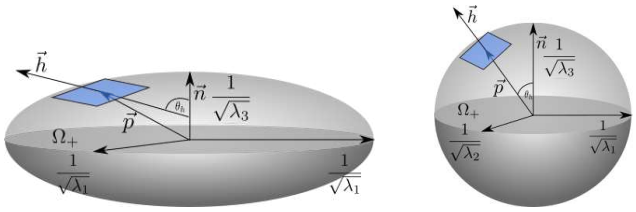

(a) specular surface (b) diffusive surface

Figure 3: The ellipsoid NDF describes that the microfacets can be re-arranged through translation to cover the upper surface of an ellipsoid. A "flatter" ellipsoid indicates that more microfacets are aligned with the surface, representing a smoother material.

3.1. Ellipsoid NDF for Isotropic Reflectance

Implicitly defining the NDF over an ellipsoid offers several algebraically appealing properties. As illustrated in Figure 3, if Ω_+ denotes an arbitrarily defined unit area of the physical surface under examination, and the microfacets can be geometrically translated to cover the upper half of a ellipsoid, then there exists a unique parametrization of surface point $\vec{p}$ by the surface normal $\frac{S\vec{p}}{|S\vec{p}|} = \vec{h}$, and the following is always satisfied:

$$\vec{p}^\intercal S\vec{p} = 1, \tag{4}$$

where S is a 3-by-3 matrix and can always be re-scaled to normalize the RHS of the equation to 1. It has the following properties to characterize the shape of the ellipsoid:

1. S is symmetric and positively definite;

2. In the case of isotropic reflectance, S denotes an ellipsoid of revolution, so its eigenvalues satisfy that $\lambda_3 \geq \lambda_2 = \lambda_1 > 0$;

3. Correspondingly, the lengths of the major and the minor axes are $\frac{1}{\sqrt{\lambda_1}}$ and $\frac{1}{\sqrt{\lambda_3}}$, respectively;

4. The arrangement of the microfacetets has to be physically consistent with the surface geometry, so the minor axis is aligned with the surface normal. Namely, $S\vec{n} = \lambda_3\vec{n}$.

Spectral theorem states that $S = \lambda_1\vec{u}\vec{u}^\intercal + \lambda_2\vec{v}\vec{v}^\intercal + \lambda_3\vec{n}\vec{n}^\intercal$, where $\vec{u}$, $\vec{v}$, and $\vec{n}$ are its eigenvectors. Correspondingly, $S^{-1} = \frac{1}{\lambda_1}\vec{u}\vec{u}^\intercal + \frac{1}{\lambda_2}\vec{v}\vec{v}^\intercal + \frac{1}{\lambda_3}\vec{n}\vec{n}^\intercal$. Also, let $|S|$ denote the determinant of S, and $K_g = |S|(\vec{h}^\intercal S^{-1}\vec{h})^2$ denote the Gaussian curvature of the microfacet-parameterized ellipsoidal surface [17], the ellipsoid NDF $D(\vec{h})$ for isotropic reflectance ($\lambda_1 = \lambda_2$) thus can be expressed as:

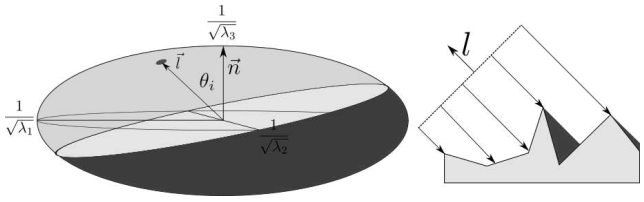

Figure 4: The shadowing function guarantees that the total area receiving illumination over a surface of unit area does not exceed $\vec{l}^\intercal\vec{n}$. In our model, we remove the restriction that the region has to be in the upper sphere Ω_+, so the entire intersected area is considered.

$$D(\vec{h}) = \frac{1}{K_g} = \frac{1}{\lambda_1^2\lambda_3(\frac{(\vec{h}^\intercal\vec{u})^2+(\vec{h}^\intercal\vec{v})^2}{\lambda_1} + \frac{(\vec{h}^\intercal\vec{n})^2}{\lambda_3})^2} = \frac{1}{\lambda_3(1-(\vec{h}^\intercal\vec{n})^2 + \frac{\lambda_1}{\lambda_3}(\vec{h}^\intercal\vec{n})^2)^2}. \quad (5)$$

