

Bi-polynomial Modeling of Low-frequency Reflectances

Boxin Shi, *Student Member, IEEE*, Ping Tan, *Member, IEEE*, Yasuyuki Matsushita, *Senior Member, IEEE*, and Katsushi Ikeuchi, *Fellow, IEEE*

Abstract—We present a bi-polynomial reflectance model that can precisely represent the low-frequency component of reflectance. Most existing reflectance models aim at accurately representing the complete reflectance domain for photo-realistic rendering purposes. In contrast, our bi-polynomial model is developed for the purpose of accurately solving inverse problems by effectively discarding the high-frequency component while retaining nonlinear variations in the low-frequency part. The bi-polynomial reflectance model is useful for estimating reflectance and shape of an object. Experimental evaluation in comparison with other parametric reflectance models demonstrates that the proposed model achieves better performance in reflectometry and photometric stereo applications.

Index Terms—Low-frequency reflectance, Parametric BRDF model, Radiometric image analysis, Reflectometry, Photometric stereo

1 INTRODUCTION

PARAMETRIC modeling of reflectances plays an important role in both rendering and inverse problems in radiometric image analysis. The vast majority of existing parametric reflectance models are developed for the purpose of photo-realistic rendering. They are designed to have an accurate representation of specular components and successfully applied to forward problems in computer graphics. However, these reflectance models are not necessarily suitable for inverse problems in computer vision, such as reflectance and shape estimation. In fact, many of these reflectance models severely complicate the inverse problems by introducing high nonlinearity when they are directly used in the computation, and as a result, the solution methods are forced to involve unstable and expensive nonlinear (or even non-convex) optimization procedures. While one could use a simplistic model to avoid such a problem, *e.g.*, the Lambert's reflectance model, the accuracy of estimates suffers from its discrepancy from the real-world reflectance. Therefore, it is desired to develop a reflectance model that well represents real-world reflectances while retaining simplicity for inverse problems.

In forward rendering problems, one of the key challenges is to accurately model specular reflection that exhibits high-frequency reflectance variations in the incident or exitant angular domain. Since the specular component significantly varies across materials, in order to faithfully represent it, the specular term of a reflectance model tends to become complex and highly nonlinear. While it is essential for photo-realistic rendering in computer graphics, explicit modeling of high-

frequency specular reflections is seldom necessary for inverse problems in computer vision, particularly when sparse lighting (such as a directional light) is used. In fact, in an image, most of the pixels exhibit low-frequency reflections (close to diffuse reflectances) for most materials under sparse lighting. For example, high-frequency (strong specular) reflections are only observed in a sparse manner at points where the surface normal is close to the bisector of viewing and lighting directions. Recent robust photometric stereo methods [1], [2] are built upon similar observation, where high-frequency reflectances are treated as outliers.

Motivated by this observation, we develop a compact parametric Bidirectional Reflectance Distribution Function (BRDF) model for radiometric image analysis using a *bi-polynomial* representation. We design this model with two goals by restricting it to isotropic BRDFs. First, it should be able to faithfully represent low-frequency reflectances of a broad class of materials. Second, it should make the solution of inverse problems tractable. Our model is built upon a factorized form of bivariate BRDF models for isotropic materials [3], [4], where the BRDF is represented as a product of two univariate functions of half and difference angles. We approximate these univariate functions by low-order polynomials. In the preliminary version of this work [5], we employed quadratic functions for these univariate functions. In this work, we extend it to a general bi-polynomial model and perform detailed analysis and validations across varying polynomial orders. We further perform comprehensive comparisons with various BRDF models by applying this model to reflectometry and photometric stereo. We show that accurate results can be obtained by analyzing the low-frequency reflectances with our proposed model.

The rest of this paper is organized as follows: In the next section, we discuss previous works in related areas. In Section 3, we define low-frequency reflectance and show its approximation by using low-intensity observations. In Section 4, we introduce the proposed bi-polynomial model. We then assess our model by fitting measured BRDFs and compare

- B. Shi and K. Ikeuchi are with Institute of Industrial Science, the University of Tokyo, Japan.
Email: {shi, ki}@cvl.iis.u-tokyo.ac.jp
- P. Tan is with Department of Electrical and Computer Engineering, National University of Singapore.
Email: eletp@nus.edu.sg
- Y. Matsushita is with Microsoft Research Asia, Beijing, China.
Email: yasumat@microsoft.com

with existing parametric models. In Section 5 and Section 6, we show applications of our model to reflectometry and photometric stereo problems. Section 7 concludes the paper.

2 RELATED WORKS

Reflectance Modeling. To precisely represent the appearance of real-world materials, various parametric BRDF models have been developed over the decades. These BRDF models can be categorized into physically-based and empirical models. Physically-based models, such as the Torrance-Sparrow [6] and Cook-Torrance [7] models, are mostly designed based on the microfacet theory. They assume that surfaces consist of shiny V -grooves with consideration of geometric attenuation (masking, shadowing, and inter-reflections) and Fresnel effects. Based on a similar microfacet theory, the Oren-Nayar model [8] captures the reflectance of rough surfaces. The empirical models such as the Phong [9] and Blinn-Phong [10] models are widely used because of their computational efficiency. The Lafortune model [11] uses generalized cosine lobes. This model can fit measured reflectance data with high precision, but with unintuitive parameters. There are also models that bridge these two categories by partly using physically motivated terms, *e.g.*, the Ward [12] and Ashikhmin models [13]. The Ward model is also based on the microfacet theory, but omits the Fresnel and geometric attenuation terms. The Ashikhmin model is an anisotropic BRDF model with a non-Lambertian diffuse term. An experimental evaluation of various models with measured data can be found in [14]. Generally, parametric models are compact and easy to use for forward problems, but they are only accurate for a limited class of materials.

A BRDF can also be represented by a 4-D discrete table indexed by lighting and viewing directions. For isotropic materials, this representation can be reduced to a 3-D table [15], [16]. The 3-D table can be re-arranged using half and difference vectors [17] through a re-parameterization as suggested in [3]. It can be further reduced to a 2-D table for a wide range of isotropic materials by omitting the rotation of difference vector [4], [18], [19]. Although high-quality rendering can be achieved using the discrete table representations, such non-parametric forms are generally unsuitable for inverse problems because the number of parameters to be estimated becomes prohibitively large. To maintain the accuracy while reducing the complexity, recent approaches use various basis functions to represent general BRDFs [20], [21], [22]. Our bi-polynomial model shares a similar goal of simplifying the BRDF representation for general materials, but we focus on modeling low-frequency reflectances with a simpler parametric form. Furthermore, we aim to solve inverse problems for radiometric image analysis rather than forward rendering.

Radiometric Image Analysis. Radiometric image analysis seeks to recover scene properties, such as reflectance and shape, from the recorded scene radiance. Here we briefly review related work in reflectometry and photometric stereo, *i.e.*, surface reflectance and shape estimation, respectively.

Most of the works in reflectometry are based on parametric reflectance models. Yu *et al.* [23] use a sparse set of photos

and assume the Ward reflectance model with spatially varying diffuse reflection and homogeneous specular reflection. Boivin and Gagalowicz [24] use the same BRDF model, but their method deals with only a single image in a hierarchical and iterative framework. Hara *et al.*'s method [25] uses multiple point light sources to estimate both illumination distribution and reflectance represented by the Torrance-Sparrow model. There are recent works of reflectometry that use non-parametric bivariate BRDFs with a discrete table representation. Romeiro *et al.* propose reflectometry with/without measured illumination [4], [19] by assuming that the isotropic BRDFs can be represented using a 2-D table. Their evaluation shows that the 2-D representation is accurate for a majority of the isotropic BRDFs.

In a shape estimation context, early photometric stereo works [26], [27] are based on the Lambert's reflectance model. Although the computation is simple, their performance degrades on many real surfaces, which often exhibit non-Lambertian reflectance. Some methods use four light sources to avoid shadow and specularities [28], [29]. By using more images and recent robust estimation techniques, the outliers that deviate from the Lambertian assumptions can be efficiently detected and discarded through rank minimization [1] and sparse regression [2]. To make use of all the observed data, more sophisticated parametric BRDF models have also been used, such as the methods based on the Torrance-Sparrow model [30], [31], the Ward model [32], [33], and other multi-lobe models [34]. However, all these methods assume the diffuse reflectance to be Lambertian, which is not true for real surfaces.

To deal with more general materials, especially those with non-Lambertian diffuse reflection, some recent methods solve the photometric stereo problem with reflectance symmetries, such as isotropy and/or reciprocity. Alldrin *et al.* [35] exploit isotropy to estimate the azimuth angle of normals. By assuming bivariate BRDFs, they estimate the elevation angle of normals and surface reflectances iteratively [18]. By further assuming that the BRDFs can be projected as a 1-D monotonic function, the elevation angle can also be estimated without using iterative optimization [36]. A theory of reflectance symmetries and its application to photometric stereo is summarized in [37]. Based on those properties, surface reconstruction methods using a special lighting rig and multi-view images are introduced in [38] and [39]. Photometric stereo can also be applied to general diffuse surfaces by considering some consensus properties [40]. Given hundreds of images, surface normal can be estimated by exploring the similarity of radiance profiles [41], [42] and attached shadow codes [43], under unknown illumination and unknown reflectance. Given thousands of images, it is even possible to apply photometric stereo to general anisotropic surfaces [44]. Although these methods can deal with a great range of general materials, they usually require special imaging setup or complicated optimization. In comparison, our bi-polynomial representation is compact yet accurate for many real materials. As a result, data capture and optimization becomes simple with our method.

Given photometric stereo images, biquadratic polynomials are useful for representing the images that may include self-

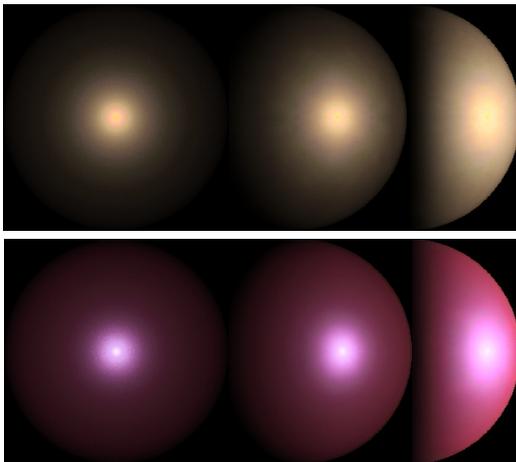


Fig. 1. An illustration of low-frequency reflectance. From left to right, spheres illuminated by distant lights from directions $[0, 0, 1]^T$, $[1/\sqrt{2}, 0, 1/\sqrt{2}]^T$, $[1, 0, 0]^T$ and viewed from $[0, 0, 1]^T$ are shown.

shadowing and interreflections as shown in polynomial texture maps [45], which are able to interpolate images and synthesize appearances under new lighting directions. A more recent approach [46] extends the polynomial texture maps to handle specularity and shadow. Our model uses a similar mathematical form for representing the low-frequency component of general isotropic reflectance.

