Computer Vision and Image Understanding 116 (2012) 793–802

Contents lists available at SciVerse ScienceDirect

Computer Vision and Image Understanding

journal homepage: www.elsevier.com/locate/cviu

Surface reflectance and normal estimation from photometric stereo ☆

Qingxiong Yang [a,*], Narendra Ahuja [b]

[a] Department of Computer Science, City University of Hong Kong, Hong Kong
[b] Department of Electrical & Computer Engineering, University of Illinois at Urbana–Champaign, IL 61801, United States

ARTICLE INFO

Article history:
Received 26 June 2009
Accepted 1 March 2012
Available online 15 March 2012

Keywords:
Reflection components separation
Specular reflection
Dichromatic reflection model
Chromaticity
Photometric stereo

ABSTRACT

In this paper, we propose a new photometric stereo method for estimating diffuse reflection and surface normal from color images. Using dichromatic reflection model, we introduce *surface chromaticity* as a matching invariant for photometric stereo, which serves as the foundation of the theory of this paper. An extremely simple and robust reflection components separation method is proposed based on the invariant. Our separation method differs from most previous methods which either assume dependencies among pixels or require segmentation. We also show that a linear relationship between the image color and the surface normal can be obtained based on this invariant. The linear relationship turns the surface normal estimation problem into a linear system that can be solved exactly or via least-squares optimization. We present experiments on both synthetic and real images, which demonstrate the effectiveness of our method.

1. Introduction

Since it was first introduced in [33], photometric stereo (PS) has received extensive attention. Woodham [33,34] used PS to recover local surface orientation of Lambertian surfaces. PS methods for non-Lambertian surfaces can be divided into three categories: one excludes highlight regions; another assumes the knowledge of the reflectance; and the third extracts the diffuse components or the specular-free images as a pre-processing step.

Coleman and Jain [3] extended Woodham's method [33] by introducing a fourth light source. Their method calculates the albedo and recovers the surface normal using three light sources that produce the smallest reflectance factor. This method assumes that the specular regions in different images do not intersect, and requires three non-specular measurements out of four measurements. Barsky and Petrou [2] refined this technique by using color information to improve the detection of specular measurements without explicitly separating the reflection components. In all these approaches, the specular measurements are treated as outliers and the specular component is assumed to have limited angular support.

Another category consists of methods that either acquire or assume reflectance information. For example, a reflectance model may be first estimated in terms of a linear basis by using a reference object [7,24]. Alternatively, an analytic BRDF model may be assumed and evaluated [8,15,25]. When these reflectance assumptions are not satisfied, the accuracy of the recovered shape can be compromised.

The last category assumes dichromatic surfaces, and explicitly recovers the diffuse and specular components or the specular-free images. Schlüns [22] assumed homogeneous dichromatic surfaces, and separated the diffuse and specular components using color histogram analysis techniques, similar to Klinker et al. [9]. Sato and Ikeuchi [21] took a similar approach, but avoided the restriction to homogeneous surfaces by using a large number of light source directions to compute a distinct color histogram at each point. The main problem of these approaches is that they require segmentation which is difficult when dealing with highly textured surfaces. Recently, a number of reflection component separation techniques [26,27,30,28] have overcome these limitations. These techniques seem to be good for PS methods since they demonstrated the ability of extracting the diffuse component from a single image. However, none of these methods appear to have been extended to shape estimation from PS. Other reflection separation approaches using multiple images captured under changing viewing conditions, such as different polarizer angles [16] and viewpoints [12], have been effective in separating diffuse and specular reflections but are not suitable for PS in which the viewing direction is assumed to be fixed. Lin [13] presented a two-view reflection separation approach under different light directions with the assumption that the specular highlights do not spatially overlap between the images. This approach seems reasonable for PS, but was only focus on reflection components separation and was not extended for PS. Instead of explicitly extracting the diffuse components from the specular images, Mallick et al. [14] computed

☆ The support of Hewlett-Packard under the Open-Innovation Research program is gratefully acknowledged.

* Corresponding author.
E-mail addresses: qiyang@cityu.edu.hk (Q. Yang), n-ahuja@illinois.edu (N. Ahuja).

1077-3142/$ - see front matter © 2012 Elsevier Inc. All rights reserved.
http://dx.doi.org/10.1016/j.cviu.2012.03.001

two-channel "pseudo-coded" images for PS and used them as the input for conventional PS methods. Differ from [14], we derived a linear relationship is directly on the image colors and surface normal. Also, instead of producing two-channel "pseudo-coded" diffuse images, our method explicitly extracts the three-channel diffuse images.

In this paper, we present a new photometric stereo method for both highlight reflection and surface normal estimation using the dichromatic reflection model. In this model, we define surface chromaticity to be the fraction of the diffuse reflection present in each color component, which is invariant to light direction and intensity. By assuming that for every pixel, at least one of the camera images is free of specular highlights, we derive a direct correlation between the color of the diffuse pixel and the diffuse component of the specular pixel in the other views based on the invariant. With additional assumptions that are common for PS, e.g., the light direction, intensity and chromaticity are known or can be estimated, we further derive a linear relationship between image color and the surface normal for non-Lambertian surfaces. Our contributions include:

- An extremely simple and robust method for reflection components separation based on the surface chromaticity invariant.
- A novel surface normal estimation method for non-Lambertian surfaces by deriving a linear relationship between the image color and the surface normal. This linear relationship turns the surface normal estimation problem into a linear system that can be solved exactly or via least-squares optimization.
- Demonstrating accurate surface reconstructions of specular objects, in a standard setup in which the illumination direction and strength can be estimated using a mirror sphere and a diffuse sphere, respectively.