3.2. The Masking-Shadowing Function

The masking-shadowing function $G(\vec{l},\vec{v})$ is introduced to impose a physical constraint that the visible and the illuminated area must not exceed the projected area along the perceived direction $\vec{v}$ or the illumination direction $\vec{l}$ (Figure 4), respectively. In the case where $\vec{v}$ is fixed, $G(\vec{l},\vec{v}) = G(\theta_i)$ has to satisfy the following for isotropic reflectance [18]:

$$\vec{n}^\intercal\vec{l} = \int_{\Omega_+} \max(\vec{h}\cdot\vec{l},0)D(\vec{h})G(\theta_i)\,d\vec{h}, \quad (6)$$

where $\max(\vec{h}\cdot\vec{l},0)$ is to ensure the the microfacet lies in the shadow vanishes. Because in a typical photometric stereo setup a large population of the microfacets are illuminated by a moving directional light over the upper hemisphere, this highly nonlinear term is only significant when a light significantly deviates away from the normal. So, with the premise that lights are distributed sufficiently, we relax this expression by removing this operator. Therefore, by plugging the widely adopted Smith Microsurface Profile [35] and the ellipsoid NDF with $\int(\vec{h}\cdot\vec{l})D(\vec{h})\,d\vec{h} = \pi\sqrt{\vec{l}^\intercal S\vec{l}|S|^{-1}}$ [39] into Equation 6, we arrive at the following derivation for the shadowing function:

$$G(\vec{l}) = \frac{\vec{l}^\intercal\vec{n}}{\pi}\frac{\lambda_1\sqrt{\lambda_3}}{\sqrt{\vec{l}^\intercal S\vec{l}}} = \frac{\vec{l}^\intercal\vec{n}}{\pi}\frac{\lambda_1\sqrt{\lambda_3}}{\sqrt{\lambda_1(1-(\vec{l}^\intercal\vec{n})^2)+\lambda_3(\vec{l}^\intercal\vec{n})^2}}. \quad (7)$$

3.3. The Fresnel Term

In theory, the Fresnel term $F(\theta_d)$ only starts to vary dramatically as $\theta_d \to \frac{\pi}{2}$, so in general it does not encode sufficient information for shape analysis unless both view and light are at the grazing angles, which only occasionally occurs when lights are located at numerous locations for photometric stereo. Therefore, this term can be safely taken as an unknown constant.

3.4. A General Reflectance Model

Combining Equation 3, 5 and 7 and letting $\lambda = \frac{\lambda_1}{\lambda_3}$ leads to:

$$I(\vec{l}) = C\frac{\lambda}{(1-(1-\lambda)(\vec{h}^\intercal\vec{n})^2)^2}\frac{\vec{l}^\intercal\vec{n}}{\sqrt{(\lambda+(\lambda-1)(\vec{l}^\intercal\vec{n})^2)}}, \quad (8)$$

where C is an unknown product subsuming the camera gain, the Fresnel term and $\frac{1}{\pi}$; more importantly, we notice that λ is a factor independent of the geometry and illumination, as well as a sole term characterizing the material's reflectance property. Being the square of the ratio of the minor axis length $\frac{1}{\sqrt{\lambda_3}}$ to the major axis length $\frac{1}{\sqrt{\lambda_1}}$, λ successfully decouples the pixel-wise material evaluation from the actual imaging process: since $\frac{1}{\sqrt{\lambda_1}}$ is the radius of a circular patch orthographically imaged to a specific pixel, λ is essentially the "normalized" shape descriptor of the ellipsoid, regardless how large the "volume" that the ellipsoid occupies, which is also independent of the camera pose and light intensity.

Algebraically, λ plays a central role in identifying the type of surface reflectance. Since microfacet arrangement has to be consistent with the surface geometry, we have $\frac{1}{\lambda_3} \le \frac{1}{\lambda_1}$, hence $\lambda \in (0,1]$. When $\lambda \to 1$, it results in a sphere, $I(\vec{l}) \to C\sqrt{\lambda}\vec{l}^\intercal\vec{n}$, which corresponds to the ideal diffusive case, because microfacets are arranged along an arbitrary direction with equal probability (Figure 3b). On contrary, when $\lambda \to 0$, $G(\vec{l}) \to 1$, the material reflectance becomes more conspicuous till only specularities are present. Since the former has been extensively studied in the existing literature, here we look into the details of the latter.

4. Physics Driven Approximate Model for Specular Reflectance

$\lambda \to 0$ leads to a description for perfect mirror reflection, and when $G(\vec{l}) \to 1$, the surface radiance has to be evaluated under two scenarios:

1. $\vec{h}^\intercal\vec{n} = 1$. By Equation 2, $I(\vec{l}) \to \int_{\Omega_+} D(\vec{h})\,d\vec{h} \to C\int_{\Omega_+}\delta_{\vec{n}}(\vec{h})\,d\vec{h} = C$, where $\delta(\vec{h})$ is the dirac delta function describing the infinite impulse due to $\frac{1}{\lambda}$;
2. $\vec{h}^\intercal\vec{n} \neq 1$. $I(\vec{l}) = \frac{C\lambda}{(1-(\vec{h}^\intercal\vec{n})^2)^2} \to 0$, which is a direct simplification from Equation 8.