3 LOW-FREQUENCY REFLECTANCE

Let us consider the reflectance as a function of lighting, viewing, and surface normal directions. We use the term *low-frequency reflectance* to denote the reflectance component that does not abruptly change with the variation of lighting directions. Note that this definition does not limit the low-frequency component to diffuse reflectance, because it can include a wide and blunt specular lobe.

On many real surfaces, strong specularity is only observed when the surface normal is close to the bisector of lighting and viewing directions. Thus, the majority of pixels in an image should present low-frequency reflectances under a directional (or sparse) lighting. We show synthetic spheres in Figure 1 rendered using the measured BRDFs ALUM-BRONZE (top) and VIOLET-ACRYLIC (bottom) from the MERL BRDF database [17]. The spheres are rendered under different lighting directions but a fixed viewpoint. The majority of pixels of these renderings have smoothly varying values¹, which we refer to as low-frequency reflectance observations. For a static scene observed from a fixed camera under a continuously moving light source, the low-frequency reflectances are observed at the pixels where intensities do not show sudden changes under varying lighting directions. We seek to model such low-frequency reflectances in a concise form by effectively discarding the high-frequency reflectances for applications to inverse problems.

1. These high-dynamic range (HDR) images are tone mapped using the method of Reinhard *et al.* [47].

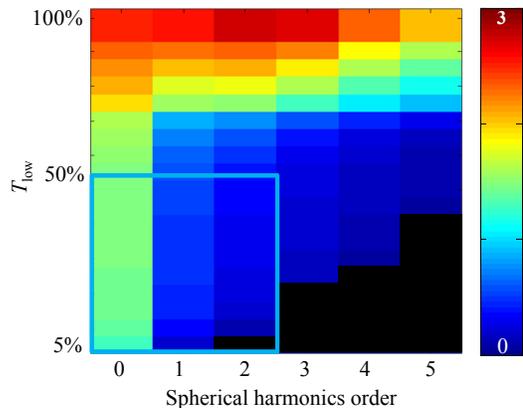


Fig. 2. Fitting errors of spherical harmonics with different orders to observations under various T_{low} . The result is square-rooted for a visualization purpose. The legend shows the mapping between the error magnitude and color. The black areas are undefined due to insufficient equations for fitting.

To capture low-frequency reflectances, we need a method to identify these observations. As we will show below, the low-frequency reflectances have strong correlation with low-intensity observations. Hence, we can simply use an intensity threshold T_{low} to extract observations of low-frequency reflectance in practice. For example, given a set of photometric stereo images, we may draw an intensity profile for each pixel under varying lighting directions. After discarding observations in shadow, we sort all the remaining observations in an ascending order and keep only those ranked below a percentage T_{low} .

To evaluate the effectiveness of this simple method, we fit spherical harmonics of different orders to our low-intensity observations and assess the fitting error. We perform fitting for the intensity profile at each pixel, where the observation is a function of lighting directions. The intensity profile y can be fit by spherical harmonics \tilde{y} represented as

$$\tilde{y}(\theta, \phi) = \sum_{l=0}^b \sum_{m=-l}^l C_{lm} Y_{lm}(\theta, \phi), \quad (1)$$

where (θ, ϕ) are elevation angle and azimuth angle of a lighting direction, b is the order of spherical harmonics, Y is spherical harmonics basis functions, and C is the coefficient. The coefficient C can be solved for by linear least squares. For experimental validation, we fit Equation (1) to the observations thresholded by T_{low} using all of the 100 measured materials in the MERL database [17]. Specifically, we fix the viewing direction as $[0, 0, 1]^T$ and sample 1620 normals by uniformly choosing 36 longitudes and 45 altitudes on the hemisphere. For each of these normals, we use 100 lighting directions randomly sampled from the hemisphere. In other words, we have 1620 intensity profiles, and the length of each is 100. We then vary T_{low} from 10% to 100% with a step of 10%, and vary the order of spherical harmonics b from 0 to 5. For each intensity profile of length f , given fixed b and T_{low} , we compute the relative Root Mean Square Error (RMSE) [4], which is defined

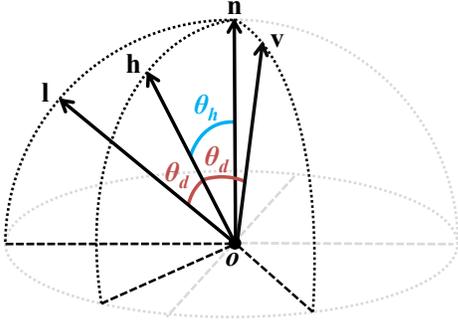


Fig. 3. The definitions of θ_h and θ_d .

as

$$\text{RMSE} = \frac{1}{f} \sqrt{\sum_{i=1}^f \frac{(y_i - \tilde{y}_i)^2}{y_i^2}}, \quad (2)$$

where y and \tilde{y} are original and spherical harmonics fit values. For each fixed b and T_{low} , we perform this evaluation for all 1620 normals and 100 BRDFs, and visualize the mean RMSE in Figure 2. From the error distribution, we can see that low-order spherical harmonics closely fit the observations selected by a small T_{low} below 50%, as indicated by the blue rectangle in the figure. Note that our representation in Equation (1) is different from the spherical harmonics based BRDF representation in [48], which is not intended for inverse problems. We use Equation (1) only to verify the effectiveness of the intensity thresholding for selecting low-frequency reflectances.

4 THE BI-POLYNOMIAL BRDF MODEL

This section describes the proposed bi-polynomial BRDF model. Given surface normal, lighting, and viewing directions \mathbf{n} , \mathbf{l} , and \mathbf{v} , we can calculate the half vector as $\mathbf{h} = (\mathbf{l} + \mathbf{v}) / \|\mathbf{l} + \mathbf{v}\|$, which is the bisector of \mathbf{l} and \mathbf{v} . Following the notations of [3], we use θ_h to denote the half angle between \mathbf{n} and \mathbf{h} , and θ_d for the difference angle between \mathbf{l} (or \mathbf{v}) and \mathbf{h} as illustrated in Figure 3. Hence, the following relationships hold: $\mathbf{n}^\top \mathbf{h} = \cos \theta_h$ and $\mathbf{l}^\top \mathbf{h} = \cos \theta_d$.

Our reflectance model is built upon the bivariate BRDF model of [3], where it is shown that most of the isotropic BRDFs can be represented as a bivariate function $\rho(\theta_h, \theta_d)$. This representation is evaluated by [4] with a large number of measured BRDFs [17] for the development of passive reflectometry. It is further discussed in [3] that any isotropic BRDF based on the microfacet theory should consist of a univariate function of θ_h , and its Fresnel term should be a univariate function of θ_d . As shown in [13] and [49], the masking and shadowing terms in a microfacet-based BRDF model vary smoothly and are actually close to constant. These analyses motivate us to further simplify the bivariate function $\rho(\theta_h, \theta_d)$ as a factorized form $\rho_1(\theta_h)\rho_2(\theta_d)$. Similar simplification has been used in [50] to assist material capturing and editing.

To obtain a compact parametric model suitable for inverse problems, we represent the factorized terms $\rho_1(\theta_h)$ and $\rho_2(\theta_d)$

as polynomial functions of $\cos \theta_h$ and $\cos \theta_d$, respectively. As a result, our BRDF model becomes a bi-polynomial function represented as

$$\begin{aligned} \rho(\theta_h, \theta_d) &\simeq \rho_1(\theta_h)\rho_2(\theta_d) = \tilde{\rho}_1(\mathbf{n}^\top \mathbf{h})\tilde{\rho}_2(\mathbf{l}^\top \mathbf{h}) \\ &= \sum_{i=0}^k A_i (\mathbf{n}^\top \mathbf{h})^i \sum_{j=0}^k B_j (\mathbf{l}^\top \mathbf{h})^j, \end{aligned} \quad (3)$$

where k is the order of polynomials. The above equation can be further expanded by the following relaxation

$$\tilde{\rho}(x, y) = \sum_{i=0}^k \sum_{j=0}^k C_{ij} x^i y^j, \quad (4)$$

where $C_{ij} = A_i B_j$ and $x = \mathbf{n}^\top \mathbf{h}$, $y = \mathbf{l}^\top \mathbf{h}$ for notation simplicity.

In this paper, we focus on discussing the bilinear, biquadratic and bicubic models, *i.e.*, $k = 1, 2$, and 3. Let us take the biquadratic model ($k = 2$) as an example. It can be expressed as

$$\begin{aligned} \tilde{\rho}_1(\mathbf{n}^\top \mathbf{h})\tilde{\rho}_2(\mathbf{l}^\top \mathbf{h}) &= (A_2(\mathbf{n}^\top \mathbf{h})^2 + A_1(\mathbf{n}^\top \mathbf{h}) + A_0) \\ &\quad (B_2(\mathbf{l}^\top \mathbf{h})^2 + B_1(\mathbf{l}^\top \mathbf{h}) + B_0), \end{aligned} \quad (5)$$

with its linear relaxation as

$$\begin{aligned} \tilde{\rho}(x, y) &= C_{22}x^2y^2 + C_{21}x^2y + C_{20}x^2 + C_{12}xy^2 + \\ &\quad C_{11}xy + C_{10}x + C_{02}y^2 + C_{01}y + C_{00}. \end{aligned} \quad (6)$$

In the biquadratic case, there are 9 reflectance parameters in the relaxed linear model in total, and we denote them in a vector form as

$$\mathbf{x} = [C_{22}, C_{21}, \dots, C_{00}]^\top \in \mathbb{R}^{9 \times 1}. \quad (7)$$

Note that the conversion from Equation (5) to Equation (6) is unique, but the other direction is not, and Equation (6) may not always have the product form of Equation (5). Equation (6) is a linear function of its parameters $[C_{22}, C_{21}, \dots, C_{00}]^\top$, while Equation (5) is a bilinear function of $[A_2, A_1, A_0]^\top$ and $[B_2, B_1, B_0]^\top$. Hence, the relaxed model is easier to fit. The bilinear and bicubic model can be defined similarly, with 4 and 16 reflectance parameters in their relaxed forms, respectively.

It is also straightforward to express the factorized terms $\tilde{\rho}_1$ and $\tilde{\rho}_2$ in Equation (3) by higher-order polynomials, or even use different orders of $\tilde{\rho}_1$ and $\tilde{\rho}_2$. But we experimentally found that models with higher orders had little advantage in modeling accuracy and were suffered from instability. Additional discussions about the choice of orders of polynomials are left for the experiment section.