In the first case the light is directly observed through the ideal mirror reflection whereas no diffusive radiance can be captured when $\vec{h}$ falls off from $\vec{n}$ in the second case.

Therefore, by letting λ take a sufficiently small value, we obtain a reflectance function for highly specular materials:

$$I(\vec{l}) \approx \frac{C\lambda}{(1-(1-\lambda)(\vec{h}^\intercal\vec{n})^2)^2}, \tag{9}$$

which can be rearranged into

$$\sqrt{\frac{I(\vec{l})}{C\lambda}}(1-(\vec{h}^\intercal\vec{n})^2+\lambda(\vec{h}^\intercal\vec{n})^2) \approx 1, \tag{10}$$

and by defining $\hat{S} = \vec{u}\vec{u}^\intercal + \vec{v}\vec{v}^\intercal + \lambda\vec{n}\vec{n}^\intercal$, we can further simplify Equation 10 as

$$((\frac{I(\vec{l})}{C\lambda})^{\frac{1}{4}}\vec{h})^T\hat{S}((\frac{I(\vec{l})}{C\lambda})^{\frac{1}{4}}\vec{h}) \approx 1. \tag{11}$$

Equation 11 is essentially the standard equation for an ellipsoid of revolution $\hat{S}$ centered at the origin, where the original ellipsoid S and $\hat{S}$ are co-axial. Fitting an ellipsoid requires at least 4 points on its surface, which can be easily satisfied for photometric stereo. After all, as the directional light relocates, distinct appearances can be obtained except for nearly perfect mirror reflection, for which only the appearance of the illuminant is directly seen from one specific location. In our analysis we rule out this extreme case.

Moreover, we expect our solver in the following section to become less accurate when applied to diffusive materials. As discussed in Section 3.4, in diffusive cases $\lambda \to 1$, the ellipsoid degenerates to a sphere without elongation (see Figure 3b). Algebraically, this means that Equation 11 can be satisfied by a set of non-unique $\hat{S}$. Fortunately, the value of λ itself serves as a good measure for estimation confidence, so we can safely "roll back" to the existing solvers implemented for low-frequency reflectance if large values for λ are detected.

5. Optimal Ellipsoid of Revolution Fitting for Normal Estimation

Given the illumination direction $\vec{l}$ (equivalently the half vector $\vec{h}$) and the image radiance $I(\vec{l})$, photometric stereo seeks to recover the surface normal $\vec{n}$. In our scenario, this boils down to fitting an unknown surface of an ellipsoid of revolution in $\mathbf{R}^3$ using points $\vec{h}$ on a spherical surface $\mathbf{S}^2$ whose lengths are re-scaled by $\left(\frac{I(\vec{l})}{C\lambda}\right)^{\frac{1}{4}}$. After the ellipsoid is determined, by detecting its elongation the surface normal can also be obtained.

Fitting an ellipsoid can be formulated into an energy minimization problem, and we are able to retrieve the global minimum from the solutions to a system of polynomials. For simplification we denote $P = \sqrt{I(\vec{l})}$ and $\omega = \sqrt{\frac{1}{C\lambda}}$. To get around the unit norm constraint on $\vec{n}$, we also let $\vec{\hat{n}} = \sqrt{(1-\lambda)\omega}\vec{n}$, hence for each of the k observations, Equation 10 can be rewritten as

$$P_i\left(\omega - \vec{\hat{n}}^T\vec{h_i}\vec{h_i}^\intercal\vec{\hat{n}}\right) = 1, i = 1, 2, \cdots, k, \tag{12}$$

where i is the observation index.

By averaging all k equations, we have $\bar{P} = \frac{\sum_{i=1}^k P}{k}$, $\bar{H} = \frac{\sum_{i=1}^k P_i\vec{h_i}\vec{h_i}^\intercal}{k}$, and

$$\omega = \frac{1+\vec{\hat{n}}^T\bar{H}\vec{\hat{n}}}{\bar{P}}. \tag{13}$$

Moreover, combining Equation 13 into Equation 12 leads to

$$\vec{\hat{n}}^T\left(P_i\vec{h_i}\vec{h_i}^\intercal - P_i\frac{\bar{H}}{\bar{P}}\right)\vec{\hat{n}} = \frac{P_i}{\bar{P}} - 1, \tag{14}$$

a quadratic polynomial with respect to $\hat{n} = [\hat{n_1}, \hat{n_2}, \hat{n_3}]^T$. Therefore, all k equations in Equation 14 can be organized into the matrix form

$$Mx = M[\hat{n}_1^2, \hat{n}_1\hat{n}_2, \hat{n}_1\hat{n}_3, \hat{n}_2^2, \hat{n}_2\hat{n}_3, \hat{n}_3^2]^T = b, \tag{15}$$

so $M \in \mathbb{R}^{k\times 6}$ and $b \in \mathbb{R}^k$ are established.