There are also other possible parameterizations of the bivariate function. We choose polynomials mainly for two reasons: 1) As discussed in [45], polynomials are good at representing smooth intensity variations (low-frequency) caused by different lightings; 2) Polynomials make the inverse problem tractable. For example, one might use the Discrete Cosine Transform (DCT) as an alternative in the domain of (θ_h, θ_d) to represent the bivariate BRDF, which also yields high modeling precision. However, when recovering the unknown surface normal \mathbf{n} , the problem becomes highly nonlinear about \mathbf{n} , because a DCT basis has the form of

$\cos(\theta_h/k) = \cos(\arccos(\mathbf{n}^\top \mathbf{h})/k)$, where k is a non-zero integer. In contrast, our polynomial model is much simpler and only involves terms like $(\mathbf{n}^\top \mathbf{h})^k$.

4.1 Relationship with other reflectance models

The bi-polynomial model can accurately represent the low-frequency component of conventional dichromatic reflectance models [51], which represent reflectance as a summation of Lambertian diffuse and specular terms. Since the specular term is mostly concentrated in the high-frequency component, the low-frequency component of these models is largely Lambertian and can be represented well by our bi-polynomial model. For example, the biquadratic model can be degraded to the Lambert's model if we set $A_2 = A_1 = B_2 = B_1 = 0$ in Equation (5). The bi-polynomial model can also represent the low-frequency component of other BRDF models that rely only on $\mathbf{n}^\top \mathbf{h}$. For instance, the Blinn-Phong model [10] can be represented with the bi-polynomial model by setting coefficients that are related to $\mathbf{l}^\top \mathbf{h}$ (except B_0) as zero.

The Cook-Torrance model [7] is widely used to represent various surface reflectances. It consists of a Lambertian diffuse, and a specular term. Its specular component S_c can be denoted as

$$S_c = \frac{k_s}{\pi} \frac{DFG}{(\mathbf{n}^\top \mathbf{l})(\mathbf{n}^\top \mathbf{v})}. \quad (8)$$

The terms D , F , and G are the microfacet distribution, Fresnel, and geometrical attenuation terms, respectively. The microfacet distribution D is represented as

$$D = \frac{1}{4m^2(\mathbf{n}^\top \mathbf{h})^4} \exp\left(\frac{\left(1 - \frac{1}{(\mathbf{n}^\top \mathbf{h})^2}\right)}{m^2}\right), \quad (9)$$

where m indicates the surface roughness, and D is clearly a function of θ_h [7]. The Fresnel term F is often simplified by the Schlick's approximation [13], [52], denoted as

$$F = k_s + (1 - k_s)(1 - \mathbf{l}^\top \mathbf{h})^5, \quad (10)$$

where k_s is a constant. Hence, F is a function of θ_d . Although the term G is relatively complicated, it varies smoothly and is close to a constant over a large range of exitant angles as evaluated in [13], [49]. Similar to the formulation of Cook-Torrance model, our model has both θ_h and θ_d terms in a product form. More importantly, as evaluated in [5], the low-frequency component of the Cook-Torrance model can be closely approximated by the biquadratic model. Other models based on the microfacet theory, such as the Ward model [12], have a similar expression for modeling the specular component, written as

$$S_w = \frac{k_s}{4\pi m^2 \sqrt{(\mathbf{n}^\top \mathbf{l})(\mathbf{n}^\top \mathbf{v})}} \exp\left(\frac{\left(1 - \frac{1}{\mathbf{n}^\top \mathbf{h}}\right)}{m^2}\right). \quad (11)$$

The low-frequency component can be purely Lambertian, or a mixture of Lambertian diffuse, and soft specular terms. When the surface roughness m is large, the specular lobe will become wide and blunt, thus the diffuse and specular cannot be separated by simply using a threshold T_{low} . Our model is designed for dealing with such cases.

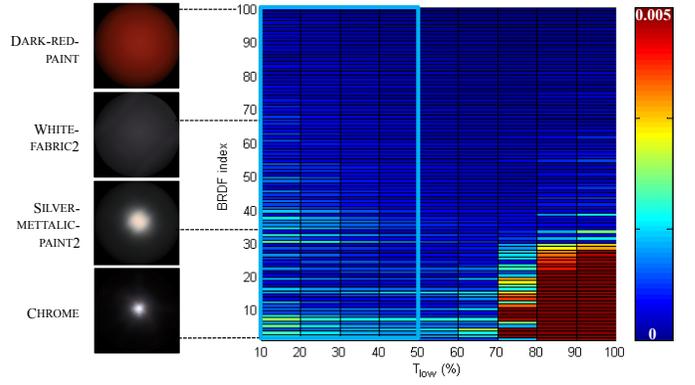


Fig. 4. BRDF fitting errors of the biquadratic model to all materials in the MERL database. The colors indicate error magnitudes. The columns vary with T_{low} , and the rows correspond to different BRDFs ordered by the mean fitting errors over columns. Some rendered spheres are displayed on the left for reference.

4.2 Model validation using measured data

To verify the representation power of the proposed model, we fit the biquadratic model to the measured BRDFs in the MERL database. We use a similar experimental setup as Section 3 with the same set of \mathbf{n} , \mathbf{l} , and \mathbf{v} . The threshold T_{low} is applied in the same manner for extracting low-frequency reflectances. Note that the fitting experiment in Section 3 is performed for each normal under varying lightings (or for each intensity profile), while here we fit the BRDF model to observations collected from different normal directions. For fitting extracted low-frequency reflectances, we first solve Equation (6) via linear least squares, through which we obtain $C_{22}, C_{21}, \dots, C_{00}$. By expanding Equation (5), we then establish a system of bilinear equations as: $C_{22} = A_2 B_2, C_{21} = A_2 B_1, \dots, C_{00} = A_0 B_0$, which can be written as $\mathbf{C} = \mathbf{a}\mathbf{b}^\top$ in a matrix form, where $\mathbf{C} \in \mathbb{R}^{3 \times 3}$ stores C_{22}, \dots, C_{00} , $\mathbf{a} = [A_2, A_1, A_0]^\top$, and $\mathbf{b} = [B_2, B_1, B_0]^\top$. This system has multiple solutions. Here, we use a singular value decomposition (SVD) as $\mathbf{C} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^\top$ for obtaining the solution as $\mathbf{a} = \mathbf{u}_1 \sqrt{\sigma_1}$ and $\mathbf{b} = \sqrt{\sigma_1} \mathbf{v}_1^\top$, where \mathbf{u}_1 is the first column vector of \mathbf{U} , \mathbf{v}_1^\top is the first row vector of \mathbf{V}^\top , and σ_1 is the greatest singular value in $\mathbf{\Sigma}$. The solution is optimal in the least squares sense. To better fit the measured data (not just to \mathbf{C}), we further perform an iterative optimization for refining \mathbf{a} and \mathbf{b} by first fixing \mathbf{a} to update \mathbf{b} , and then \mathbf{b} to update \mathbf{a} .

The RMSEs of all BRDFs (rows) varying with T_{low} (columns) are shown in Figure 4. From the region within the light blue rectangle, we can see that our model closely fits low-frequency reflectances of different materials. As a general tendency, our model has smaller fitting errors for materials with broader specular lobes.

4.3 Comparison with other parametric models

There are parametric models with general diffuse terms, such as the Lafortune and Ashikhmin² models. In the Lafortune

2. We only consider the non-Lambertian diffuse terms of these two models.

model [11], a rotationally-symmetric diffuse component D_L is written as

$$D_L = C_d (\mathbf{n}^\top \mathbf{1})^k (\mathbf{n}^\top \mathbf{v})^k, \quad (12)$$

where C_d and k are model parameters. The general diffuse term of the Ashikhmin model [13] D_A is defined as

$$D_A = R \left(1 - \left(1 - \frac{\mathbf{n}^\top \mathbf{1}}{2} \right)^5 \right) \left(1 - \left(1 - \frac{\mathbf{n}^\top \mathbf{v}}{2} \right)^5 \right), \quad (13)$$

where R is the model parameter.

Several parametric models assume a Lambertian diffuse term, and use a microfacet-based specular component, such as the Cook-Torrance and Ward models. According to Equation (8) and Equation (11), these models can be represented as

$$\rho_S = \frac{k_d}{\pi} + k_s S(\mathbf{n}, \mathbf{l}, \mathbf{v}, m), \quad (14)$$

where k_d and k_s are model parameters representing the strength of diffuse and specular terms, respectively; S is a nonlinear function with m as another model parameter. In the Cook-Torrance model, m is encoded in the term D in Equation (9).

We again use the MERL database for evaluation by setting the threshold $T_{low} = 25\%$ to extract the low-frequency reflectances. Other experiment settings are the same as those used in model validation in Section 4.2. Fitting the Ashikhmin model in Equation (13) is straightforward. For fitting the Lafortune model in Equation (12), we take the logarithm at both sides of the equation and estimate the log parameters using linear least squares. For fitting the Cook-Torrance and Ward models of Equation (14), we adopt a similar strategy as [14] and use a Matlab function “lsqnonlin” to solve the nonlinear optimization.

The fitting errors are summarized in Figure 5, and mean errors across all materials are listed in the first row of Table 1³. We take the biquadratic case as an example to test both the relaxed model of Equation (6) and the original model of Equation (5). Their average fitting errors are 7.18×10^{-4} and 7.97×10^{-4} , respectively, which suggests their approximated equivalence in accuracy. Since the relaxed form is more efficient in computation, for the bi-polynomial model we use the relaxed form in the rest of the experiments.

Our model has a consistently smaller fitting error than the general diffuse terms of the Lafortune and Ashikhmin models. The Cook-Torrance and Ward models underperform our model in representing the low-frequency reflectance, although specular terms are included in these models. This is mainly because the Cook-Torrance and Ward models behave very similarly as the Lambert’s model in low-frequency reflectance domain. The bicubic model has the highest modeling accuracy for low-frequency reflectance, while the bilinear model is less accurate than the Lafortune model. In terms of modeling complexity, the bi-polynomial model has simpler analytic forms than other parametric models, and only a linear least squares fitting is

3. According to Figure 5, the result of the Ward model is very close to the Cook-Torrance model, and error of the Ashikhmin model is much larger than others. Hence, these two models are omitted thereafter.

TABLE 1
BRDF fitting comparison.

The table shows average BRDF fitting errors (RMSE $\times 10^{-4}$) over all materials for various reflectance models. The result of noise-free data is shown in the first row, and the other rows are the results with variant levels of additive noise besides quantization. The two parameters, μ ($\times 10^{-7}$) and λ ($\times 10^{-4}$), are the weights for signal-independent and signal-dependent noise, respectively.

μ/λ	Bicub.	Biquad.	Bilin.	Laf.	C.-T.	Lamb.
No Noise	6.71	7.18	9.01	11.01	11.97	13.00
0/0	16.04	16.79	18.47	18.01	29.23	23.01
1/5	21.00	21.82	22.98	24.08	39.98	29.81
5/5	24.43	25.57	26.04	26.24	32.93	31.95
5/10	49.06	50.13	51.94	49.88	47.61	59.12
5/30	174.88	177.57	187.40	179.92	128.11	211.84
10/5	26.27	26.92	30.57	28.48	42.53	36.37

required to estimate the reflectance parameters. We have also tested a higher-order bi-polynomial model, *i.e.*, the biquartic model. However, the performance gain was rather limited and showed an almost identical error curve with the bicubic model (which is omitted in Figure 5) with RMSE of 5.60×10^{-4} . Therefore, we limit our discussion to polynomials up to the third order.

Note that the BRDF fittings here are performed using only low-frequency reflectances. The modeling accuracy of the bi-polynomial model will be deteriorated at high-frequency reflectances for materials with specularity. In such a case, the models with the specular terms such as the Cook-Torrance model may outperform our model. We refer the readers to [14] for evaluations of various parametric BRDF models in the complete BRDF domain (both low-frequency and high-frequency).

Evaluation using noisy data. The previous experiments are performed using a carefully measured data (MERL BRDFs), where a double precision is used for data storage. However, in practical scenarios, the image formation process involves various types of noise and quantization errors. Therefore, we simulate these factors and evaluate their influences on the performances of different BRDF models. Here, we consider the additive noise and 16-bit image quantization⁴. We apply both signal-independent and signal-dependent noise by adding them to the original signal as $\tilde{y} = y + (\mu + \lambda\sqrt{y})X$, which is a commonly used noise model of an imaging sensor [53]. Here, \tilde{y} and y are data with and without noise, μ and λ are weighting factors for signal-independent and signal-dependent noise, respectively, and $X \sim \mathcal{N}(0, 1)$ is a random variable following a Gaussian distribution with mean and standard deviation of 0 and 1. To apply quantization, we first cap the data with a lower bound 10^{-6} and an upper bound 1.0. The capped data are then uniformly quantized to 2^{16} levels. The original data are first corrupted by additive noise before performing quantization. We vary μ and λ to change the noise levels ($\mu = \lambda = 0$ indicates the case where only quantization noise is applied).

The average BRDF fitting errors for noisy input are summa-

4. Usually, when dealing with the general BRDFs, HDR images are used as done in [18], [39]. So we simulate 16-bit quantization instead of 8-bit LDR images.

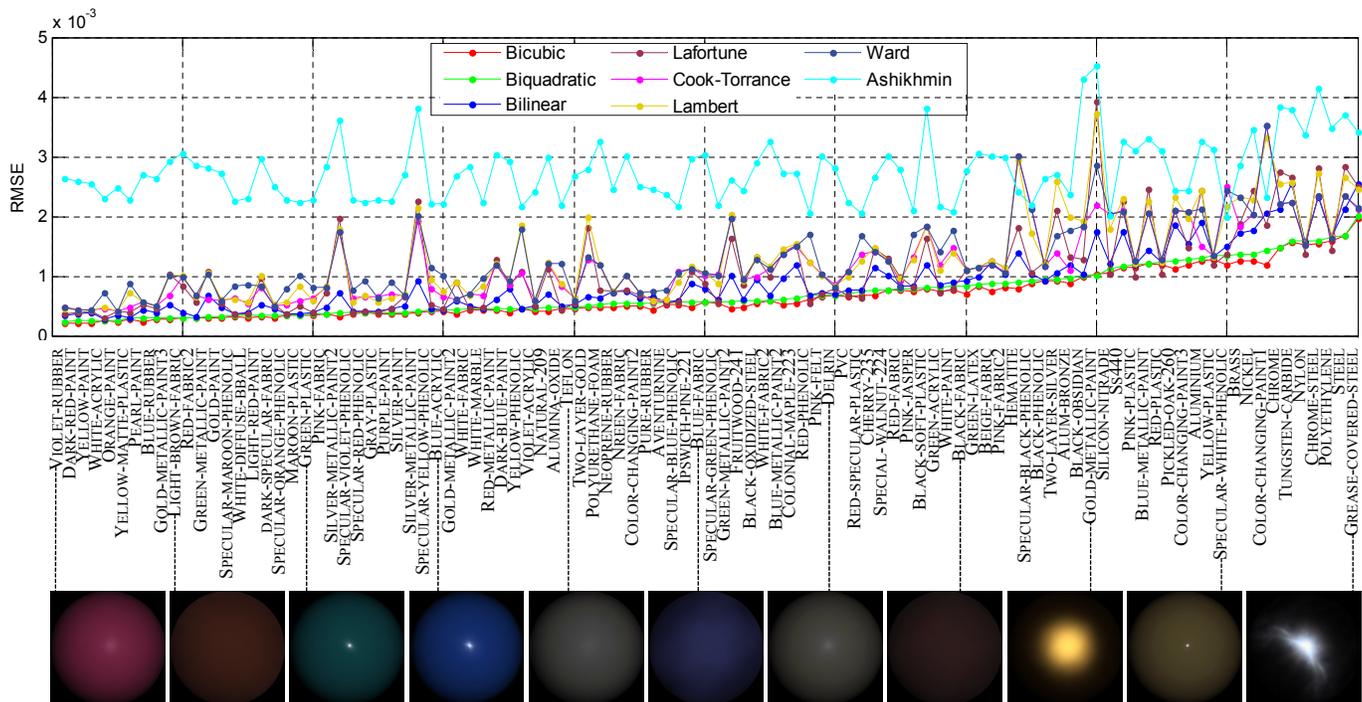


Fig. 5. BRDF fitting comparison for various reflectance models and all materials in the MERL database. The Y-axis shows the RMSE values; the X-axis shows BRDF names ordered by the fitting errors of the biquadratic model, with some renderer spheres below for a visualization purpose.

rized in the second through seventh rows of Table 1. Compared with the noise-free result in the first row, the errors become larger with the increasing noise level as anticipated. The bicubic model shows the best performance in the majority of the cases, unless the noise is too strong (fifth and sixth rows).

5 APPLICATION TO REFLECTOMETRY

Although the bi-polynomial model is designed to represent low-frequency reflectances, it can also be used for reflectometry of materials without significant specular spikes. Using the linear representation of bi-polynomial model in Equation (4), reflectometry under directional light sources only requires solving a linear equation $\mathbf{A}\mathbf{x} = \mathbf{i}$, where \mathbf{i} records radiance values. For each observation, we can calculate the matrix \mathbf{A} from \mathbf{n} , \mathbf{l} , and \mathbf{v} when the shape and lighting are all calibrated. The matrix \mathbf{A} has 4, 9 and 16 columns for bilinear, biquadratic, and bicubic models respectively according to Equation (4). In the biquadratic case for instance, from p ($p \geq 9$) independent samples, the matrix $\mathbf{A} \in \mathbb{R}^{p \times 9}$ and observations $\mathbf{i} \in \mathbb{R}^{p \times 1}$ are constructed. The model parameter \mathbf{x} can be determined by simply solving the linear system as $\mathbf{x} = (\mathbf{A}^T \mathbf{A})^{-1} \mathbf{A}^T \mathbf{i}$. If only one image of a curved surface under a directional light source is available, the reflectance shows only variations along the half angle, our model reduces to a univariate function as $\rho(\theta_h) \simeq \sum_{i=0}^k A_i (\mathbf{n}^T \mathbf{h})^i$. In such a case, the estimated reflectance only shows the variations along the half angle.

Reflectometry using measured BRDFs. To verify the method, we select some materials from the MERL database (e.g., fabric, matte-paint, rubber, etc.) that do not contain

strong specularities. We render a single image of a sphere under a directional light source as input to estimate the BRDF. We test this simple reflectometry method using various parametric BRDF models (i.e., bilinear, biquadratic, bicubic, Lafortune, Cook-Torrance, and Lambert). We then reconstruct images under the same lighting and viewing directions using the estimated reflectance parameters and evaluate the reconstruction errors. The reconstruction errors are defined as the mean of pixel-wise absolute difference.

Here we show the BLUE-FABRIC and GREEN-LATEX results as two examples in top two rows of Figure 6. We show three rendered spheres using the ground truth BRDF, fittings to the biquadratic model, and to the Lambert’s model for each example. The ground truth and estimated BRDFs of various models are visualized as 2-D polar plots. As summarized in the legends of Figure 6, the reconstruction errors show a similar performance ordering to the result in Figure 5 and Table 1, where the bicubic model performs the best, and the Lambert’s model performs the worst. From the plotted BRDF slices, we can see that the bicubic and biquadratic models fit closely to the measured data even though they only require a few parameters. The biquadratic model is a little bit worse than the bicubic model, but outperforms all other parametric models. The bilinear model shows some inaccuracy, and is worse than the Cook-Torrance model but better than the Lafortune model. Notice that this experiment setting is slightly different from that in Figure 5. Here, we fit the models to the complete BRDF domain for materials without strong highlight, while in Figure 5 we fit the models only to the low-frequency component extracted by T_{low} .

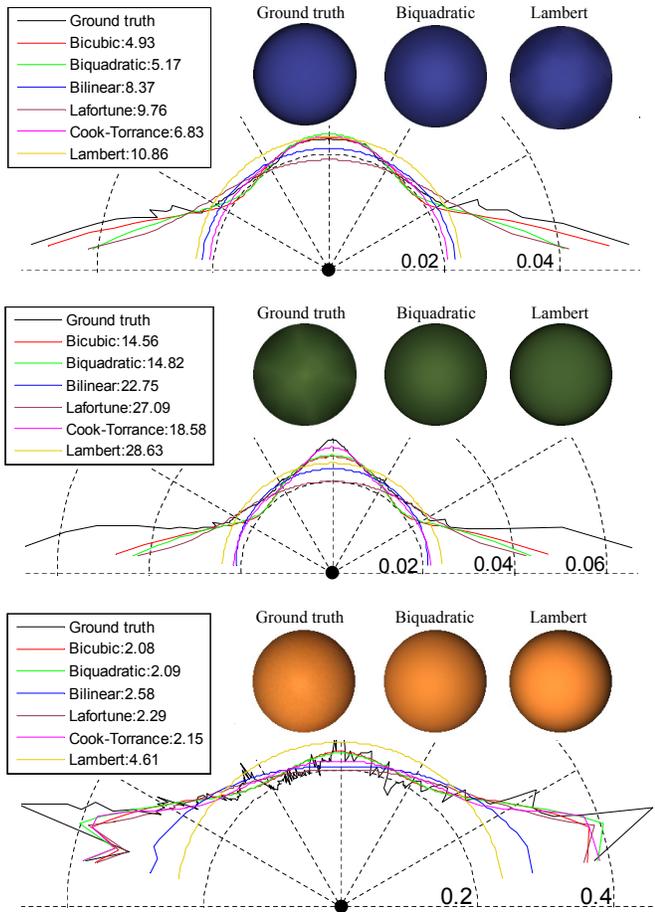


Fig. 6. Reflectometry results. Top two rows are from synthetic data, and the bottom row is from real data. BRDF plots for all models and the rendered spheres using measured data, the biquadratic model, and the Lambert’s model are shown. The reconstruction errors ($\times 10^{-4}$ for top two examples, $\times 10^{-2}$ for the bottom example) for each model are summarized in the legends. Each BRDF is visualized as a 2-D curve, which is a polar plot with angle as the elevation angle of surface normal and radius as the reflectance magnitude. The surface normal with zero azimuth angle is selected for a visualization in 2-D.