However, due to measurement noise and model approximation, we do not enforce the equality in Equation 15 to hold strictly, instead we try to find the optimal $\hat{n} = [\hat{n_1}, \hat{n_2}, \hat{n_3}]^T$ that minimizes the following energy function

$$f(\hat{n}) = \|Mx - b\|_2^2 = x^TM^TMx - 2b^TMx + b^Tb. \tag{16}$$

Since the cost function in Equation 16 is nonconvex, we try to find its global minimizer by retrieving all its stationary points. Specifically, we solve the three-variable cubic equations defined by the partial derivatives as

$$\frac{\partial f}{\partial \hat{n}_1} = 0, \frac{\partial f}{\partial \hat{n}_2} = 0, \frac{\partial f}{\partial \hat{n}_3} = 0, \tag{17}$$

which is a three-variable cubic polynomial system that has 27 solutions. Since the system is homogeneous, the solutions are positive-negative symmetric. Therefore, we only need to examine 13 independent solutions. These facts motivate us to develop a solver based on the symmetric Gröbner basis [24]. To our best knowledge, this is the first example for such technique to apply to photometric stereo. Finally, λ and C can be determined consecutively using the length of $\vec{\hat{n}}$.

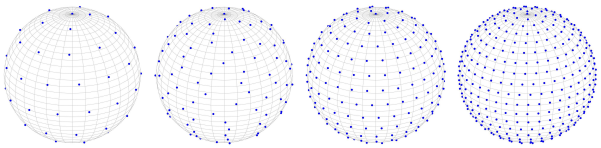

(a) 60 lights (b) 150 lights (c) 250 lights (d) 500 lights

Figure 5: Distribution of lights with various densities. As illustrated, the lights from the behind make no effect on image formation.

6. Experiments

We validate our model and evaluate our method on both synthetic and real images. Since shadowing and interreflections are not considered, we use PBRT [29] and MERL BRDF [25] to render spheres as our synthetic input. In the experiments, we compare our results with the results produced by the Least Square (*LS*), Constrained Bivariate Regression (*CBR*) [21] and Biquadratic (*Bi-Quad*) [33], where the latter two offer the state-of-the-art performance according to the benchmark test in [34]. We locate the direction of the lights using a set of spiral points [31]. We also apply our method on the helmet model from the USC "Light Stage Data Gallery" [13] to evaluate its stability on specular materials, as well as on the "DiLiGenT" benchmark data set [34] to examine its performance in diffusive cases.

6.1. Evaluation with Synthetic Images

We perform two types of experiments on synthetic images. First, we compare the angular estimation accuracy obtained by the four methods under sparse (*i.e.* 60 lights) and dense (*i.e.* 500 lights) light distributions, respectively, then we investigate how the light density interacts with our solver by applying it to 4 different lighting configurations (*i.e.* 60, 150, 250, 500 lights, Figure 5). For CBR, we set $N_1 = 2$ and $N_2 = 4$, and "retroreflective" on. Unless otherwise stated, in our setup all methods take the original input as is, and only positive pixel values are considered. This means that only a subset of the lights contribute the actual computation for a specific pixel.

6.1.1 Complementary Accuracy Traces

With respect to materials, our solution produces an accuracy trace complementary to those produced by the other three solutions targeting on low-frequency reflectance. Figure 6a and 6b present the angular estimation errors in degrees over the 100 materials in MERL, under both dense light distribution (500 lights) and sparse light distribution (60 lights). From the plots we draw two observations: (1) All methods perform inferiorly over a specific set of materials, and this division performance over the materials is almost independent of light density; (2) The peer methods outperform on surface with diffusive reflectance, and our proposed solver delivers better performance for specular surfaces with a few exceptions (Section 6.1.2), which is consistent with our prediction made in Section 3.4.

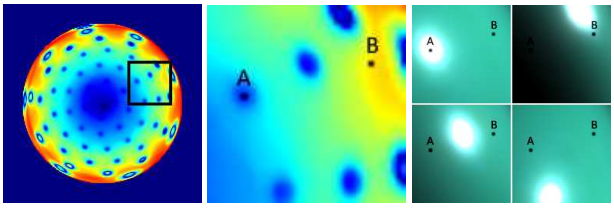

Figure 7: Some materials exhibit highly localized specularity, so its appearance is sensitive to light density. In terms of model fitting, the pixels that carry specular signals (*e.g.* point A) are more likely to be correctly estimated by our solver than those do not (*e.g.* point B). Left: Estimation error for specular green phenolic. Middle: A closer view over the region showing both accurate and inaccurate estimations. Right: Appearances produced by four distinct lights. Point B is "by-passed" by all lights so it does not carry specular signals.

6.1.2 Impact of Light Density

We also study how the light density affects the accuracy of our solver. Figure 6c compares the accuracy obtained under various illumination densities: 500, 250, 150, 60 lights, respectively. The main observation we make is that our solution produces stable output for various lighting densities, but the denser the distribution, the higher the accuracy is achieved. More importantly, the gain brought by applying denser light distribution is more significant for some particular specular surfaces, which are highlighted in red. For example, "specular green phenolic", as indicated in Figure 7, exhibits extremely localized specularity with large estimation error. This is because when light distribution is sufficiently dense, many pixels (*e.g.* point B) do not exhibit specular property at all, as compared with the pixels (*e.g.* point A) covered by specularities, and they are less accurately estimated as a consequence of mis-fitting the model. Though uncommon in a photometric stereo setup, we expect to see that with a denser light distribution, estimation on the specular materials highlighted in red in Figure 6 shall continue to improve.

6.2. Evaluation with Real Images

We apply our solution to two image sets, "helmet side right" from "Light Stage Data Gallery" and "DiLiGenT". The "helmet" image set contains specular appearances captured under 253 directional lights, and the "DiLiGenT" data

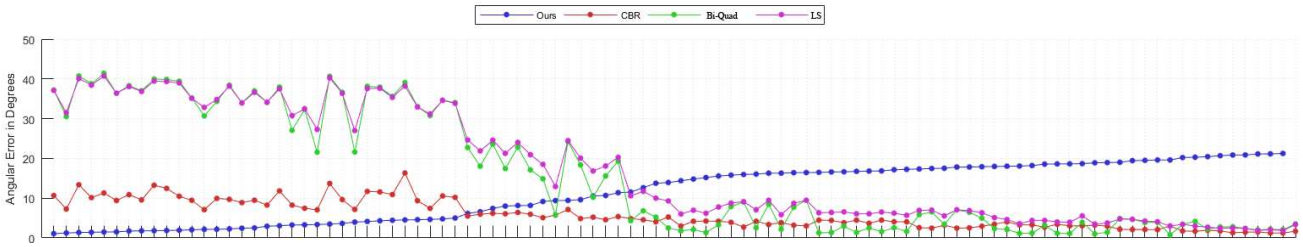

(a) Mean estimation error in degrees produced by the four methods on sampled MERL reflectance using 60 lights.

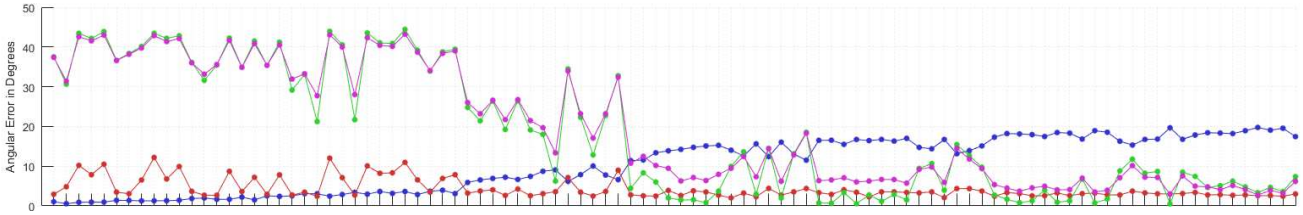

(b) Mean estimation error in degrees produced by the four methods on sampled MERL reflectance using 500 lights.