Reflectometry using real data. We also test the BRDF estimation capability using real data. A sphere under a directional lighting is recorded. We show the result in the bottom row of Figure 6, which is consistent with the synthetic test. The bicubic model performs the best with a slight advantage over the biquadratic model; modeling accuracy of the biquadratic model is much better than the bilinear model, and is close to the bicubic model. Higher-order bi-polynomial models, such as bicubic and biquartic, show only negligible improvement in our experiments, hence the biquadratic model is a good trade-off in terms of modeling accuracy and simplicity.

6 APPLICATION TO PHOTOMETRIC STEREO

In this section, we apply the bi-polynomial reflectance model to photometric stereo for estimating surface normals from

images captured by a fixed camera under varying lightings. We assume an orthographic camera and directional lightings. The camera-centered coordinate system is chosen such that $\mathbf{v} = [0, 0, 1]^T$. To use the bi-polynomial model for photometric stereo problems, we fit the reflectance model at each pixel independently. Like previous methods that deal with spatially varying BRDFs [18], our approach determines the BRDF for each pixel from its intensity observations, lighting directions, and estimated surface normals. Therefore, our method is able to handle spatially varying BRDFs.

6.1 An iterative normal estimation method

The bilinear, biquadratic, and bicubic models can be used in the same manner for solving photometric stereo. Here we use the biquadratic model as an example. From the photometric stereo images, we observe multiple radiance intensities at each pixel. For each pixel, we first use a very small intensity threshold (10^{-6} in our experiment for synthetic data) to neglect shadows. Then we sort remaining observations in an ascending order and keep only those ranked below the percentage T_{low} , which is empirically determined within [15%, 50%]. We use \mathbf{i}_{low} to denote the concatenated vector of these remaining observations and stack their corresponding lighting directions to form a matrix \mathbf{L}_{low} . We finally obtain the following equation:

$$\mathbf{i}_{low} = \bar{\rho}(\mathbf{n}, \mathbf{L}_{low}) \circ (\mathbf{n}^T \mathbf{L}_{low}), \quad (15)$$

where “ \circ ” indicates element-wise multiplication. $\bar{\rho}$ encodes the reflectance parameter \mathbf{x} in the same manner as $\tilde{\rho}$ in Equation (6), but operates on each \mathbf{n} and \mathbf{l} . The reflectance parameter \mathbf{x} represents the 9 polynomial coefficients $[C_{22}, C_{21}, \dots, C_{00}]^T$ for the biquadratic model, as defined in Equation (7). We use the relaxed biquadratic model of Equation (6) for computational efficiency. Surface normal \mathbf{n} and the BRDF parameter \mathbf{x} can be determined by iteratively optimizing the following objective function:

$$(\mathbf{n}^*, \mathbf{x}^*) = \underset{\mathbf{n}, \mathbf{x}}{\operatorname{argmin}} |\bar{\rho}(\mathbf{n}, \mathbf{L}_{low}) \circ (\mathbf{n}^T \mathbf{L}_{low}) - \mathbf{i}_{low}|_2. \quad (16)$$

At each iteration, we first fix the normal direction \mathbf{n} and refine \mathbf{x} by computing a linear least squares. We then substitute \mathbf{x} and \mathbf{n} to determine $\bar{\rho}$. Once $\bar{\rho}$ is calculated, we update \mathbf{n} again by a linear least squares. A normalization to unit-norm is performed immediately after each time \mathbf{n} is solved. To initialize this iterative optimization, we first apply Lambertian photometric stereo [26] using \mathbf{i}_{low} and \mathbf{L}_{low} to estimate the initial normal. The iterative optimization stops when the residual of Equation (16) does not change. In our implementation, we stop it when the change becomes less than 10^{-7} or a maximum iteration of 100 times is exceeded. The normal estimation algorithm is summarized as Algorithm 1. Note that we take the biquadratic model as an example in Algorithm 1, and the $\bar{\rho}$ can be replaced by other parametric models. Theoretically, to optimize \mathbf{n} with fixed \mathbf{x} , minimizing a multivariate polynomial system using Gröbner basis [54] is one of the solutions. But through experimentation, we find our simpler approach outlined above produces equivalent convergence with much lighter computation.

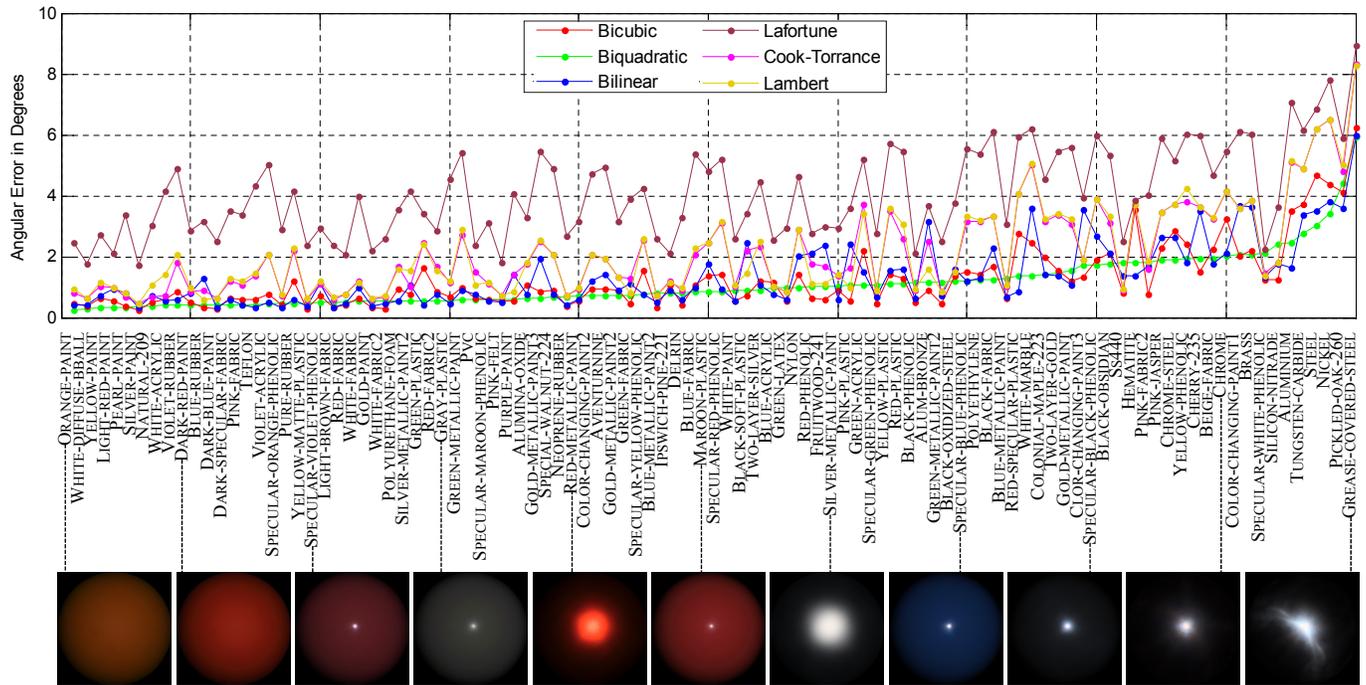


Fig. 7. Photometric stereo results comparison for various reflectance models and all materials in the MERL database. The Y-axis shows mean angular errors (degree); the X-axis shows BRDF names ordered by the mean angular errors of the biquadratic model, with some selected rendered spheres below for a visualization purpose.

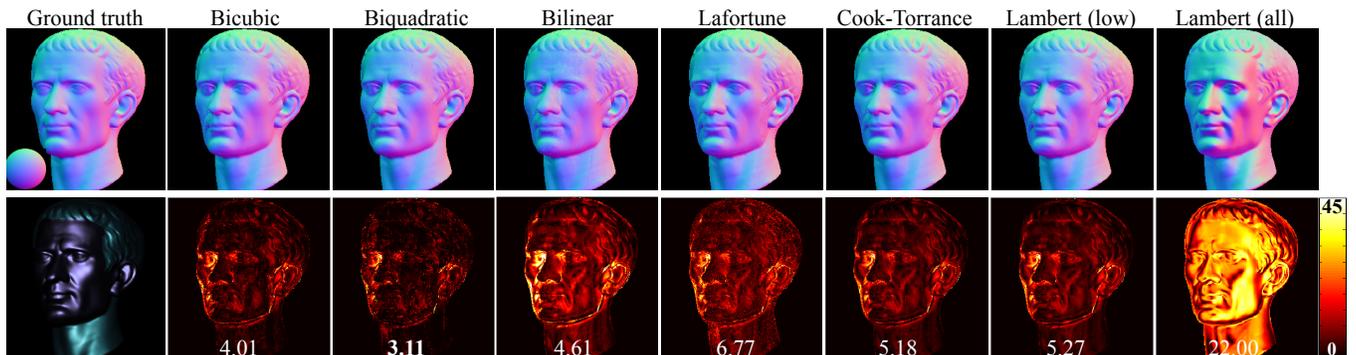


Fig. 8. Photometric stereo results using synthetic data. One of the input images is shown under the ground truth normal map. Normal map estimates using different BRDF models are shown in the top row; the bottom row shows angular difference maps w.r.t. the ground truth. The numbers on the difference maps show mean angular errors (degree).

6.2 Surface normal estimation results

Using the above iterative solution method, we perform surface normal estimation using the bi-polynomial model in comparison with other parametric BRDF models. For other parametric BRDF models, such as the Cook-Torrance and Lafortune models, we use the same iterative solution method for deriving the surface normal (simply replacing $\bar{\rho}$ with a designated model). We use the same dataset as used in Section 4.2 and T_{low} is set to 25%. Note that although the intensity thresholding is applied in the same way, for photometric stereo the BRDF fitting is performed for each pixel (intensity profile), while for the experiment in Section 4.2 the BRDF fitting is performed for all pixels with their corresponding lightings.

When the Cook-Torrance and Lafortune models are used, the estimation of their BRDF parameters by fixing \mathbf{n} becomes highly nonlinear. This causes some numerical instability. A similar issue arises when a high-order bi-polynomial model is used (empirically, higher than cubic).

The results for all 100 materials are summarized in Figure 7 with mean errors listed in the first row of Table 2 as a quantitative evaluation. Among all the tested models, the biquadratic model performs the best on the average over 100 materials. The mean error from the bilinear model is larger than the biquadratic case due to the model's poor accuracy, and the bicubic case also has larger errors than the biquadratic case due to the instability caused by high-order polynomial fitting.

Algorithm 1 Normal estimation

INPUT: Scene radiance values \mathbf{i} , lighting directions \mathbf{L} , threshold T_{low} .

for each pixel **do**

Extract \mathbf{i}_{low} and \mathbf{L}_{low} using T_{low} ;

Solve for initial \mathbf{n} with $\bar{\rho}$ as constant (Equation (15));

while resid. Equation (16) $> 10^{-7}$ **OR** #iter. ≤ 100 **do**

Update $\bar{\rho}$ by fixing \mathbf{n} (e.g., Equation (6));

Update \mathbf{n} by fixing $\bar{\rho}$ (Equation (15));

end while

end for

OUTPUT: Estimated surface normal \mathbf{n} for all pixels.

TABLE 2
Photometric stereo result comparison.

The table shows the mean angular errors (degree) of surface normal estimates over all materials for various reflectance models. The result of noise-free data is shown in the first row, and the other rows are the results with variant levels of additive noise besides quantization. The two parameters, μ ($\times 10^{-7}$) and λ ($\times 10^{-4}$), are the weights for signal-independent and signal-dependent noise respectively.

μ/λ	Bicub.	Biquad.	Bilin.	Laf.	C.-T.	Lamb.
No Noise	1.25	1.12	1.37	4.07	2.13	2.14
0/0	1.42	1.34	1.58	4.17	3.18	2.28
1/5	1.48	1.39	1.60	4.17	3.19	2.28
5/5	1.85	1.81	1.98	4.30	3.40	2.46
5/10	1.98	1.94	2.03	4.33	3.43	2.49
5/30	2.79	2.68	2.35	4.55	3.69	2.73
10/5	6.38	6.57	6.56	6.47	6.13	5.04

For the Cook-Torrance and Lafortune models, their errors are larger than the bi-polynomial model partly due to their lower modeling accuracy of low-frequency reflectances. With only having a Lambertian diffuse term, the Cook-Torrance model hardly improves the accuracy from the initial result of the Lambert’s model. In addition, these two nonlinear models cause optimization difficulties, which also explains the larger error of the Lafortune model.

Our model allows a simple alternating optimization while other reflectance models require a more sophisticated optimization technique due to their highly nonlinear nature. In our experiments, the bilinear, biquadratic, and bicubic models usually converge in a similar manner with fewer than 10 iterations. The Cook-Torrance model shows either a quick convergence after one or two iterations, or shows no decrease from the initial value, since its low-frequency term is the same as the Lambert’s model. The Lafortune model shows instability during the iterations, which is caused by both the nonlinearity of the model and the simple optimization technique that is employed. Though there is no theoretical guarantee for the convergence in our optimization method, we empirically find that this simple optimization technique shows better convergence for our model than for other models.

Evaluation using noisy data. We evaluate the effect of noisy input to photometric stereo estimates across different parametric BRDF models. The noise is applied in the same way as in Section 4.3. We summarize the mean angular errors of photometric stereo across all materials for various

reflectance models by applying different noise levels in the second through seventh rows of Table 2. The results are generally consistent with the noise-free case (first row). When the signal-dependent noise level is very strong (sixth row), the bilinear model performs best due to its simpler form. As shown in the seventh row when the magnitude of signal-independent noise reaches the shadow threshold (10^{-6}), all models show worse performances than the initial results (Lambert). This is because shadowed observations are mis-classified as low-frequency reflectances and results in unpredictable errors. Increasing the shadow threshold might solve this problem, but how to choose the optimal threshold remains difficult.

Results on spatially varying BRDFs. We also generate synthetic images with spatially varying BRDFs to show the photometric stereo results. We use two different BRDFs, namely, the BLUE-METALLIC-PAINT and GREEN-METALLIC-PAINT materials from the MERL database on a CAESAR model, and render 100 images under varying illumination as input. The estimated normal maps and their angular errors are shown in Figure 8. We linearly map x, y, z components of surface normal to the R, G, B color channels, as indicated by a reference sphere shown in the ground truth normal map. We add the results of the traditional Lambertian photometric stereo [26] with all images (“Lambert (all)”) for observing its failure mode on those materials. By only using the low-frequency reflectance, photometric stereo works much better on general reflectances even using the Lambert’s model (“Lambert (low)”), thus it can provide reasonable initial normals used for fitting all other parametric BRDF models. The performance ranking of models by mean angular errors here is consistent with Figure 7 and Table 2; the biquadratic model yields the most accurate normal estimates.

6.3 Effect of varying numbers of lightings

We perform a similar experiment as Section 6.2 with varying numbers of lighting directions (images) to observe its effect on the accuracy. We vary the number of lighting directions from 25 to 250 with a step of 25, and perform photometric stereo using various BRDF models. The threshold T_{low} is fixed to 25% in this experiment. The mean angular errors over 100 materials of MERL database are plotted in Figure 9⁵. Here we do not compare with the Lafortune model because of its unstable convergence. Generally, a larger number of images improves the results. Compared with the Cook-Torrance model, the bi-polynomial model shows greater improvement on accuracy with the increasing number of light directions, but the Cook-Torrance model still shows a similar performance to the Lambert’s model. When the number of lighting directions becomes greater, the Cook-Torrance model even shows a slightly worse result than the Lambert’s model due to its numerical instability during the optimization. Note that in our experiment, the fitting is performed using the low-frequency reflectance data.

5. The biquadratic and bicubic models need at least 9 and 16 equations for fitting, therefore their curves start from using 50 and 75 images, respectively, when $T_{low} = 25\%$.

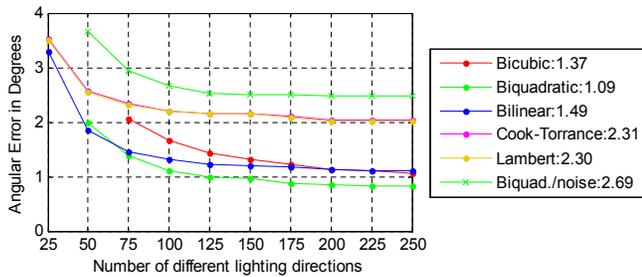


Fig. 9. Angular error (degree) varying with number of lightings. The numbers in the legend are the mean values over X -axis. The “Biquad./noise” case corresponds to the same noise level as the sixth row of Table 2.

With more input images, the low-frequency reflectances can be more reliably extracted using T_{low} . Therefore, regardless of the orders of polynomials, the bi-polynomial results generally become better as the number of input images increases. Empirically, about 100 images are sufficient for photometric stereo with the biquadratic model to produce accurate results for general isotropic reflectance. On average, the angular error becomes about 1° for the materials in the MERL database. The major reason for requiring many images is that the method involves the estimation of BRDFs, and it requires a sufficient sampling resolution in the angular domain. Similar numbers of images have been used for other state-of-the-art techniques such as [18], [36], [42]. When there is noise in the input data (curve “Biquad./noise” in Figure 9), the angular errors increase in all settings of different numbers of input lightings, but in general more lights yield higher accuracy.

6.4 Analysis on intensity threshold T_{low}

Extracting the low-frequency reflectance observations plays an important role in our BRDF modeling. We now analyze the performance variation using various T_{low} to see its influence on the photometric stereo results.

Performance variation with diverse T_{low} . Again, we perform similar experiments as Section 6.2 and Section 6.3 with the number of lighting directions fixed to 100. We plot the mean angular errors of 100 materials with varying T_{low} from 5% to 100% with a step of 5% in Figure 10⁶. The overall tendency is that all models show larger errors with increasing T_{low} . This shows that as more observations of high-frequency reflectances are involved, the problem of photometric stereo becomes more difficult. This observation agrees well with our motivation of focusing only on the low-frequency reflectance. Degradation of the Lambertian photometric stereo results with increasing T_{low} also influences the results of all other models, because they all rely on the Lambertian photometric stereo for initialization. For the Cook-Torrance model, we find that it performs very similar to the Lambert’s model when T_{low} is smaller than 30%, but it outperforms the bi-polynomial model when T_{low} is larger than 75%. For the latter case, the specularity becomes

6. The biquadratic and bicubic models need at least 9 and 16 equations for fitting, therefore their curves start from using $T_{low} = 10\%$ and 20% , respectively, when 100 images are given as input.

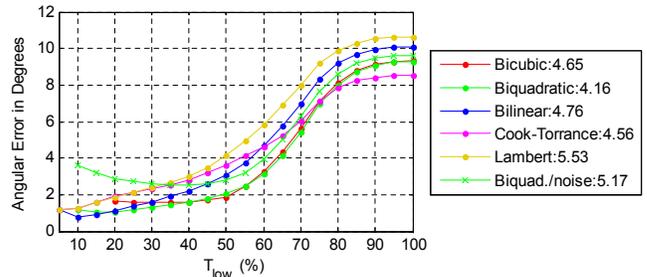


Fig. 10. Angular error (degree) varying with T_{low} . The numbers in the legend are the mean values over X -axis. The “Biquad./noise” case corresponds to the same noise level as the sixth row of Table 2.

significant, and it becomes necessary to model high-frequency components with explicit specular terms.

As for the difference between the bilinear, biquadratic, and bicubic models, we find that the biquadratic model still performs best on average across varying T_{low} . However, when T_{low} is larger than 55%, the bicubic model improves due to its greater representation ability. When T_{low} is very small (below 20%), the bilinear model can be more stably estimated due to its simpler form, and the reason here is similar to the increased noise cases in the bottom three rows of Table 2. According to Figure 10, a T_{low} around 20% is a good choice for the bi-polynomial model with ideal data.

When there is noise in the input data (curve “Biquad./noise” in Figure 10), the angular errors increase accordingly. To handle the input data with such large noise, a larger T_{low} is desired for obtaining the optimal performance. Fortunately, as discussed above, our method is not sensitive to the choice of T_{low} in the range of [20%, 50%], therefore we can pick a T_{low} that could be larger than necessary.