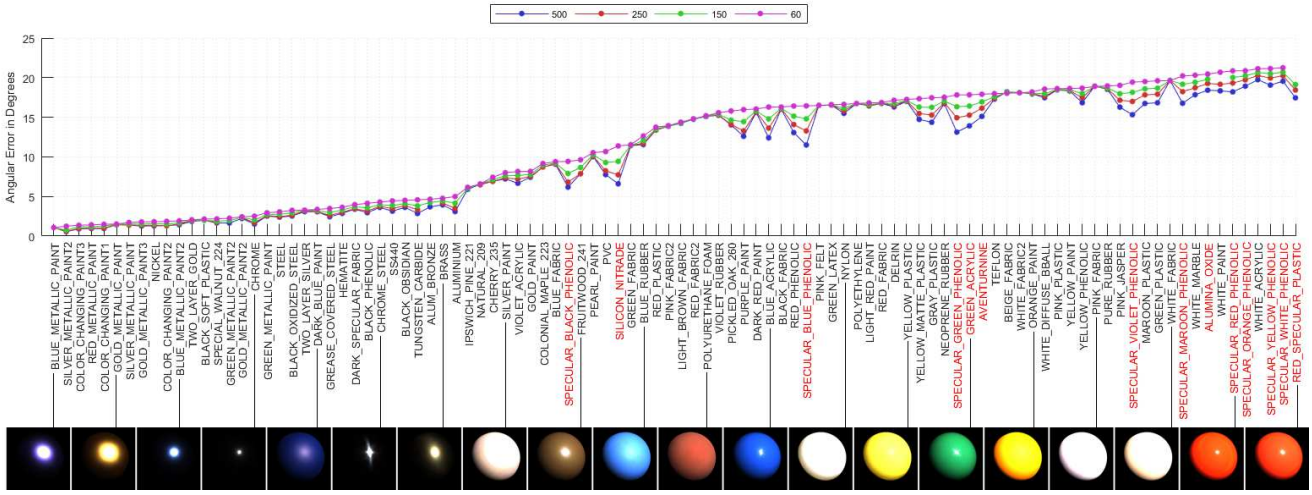

(c) Mean estimation error in degrees produced by the our method on sampled MERL reflectance using 60, 150, 250 and 500 lights.

Figure 6: Performance evaluation on synthetic images rendered by PBRT using MERL. Materials that exhibit extremely localized specularities are highlighted in red, the estimate accuracy of which rely on the light density for our method.

set mainly represents diffusive materials.

6.2.1 Performance on Specular Surfaces

Figure 8 visualizes the estimated normals with their respective $+x$, $+y$ and $+z$ components, with (1) upper sphere lights only and (2) all 253 lights. We do not make qualitative comparison because ground truth is unavailable, but qualitatively our method shows reasonable and consistent results, indicating its stability. Besides, the convex shape of the model is clearly illustrated by $+x$, $+y$ and $+z$ components together. We also note that *CBR* delivers much more reasonable results when only upper sphere lights are selected (Figure 9). In general, light distribution makes a major impact on estimation accuracy. It is reported that existing approaches shall perform better with properly adjusted "position threshold" [34], so we will look into the details of this factor in our future work.

6.2.2 Detection of Diffusive Reflectance

Among the ten models in "DiLiGenT", "ball", "reading", "cow" and "harvest" represent relatively more specular materials. Figure 10 compares the estimation error produced by *CBR*, *Bi-Quad* and ours, together with the median value of λ we obtained for each model. Here we assume that each object is made of homogenous material, so a rough cross-pixel analysis is allowed. It is interesting to observe that except for "ball" and "reading", the lower the λ value we detect, the better the performance our method delivers. The

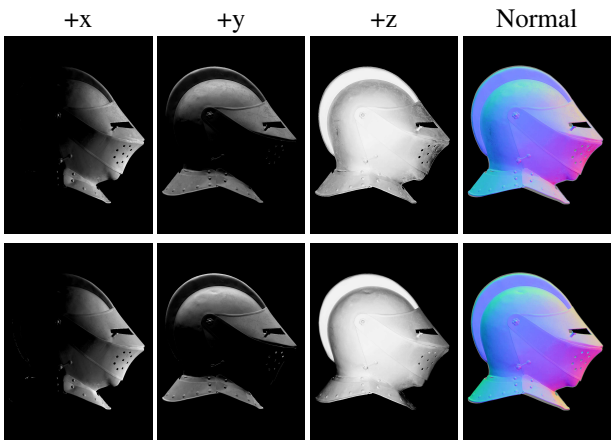

Figure 8: Normal maps obtained by our method for "helmet front left". Row 1: with upper sphere lights only. Row 2: with all 253 lights.

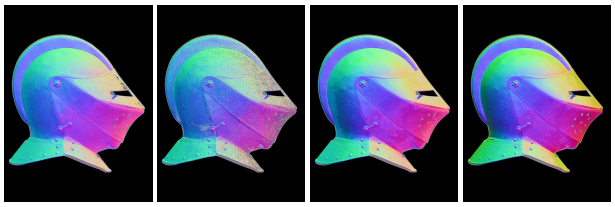

Figure 9: Normal maps obtained by *CBR* and *Bi-Quad*. From left to right: *CBR* with upper sphere lights only and with all 253 lights; *Bi-Quad* with upper sphere lights only and with all 253 lights.

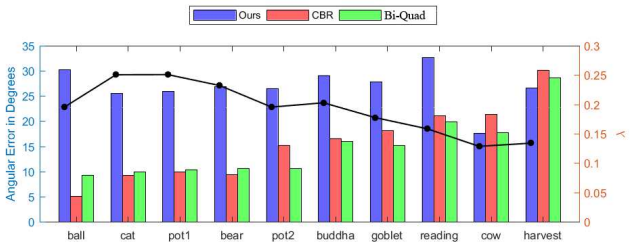

Figure 10: Average estimation error in degrees produced by the three methods. The black dotted line indicates the corresponding λ obtained for each model. Our model predicts that the trace of λ should be consistent with the trace of estimations error for homogeneous materials.

inferior performance on "ball" is due to limited light distribution as discussed in Section 6c, and "reading" is inaccurately estimated because surface non-convexity has caused a significant amount of specular inter-reflections.