What materials are (in)sensitive to T_{low} ? From Figure 10 we can see that a T_{low} below 50% is safe to use for various BRDFs. However, to determine the best T_{low} for each material is not easy in our empirical model. If we plot and check the curve of angular errors varying with T_{low} for each material, we can find some material-related properties. We again take the biquadratic model as an example and plot the curves in Figure 11. We plot thin curves for all materials and use a thick curve for their average in the same figure. Based on the shape of these curves, we find that the materials can be categorized into two groups according to the sensitivity of normal estimation error to T_{low} . The sensitive materials show a sudden increase of errors when T_{low} is greater than 50%, while the insensitive materials show almost constant error curves. To check the type of materials in each group, we look at the point of $T_{low} = 70\%$ and the corresponding materials with large and small errors. We list the best nine and worst nine materials in Figure 11 with rendered spheres and material names. Interestingly, the errors of metal-like materials are sensitive to T_{low} and that of fabric-like materials are insensitive. Our method is not suitable for dealing with materials having relatively broad and strong highlight, like some metallic paint. But when the highlight is very focused or it is very weak, our method works well

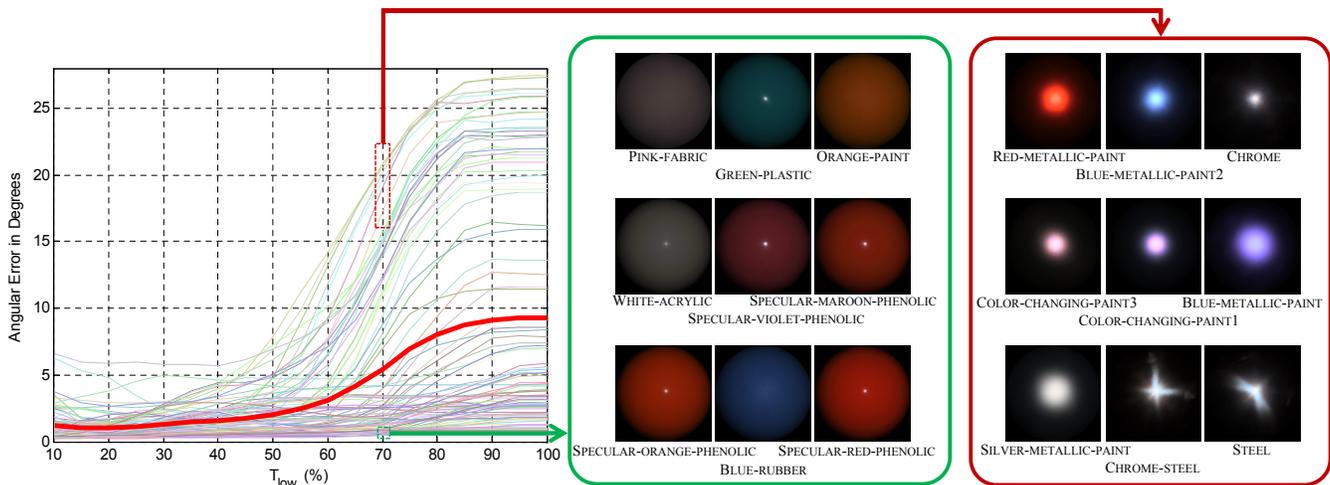


Fig. 11. Angular error (degree) varying with T_{low} for the biquadratic model. Each thin curve represents the result of one material; the thick red curve is the average of 100 materials. The materials in dark green and red frames are insensitive and sensitive material examples to T_{low} .

with relatively arbitrary selections of T_{low} . This is consistent with the above discussions, since wide and strong specularities contain significant high-frequency reflectances which cannot be easily discarded by T_{low} , while spiky specular lobes are easily separable.

6.5 Results using real-world data

We show the results of real-world data in Figure 12. We use the biquadratic model, since it is found to be the optimal one according to the synthetic test. We compare with the method from Alldrin *et al.* [18] using their datasets in the top three rows. We refer to these scenes as (from top to bottom row) GOURD1 (102), GOURD2 (98), and APPLE (112), with the number of input images in the parenthesis. Since we do not have the ground truth of these data, a quantitative evaluation cannot be performed. For a qualitative evaluation, we reconstruct the surfaces using the estimated normals and the method in [55]. In Figure 12, the left column shows one of the input images to the photometric stereo algorithm and a reference image (not used in calculation) of rendered results from [18]. The middle and right columns show the estimated surface normals and recovered surfaces using our method and the Lambertian photometric stereo. The recovered surfaces are aligned to the reference views, and our reconstruction agrees closely with the result of [18]. The data in the bottom two rows (named POST (91) and TEAPOT (73)) were captured with a Sony XCD-X710CR camera with a linear response function. While we did not carefully control exposures to avoid saturation, our method can naturally skip undesired strong specular and saturation regions by setting T_{low} . The consistency of the reconstructed surfaces using estimated normals with the pictures of the objects taken from another viewpoint indicates the effectiveness of the proposed approach.

7 CONCLUSIONS

We present the bi-polynomial reflectance model for representing low-frequency reflectances of isotropic surfaces. The

proposed reflectance model accurately captures the general diffuse reflectance that spans in the low-frequency domain in comparison with other parametric BRDF models and is useful for inverse problems such as reflectometry and surface normal recovery. We make comparisons with existing parametric models and demonstrate the usefulness of the proposed BRDF model. We also discuss the choices of different orders of polynomials, and conclude that the biquadratic model is generally the most suitable. The current model is limited to isotropic materials. In future work, we hope to analyze the characteristics of the low-frequency component of anisotropic reflectance and extend our bi-polynomial model to a wider variety of materials.

ACKNOWLEDGEMENTS

The authors would like to thank Yinqiang Zheng for helpful discussion. P. Tan was supported by the Singapore ASTAR PSF grant R-263-000-698-305.

REFERENCES

- [1] 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 (ACCV)*, 2010, pp. 703–717.
- [2] S. Ikehata, D. Wipf, Y. Matsushita, and K. Aizawa, "Robust photometric stereo using sparse regression," in *Proc. of IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2012, pp. 318–325.
- [3] S. Rusinkiewicz, "A new change of variables for efficient BRDF representation," in *Rendering Techniques (Proc. of Eurographics Workshop on Rendering)*, 1998, pp. 11–22.
- [4] F. Romeiro, Y. Vasilyev, and T. Zickler, "Passive reflectometry," in *Proc. of European Conference on Computer Vision (ECCV)*, 2008, pp. 859–872.
- [5] B. Shi, P. Tan, Y. Matsushita, and K. Ikeuchi, "A biquadratic reflectance model for radiometric image analysis," in *Proc. of IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2012, pp. 230–237.
- [6] K. Torrance and E. Sparrow, "Theory for off-specular reflection from roughened surfaces," *Journal of the Optical Society of America*, vol. 57, no. 9, pp. 1105–1114, 1967.
- [7] R. Cook and K. Torrance, "A reflectance model for computer graphics," *ACM Trans. on Graph.*, vol. 1, no. 1, pp. 7–24, 1982.

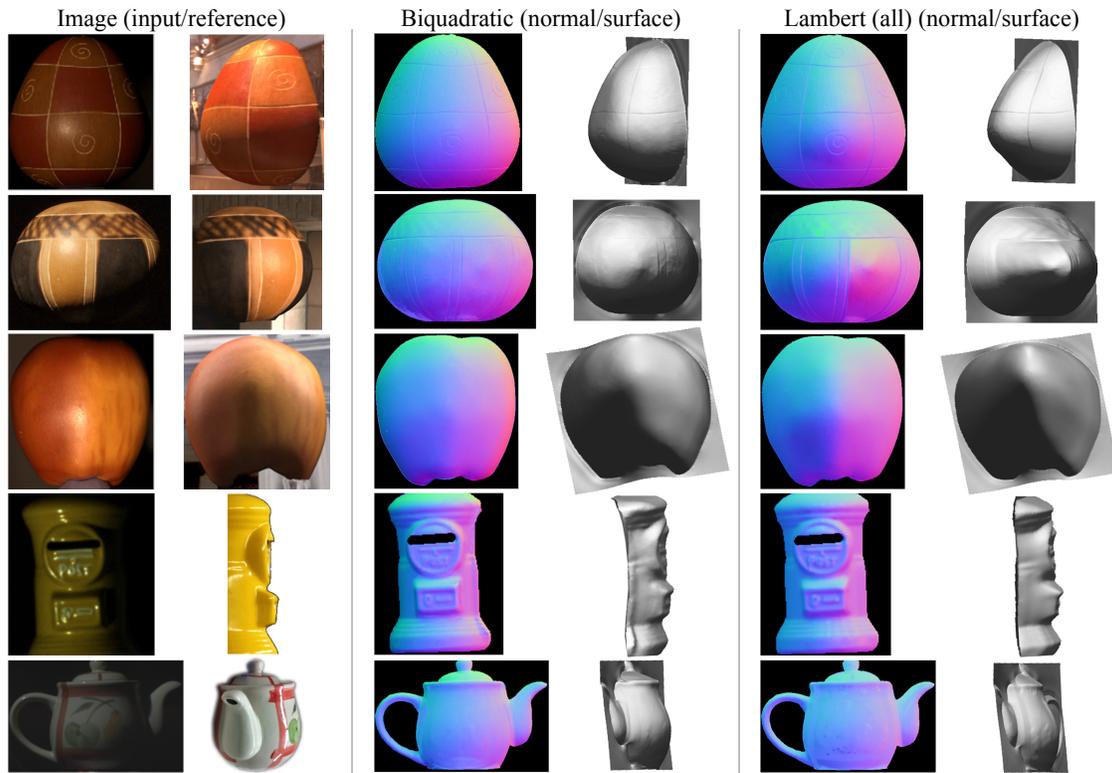


Fig. 12. Photometric stereo results using real-world data in comparison with the results of [18] and Lambertian photometric stereo. The top three rows of data are courtesy of Alldrin *et al.* [18].