However, we would also like to point out that though λ can correctly indicate the smoothness of the surface, a fail-safe "switch" that allows us to roll back to the existing solutions for low-frequency reflectance remains to be absent. So, how to properly incorporate λ into a solution that handles general reflectance remains as part of our future work.

7. Conclusion

In this paper, we derive a novel analytical microfacet-based isotropic reflectance model from the ellipsoid normal distribution function. We also introduce a physically interpretable approximate of our model that is particularly serviceable for specular reflectance analysis. With this approximate, we identify that the problem for specular surface normal recovery is essentially an ellipsoid of revolution fitting problem, where the latter can be described by a system of polynomials. And in order to solve this problem, we also devise a fast, non-iterative and globally optimal solver. Experiments on both synthetic and real images demonstrates the superiority of our model and algorithm on the targeted specular surfaces.

Currently, our approximate model and its polynomial system solver are tailored for highly specular surfaces. To extend the success into shape recovery of general isotropic surfaces is left as our future work.

Acknowledgement

This work was supported in part by JSPS KAKENHI Grant Number JP15H05918 and a project commissioned by the New Energy and Industrial Technology Development Organization (NEDO).

References

[1] N. Alldrin, T. Zickler, and D. Kriegman. Photometric stereo with non-parametric and spatially-varying reflectance. In *CVPR*, pages 1–8, 2008.

[2] N. G. Alldrin and D. J. Kriegman. Toward reconstructing surfaces with arbitrary isotropic reflectance: A stratified photometric stereo approach. In *ICCV*, pages 1–8, 2007.

[3] M. Ashikhmin and S. Premoze. Distribution-based brdfs. *Unpublished Technical Report, University of Utah*, 2:6, 2007.

[4] M. Ashikhmin and P. Shirley. An anisotropic phong brdf model. *Journal of graphics tools*, 5(2):25–32, 2000.

[5] M. Ashikmin, S. Premože, and P. Shirley. A microfacet-based brdf generator. In *Proceedings of the 27th annual conference on computer graphics and interactive techniques*, pages 65–74, 2000.

[6] M. M. Bagher, J. Snyder, and D. Nowrouzezahrai. A non-parametric factor microfacet model for isotropic brdfs. *ACM Transactions on Graphics*, 35(5):159, 2016.

[7] J. F. Blinn. Models of light reflection for computer synthesized pictures. In *ACM SIGGRAPH Computer Graphics*, volume 11, pages 192–198, 1977.

[8] B. Burley and W. D. A. Studios. Physically-based shading at disney. In *ACM SIGGRAPH*, pages 1–7, 2012.

[9] H.-S. Chung and J. Jia. Efficient photometric stereo on glossy surfaces with wide specular lobes. In *CVPR*, pages 1–8, 2008.

[10] R. L. Cook and K. E. Torrance. A reflectance model for computer graphics. *ACM Trans. Gr.*, 1(1):7–24, 1982.

[11] Y. Dong, J. Wang, X. Tong, J. Snyder, Y. Lan, M. Ben-Ezra, and B. Guo. Manifold bootstrapping for svbrdf capture. In *ACM Transactions on Graphics*, volume 29, page 98, 2010.

[12] Z. Dong, B. Walter, S. Marschner, and D. P. Greenberg. Predicting appearance from measured microgeometry of metal surfaces. *ACM Transactions on Graphics*, 35(1):9, 2015.

[13] P. Einarsson, C.-F. Chabert, A. Jones, W.-C. Ma, B. Lamond, T. Hawkins, M. T. Bolas, S. Sylwan, and P. E. Debevec. Relighting human locomotion with flowed reflectance fields. *Rendering techniques*, 2006:17th, 2006.

[14] A. Gardner, C. Tchou, T. Hawkins, and P. Debevec. Linear light source reflectometry. In *ACM Transactions on Graphics*, volume 22, pages 749–758, 2003.

[15] A. S. Georghiades. Incorporating the torrance and sparrow model of reflectance in uncalibrated photometric stereo. In *CVPR*, pages 816–823, 2003.

[16] D. B. Goldman, B. Curless, A. Hertzmann, and S. M. Seitz. Shape and spatially-varying brdfs from photometric stereo. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 32(6):1060–1071, 2010.

[17] R. Goldman. Curvature formulas for implicit curves and surfaces. *Computer Aided Geometric Design*, 22(7):632–658, 2005.

[18] E. Heitz. Understanding the masking-shadowing function in microfacet-based brdfs. *Journal of Computer Graphics Techniques*, 3(2):32–91, 2014.