- [8] M. Oren and S. Nayar, "Diffuse reflectance from rough surfaces," in *Proc. of IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 1993, pp. 763–764.
- [9] B. Phong, "Illumination for computer generated pictures," in *Proc. of SIGGRAPH*, 1975, pp. 117–126.
- [10] J. Blinn, "Models of light reflection for computer synthesized pictures," in *Proc. of SIGGRAPH*, 1977, pp. 192–198.
- [11] E. Lafortune, S. Foo, K. Torrance, and D. Greenberg, "Non-linear approximation of reflectance functions," in *Proc. of SIGGRAPH*, 1997.
- [12] G. Ward, "Measuring and modeling anisotropic reflection," *Computer Graphics*, vol. 26, no. 2, pp. 265–272, 1992.
- [13] M. Ashikhmin and P. Shirley, "An anisotropic Phong BRDF model," *Journal of Graphics Tools*, vol. 5, no. 2, pp. 25–32, 2000.
- [14] A. Ngan, F. Durand, and W. Matusik, "Experimental analysis of BRDF models," in *Proc. of Eurographics Symposium on Rendering (EGSR)*, 2005, pp. 117–126.
- [15] K. Dana, B. Ginneken, S. Nayar, and J. Koenderink, "Reflectance and texture of real world surfaces," *ACM Trans. on Graph.*, vol. 18, no. 1, pp. 1–34, 1999.
- [16] S. Marschner, S. Westin, E. Lafortune, K. Torrance, and D. Greenberg, "Image-based BRDF measurement including human skin," in *Proc. of Eurographics Symposium on Rendering (EGSR)*, 1999, pp. 139–152.
- [17] W. Matusik, H. Pfister, M. Brand, and L. McMillan, "A data-driven reflectance model," in *Proc. of SIGGRAPH (ACM Trans. on Graph.)*, 2003, pp. 759–769.
- [18] N. Alldrin, T. Zickler, and D. Kriegman, "Photometric stereo with non-parametric and spatially-varying reflectance," in *Proc. of IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2008.
- [19] F. Romeiro and T. Zickler, "Blind reflectometry," in *Proc. of European Conference on Computer Vision (ECCV)*, 2010, pp. 45–58.
- [20] T. Zickler, R. Ramamoorthi, S. Enrique, and P. Belhumeur, "Reflectance sharing: Predicting appearance from a sparse set of images," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 28, no. 8, pp. 1287–1302, 2006.
- [21] K. Nishino, "Directional statistics BRDF model," in *Proc. of International Conference on Computer Vision (ICCV)*, 2009, pp. 476–483.
- [22] M. Chandraker and R. Ramamoorthi, "What an image reveals about material reflectance," in *Proc. of International Conference on Computer Vision (ICCV)*, 2011, pp. 1076–1083.
- [23] Y. Yu, P. Debevec, J. Malik, and T. Hawkins, "Inverse global illumination: Recovering reflectance models of real scenes from photographs," in *Proc. of SIGGRAPH*, 1999, pp. 215–224.
- [24] S. Boivin and A. Gagalowicz, "Image-based rendering of diffuse, specular and glossy surfaces from a single image," in *Proc. of SIGGRAPH*, 2001, pp. 107–116.
- [25] K. Hara, K. Nishino, and K. Ikeuchi, "Mixture of spherical distributions for single-view relighting," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 30, no. 1, pp. 25–35, 2008.
- [26] R. Woodham, "Photometric method for determining surface orientation from multiple images," *Optical Engineering*, vol. 19, no. 1, pp. 139–144, 1980.
- [27] W. Silver, "Determining shape and reflectance using multiple images," *Master's thesis, MIT*, 1980.
- [28] E. Coleman and R. Jain, "Obtaining 3-dimensional shape of textured and specular surfaces using four-source photometry," *Computer Graphics and Image Processing*, vol. 18, no. 4, pp. 309–328, 1982.
- [29] F. Solomon and K. Ikeuchi, "Extracting the shape and roughness of specular lobe objects using four light photometric stereo," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 18, no. 4, pp. 449–454, 1996.
- [30] G. Kay and T. Caelly, "Estimating the parameters of an illumination model using photometric stereo," *Graphical Models and Image Processing*, vol. 57, no. 5, pp. 365–388, 1995.
- [31] A. Georghiadis, "Incorporating the Torrance and Sparrow model of reflectance in uncalibrated photometric stereo," in *Proc. of International Conference on Computer Vision (ICCV)*, 2003, pp. 816–823.
- [32] H. Chung and J. Jia, "Efficient photometric stereo on glossy surfaces with wide specular lobes," in *Proc. of IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2008.
- [33] D. Goldman, B. Curless, A. Hertzmann, and S. Seitz, "Shape and spatially-varying BRDFs from photometric stereo," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 32, no. 6, pp. 1060–1071, 2010.
- [34] P. Tan, S. Lin, and L. Quan, "Subpixel photometric stereo," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 30, no. 8, pp. 1460–1471, 2008.
- [35] N. Alldrin and D. Kriegman, "Toward reconstructing surfaces with arbitrary isotropic reflectance: A stratified photometric stereo approach," in *Proc. of International Conference on Computer Vision (ICCV)*, 2007.
- [36] B. Shi, P. Tan, Y. Matsushita, and K. Ikeuchi, "Elevation angle from

reflectance monotonicity: Photometric stereo for general isotropic reflectances,” in *Proc. of European Conference on Computer Vision (ECCV)*, 2012, pp. 455–468.

- [37] P. Tan, L. Quan, and T. Zickler, “The geometry of reflectance symmetries,” *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 33, no. 12, pp. 2506–2520, 2011.
- [38] M. Chandraker, J. Bai, and R. Ramamoorthi, “A theory of photometric reconstruction for unknown isotropic reflectances,” in *Proc. of IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2011, pp. 2505–2512.
- [39] Z. Zhou, Z. Wu, and P. Tan, “Multi-view photometric stereo with spatially varying isotropic materials,” in *Proc. of IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2013.
- [40] T. Higo, Y. Matsushita, and K. Ikeuchi, “Consensus photometric stereo,” in *Proc. of IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2010, pp. 1157–1164.
- [41] I. Sato, T. Okabe, Q. Yu, and Y. Sato, “Shape reconstruction based on similarity in radiance changes under varying illumination,” in *Proc. of International Conference on Computer Vision (ICCV)*, 2007.
- [42] F. Lu, Y. Matsushita, I. Sato, T. Okabe, and Y. Sato, “Uncalibrated photometric stereo for unknown isotropic reflectances,” in *Proc. of IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2013.
- [43] T. Okabe, I. Sato, and Y. Sato, “Attached shadow coding: Estimating surface normals from shadows under unknown reflectance and lighting conditions,” in *Proc. of International Conference on Computer Vision (ICCV)*, 2009, pp. 1693–1700.
- [44] M. Holroyd, J. Lawrence, G. Humphreys, and T. Zickler, “A photometric approach for estimating normals and tangents,” in *Proc. of SIGGRAPH Asia (ACM Trans. on Graph.)*, 2008, pp. 133–141.
- [45] T. Malzbender, D. Gelb, and H. Wolters, “Polynomial texture maps,” in *Proc. of SIGGRAPH*, 2001, pp. 519–528.
- [46] M. Drew, Y. Hel-Or, T. Malzbender, and N. Hajari, “Robust estimation of surface properties and interpolation of shadow/specularity components,” *Image and Vision Computing*, vol. 30, no. 4-5, pp. 317–331, 2012.
- [47] E. Reinhard, M. Stark, P. Shirley, and J. Ferwerda, “Photographic tone reproduction for digital images,” in *Proc. of SIGGRAPH (ACM Trans. on Graph.)*, 2002, pp. 267–276.
- [48] R. Ramamoorthi and P. Hanrahan, “Frequency space environment map rendering,” in *Proc. of SIGGRAPH (ACM Trans. on Graph.)*, 2002, pp. 517–526.
- [49] S. Nayar, K. Ikeuchi, and T. Kanade, “Surface reflection: Physical and geometrical perspectives,” *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 13, no. 7, pp. 611–634, 1991.
- [50] J. Lawrence, A. Ben-Artzi, C. DeCoro, W. Matusik, H. Pfister, R. Ramamoorthi, and S. Rusinkiewicz, “Inverse shade trees for non-parametric material representation and editing,” in *Proc. of SIGGRAPH (ACM Trans. on Graph.)*, 2006, pp. 735–745.
- [51] S. Shafer, “Using color to separate reflection components,” *Color Research*, vol. 10, no. 4, pp. 210–218, 1985.
- [52] C. Schlick, “An inexpensive BRDF model for physically-based rendering,” *Computer Graphics Forum*, vol. 13, no. 3, pp. 233–246, 1994.
- [53] N. Ratner and Y. Schechner, “Illumination multiplexing within fundamental limits,” in *Proc. of IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2007.
- [54] Z. Kukulova, M. Bujnak, and T. Pajdla, “Automatic generator of minimal problem solvers,” in *Proc. of European Conference on Computer Vision (ECCV)*, 2008, pp. 302–315.
- [55] P. Kovsi, “Shapelets correlated with surface normals produce surfaces,” in *Proc. of International Conference on Computer Vision (ICCV)*, 2005, pp. 994–1001.

Boxin Shi received the B.E. degree from Beijing University of Posts and Telecommunications in 2007 and M.E. degree from Peking University in 2010. He joined the graduate school of information science and technology at the University of Tokyo in 2010, where currently he is a Ph.D. candidate under the supervision of Prof. Katsushi Ikeuchi. His research interests include photometric methods in computer vision, reflectance and illumination modeling, and 3D reconstruction. He is a student member of IEEE.



Pin Tan received the BS degree in Applied Mathematics from the Shanghai Jiao Tong University in China in 2000 and the PhD degree in Computer Science and Engineering from the Hong Kong University of Science and Technology in 2007. He joined the Department of Electrical and Computer Engineering at the National University of Singapore as an assistant professor in 2007. He received the inaugural MIT TR35@Singapore award in 2012, and the Image and Vision Computing Outstanding Young Researcher Honorable Mention Award in 2012. He is an editorial board member of the International Journal of Computer Vision (IJCV), Machine Vision and Applications (MVA). His research interests include computer vision and computer graphics. He is a member of the IEEE and ACM.



Yasuyuki Matsushita received his B.S., M.S., and Ph.D. degrees in EECS from the University of Tokyo in 1998, 2000, and 2003, respectively. He joined Microsoft Research Asia in April 2003. He is a Lead Researcher in Visual Computing Group. His areas of research are computer vision (photometric techniques, such as radiometric calibration, photometric stereo, shape-from-shading), computer graphics (image relighting, video analysis and synthesis). He is on the editorial board member of IEEE Transactions on

Pattern Analysis and Machine Intelligence (TPAMI), International Journal of Computer Vision (IJCV), IPSJ Journal of Computer Vision and Applications (CVA), The Visual Computer Journal, and Encyclopedia of Computer Vision. He served/is serving as a Program Co-Chair of PSIVT 2010, 3DIMPVT 2011, and ACCV 2012 and a General Co-Chair for ACCV 2014. He is appointed as a Guest Associate Professor at Osaka University (April 2010-), Visiting Associate Professor at National Institute of Informatics (April 2011-) and Tohoku University (April 2012-), Japan. He is a senior member of IEEE.



Katsushi Ikeuchi received the BE degree from Kyoto University in 1973 and the PhD degree from the University of Tokyo in 1978. After working at the Massachusetts Institute of Technology (MIT) Artificial Intelligence (AI) Laboratory for three years, ETL for five years, and Carnegie Mellon University (CMU) Robotics Institute for 10 years, he joined the University of Tokyo in 1996 and is currently a full professor. His research interest spans computer vision, robotics, and computer graphics. In these research fields, he

has received several awards, including the David Marr Prize in computational vision for the paper “Shape from Interreflection”, and IEEE Robotics and Automation Society K. S. Fu memorial best transaction paper award for the paper “Toward Automatic Robot Instruction from Perception-Mapping Human Grasps to Manipulator Grasps”. In addition, in 1992, his paper “Numerical Shape from Shading and Occluding Boundaries” was selected as one of the most influential papers to have appeared in the Artificial Intelligence Journal within the past 10 years. His IEEE activities include general chair of IROS '95, ITSC '00, IV '01; program chair of the CVPR '96 and ICCV '03; associate editor of IEEE Transactions on Robotics and Automation, IEEE Transactions on Pattern Analysis and Machine Intelligence; and a distinguished lecturer of the Signal Processing Society in 2000-2002 and Robotics and Automation Society in 2004-2006. He was elected an IEEE fellow in 1998. He is the editor-in-chief of the International Journal of Computer Vision. He received the Computer Vision Significant Researcher Award in ICCV '11.