[19] T. Higo, Y. Matsushita, and K. Ikeuchi. Consensus photometric stereo. In *CVPR*, 2010.

[20] Z. Hui and A. C. Sankaranarayanan. A dictionary-based approach for estimating shape and spatially-varying reflectance. In *ICCP*, pages 1–9, 2015.

[21] S. Ikehata and K. Aizawa. Photometric stereo using constrained bivariate regression for general isotropic surfaces. In *CVPR*, pages 2179–2186, 2014.

[22] S. Ikehata, D. Wipf, Y. Matsushita, and K. Aizawa. Photometric stereo using sparse bayesian regression for general diffuse surfaces. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 36(9):1816–1831, 2014.

[23] E. P. Lafortune, S.-C. Foo, K. E. Torrance, and D. P. Greenberg. Non-linear approximation of reflectance functions. In *Proceedings of the 24th annual conference on Computer graphics and interactive techniques*, pages 117–126, 1997.

[24] V. Larsson and K. Åström. Uncovering symmetries in polynomial systems. In *ECCV*, pages 252–267. Springer, 2016.

[25] W. Matusik and M. Brand. A data-driven reflectance model. *ACM Transactions on Graphics*, 22(3):759–769, 2003.

[26] G. Nam, J. H. Lee, H. Wu, D. Gutierrez, and M. H. Kim. Simultaneous acquisition of microscale reflectance and normals. *ACM Transactions on Graphics*, 35(6), 2016.

[27] A. Ngan, F. Durand, and W. Matusik. Experimental analysis of brdf models. *Rendering Techniques*, 2005(16th):2, 2005.

[28] M. Oren and S. K. Nayar. Generalization of lambert's reflectance model. In *Proceedings of the 21st annual conference on Computer graphics and interactive techniques*, pages 239–246, 1994.

[29] M. Pharr, W. Jakob, and G. Humphreys. *Physically based rendering: From theory to implementation*. Morgan Kaufmann, 2016.

[30] B. T. Phong. Illumination for computer generated pictures. *Communications of the ACM*, 18(6):311–317, 1975.

[31] E. A. Rakhmanov, E. Saff, and Y. Zhou. Minimal discrete energy on the sphere. *Math. Res. Lett*, 1(6):647–662, 1994.

[32] B. Shi, P. Tan, Y. Matsushita, and K. Ikeuchi. Elevation angle from reflectance monotonicity: Photometric stereo for general isotropic reflectances. In *ECCV*, pages 455–468, 2012.

[33] B. Shi, P. Tan, Y. Matsushita, and K. Ikeuchi. Bi-polynomial modeling of low-frequency reflectances. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 36(6):1078–1091, 2014.

[34] B. Shi, Z. Wu, Z. Mo, D. Duan, S.-K. Y. Yeung, and P. Tan. A benchmark dataset and evaluation for non-lambertian and uncalibrated photometric stereo. In *CVPR*, 2016.

[35] B. Smith. Geometrical shadowing of a random rough surface. *IEEE Transactions on Antennas and Propagation*, 15(5):668–671, 1967.

[36] K. E. Torrance and E. M. Sparrow. Theory for off-specular reflection from roughened surfaces. *Journal of the Optical Society of America*, 57(9):1105–1112, 1967.

[37] S. Tozza, R. Mecca, M. Duocastella, and A. D. Bue. Direct differential photometric stereo shape recovery of diffuse and specular surfaces. *Journal of Mathematical Imaging and Vision*, 56:57–76, 2016.

[38] T. S. Trowbridge and K. P. Reitz. Average irregularity representation of a rough surface for ray reflection. *Journal of the Optical Society of America*, 65(5):531–536, 1975.

[39] G. Vickers. The projected areas of ellipsoids and cylinders. *Powder technology*, 86(2):195–200, 1996.

[40] B. Walter, Z. Dong, S. Marschner, and D. P. Greenberg. The ellipsoid normal distribution function. *Supplementary material*, http://www.cs.cornell.edu/Projects/metalappearance/, 2016.

[41] B. Walter, S. R. Marschner, H. Li, and K. E. Torrance. Microfacet models for refraction through rough surfaces. In *Proceedings of the 18th Eurographics conference on Rendering Techniques*, pages 195–206, 2007.

[42] G. J. Ward. Measuring and modeling anisotropic reflection. *ACM SIGGRAPH Computer Graphics*, 26(2):265–272, 1992.

[43] R. J. Woodham. Photometric method for determining surface orientation from multiple images. *Optical engineering*, 19(1):191139–191139, 1980.

[44] 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 Conference on Computer Vision*, pages 703–717, 2010.

[45] Z. Wu and P. Tan. Calibrating photometric stereo by holistic reflectance symmetry analysis. In *CVPR*, pages 1498–1505, 2013.