2. Approach

In this section, we first give a brief overview of the existing BRDF models (Section 2.1) and the dichromatic reflection model we adopted (Section 2.1), and then derive solutions for reflection components separation (Section 2.3) and photometric stereo on non-Lambertian surfaces (Section 2.4).

2.1. BRDF Models

The bidirectional reflectance distribution function (BRDF) was first defined by Nicodemus [18]. It is a function that defines how light is reflected at an opaque surface. The function takes an incoming light direction w_i, and outgoing direction w_o, and returns the ratio of reflected radiance exiting along w_o to the irradiance incident on the surface from direction w_i. Existing BRDF models can be separated into two categories: empirical and physical models. Both empirical and physical BRDF models are only approximations of the reflectance properties of real materials. Many of these models are based on material parameters that in principle could be measured, but in practice are difficult to acquire.

Lambertian reflection is often used as a model for diffuse reflection. An ideal diffuse reflecting surface will have equal luminance from all directions in the hemisphere surrounding the surface. Phong reflection model [19] is neither theoretically plausible or experimentally correct. Nevertheless, it is still one of the most widely used because of its mathematically simplicity which is indispensable in computer graphics. Ward [32] developed a model that is capable of describing most significant reflection phenomena. His model obeys the most basic of physical laws (reciprocity and energy conversation) and it is relatively simple compared to most other analytical reflectance models. The Lafortune model [10] is a function with nonlinear parameters capable of approximating a broad class of BRDFs. The model is simple, compact and computationally efficient. It is also physically correct in the sense that it is reciprocal and energy-conserving, and can describe a variety of physical effects.

A comprehensive but computationally expensive model based on physical theory was developed by He et al. [6]. Torrance and Sparrow [31], Cook and Torrance [4] and Ashikhmin et al. [1] have presented BRDF models based on surfaces made up of small planar microfacets with varying angle and size. The Torrance–Sparrow and Cook–Torrance models are physically-based models that assume the microfacets on the surface form "V-cavities". In contrast, Ashikhmin et al. is an empirical model. These models account for masking and self-shadowing effects and predict off-specular reflection. Cook–Torrance [4] and Ashikhmin [1] models have been demonstrated to match real data quite well [17].

2.2. Dichromatic reflection model

The characteristics of the BRDF will determine what "type" of material the viewer thinks the displayed object is composed of, so the choice of BRDF model and its parameters is important. However, in this paper, we are not focus on scene rendering but the removing of specular highlights, thus we adopt a very simple BRDF model – dichromatic reflection model – which is a classical sum of a diffuse term and a specular term:

$$I_c = I_c^{Diffuse} + I_c^{Specular} \tag{1}$$

$$= M^{Diffuse}\Lambda_c + M^{Specular}\Gamma_c, \tag{2}$$

where $c \in \{R,G,B\}$ is the color component,

$$M^{Diffuse} = \sum_{c\in\{R,G,B\}} I_c^{Diffuse} \tag{3}$$

and

$$M^{Specular} = \sum_{c\in\{R,G,B\}} I_c^{Specular} \tag{4}$$

are the total diffuse and specular reflection over all colors (at some pixel location), respectively, and

$$I_{sum} = \sum_{c\in\{R,G,B\}} I_c = M^{Diffuse} + M^{Specular}, \tag{5}$$

is the total reflection over all colors,

$$\Lambda_c = I_c^{Diffuse} \Big/ \left(\sum_{u\in\{R,G,B\}} I_u^{Diffuse} \right) = I_c^{Diffuse}/M^{Diffuse} \tag{6}$$

is the *surface/diffuse chromaticity* denoting the fraction of the color component c presented in the diffuse reflection, and

$$\Gamma_c = I_c^{Specular} \Big/ \left(\sum_{u\in\{R,G,B\}} I_u^{Specular} \right) = I_c^{Specular}/M^{Specular} \tag{7}$$

is the *illumination chromaticity* denoting the fraction of the color component c presented in the specular reflection.

In this paper, we assume that a single light source is used and the images will be captured by changing the light source location. According to the definition of the dichromatic reflection model, the specular highlight only reflects the color of the light source, not the color of the reflecting object, thus Γ_c is the same throughout the whole image and is invariant to the light direction and intensity. The use of dichromatic reflection model also confine surface chromaticity Λ_c to be invariant to the light direction and intensity. The dichromatic reflection model is invalid for metallic materials such as gold as the color of the specular highlight will reflect the color of the material.

2.3. Reflection estimation from photometric stereo

As discussed in Section 2.2, the surface chromaticity Λ_c and illumination chromaticity Γ_c are invariant to light direction and intensity. If I and J are any two color images captured for photometric stereo, $M_I^{Specular}$ and $M_J^{Specular}$ are defined as the sum of all specular color components of I and J, respectively, then according to the definition of surface chromaticity in Eqs. (6) and (2),

$$\Lambda_c = \frac{I_c - M_I^{Specular}\Gamma_c}{\sum_{u\in\{r,g,b\}}\left(I_u - M_I^{Specular}\Gamma_u\right)} = \frac{J_c - M_J^{Specular}\Gamma_c}{\sum_{u\in\{r,g,b\}}\left(J_u - M_J^{Specular}\Gamma_u\right)}. \quad (8)$$

We assume that for every pixel $\mathbf{x}$, at least one of the color images is diffuse and let it be image J. More specifically, let $M_J^{Specular}(\mathbf{x}) = 0$. We refer to this assumption as *specular-free observability assumption*. Under this assumption, we show that there is a direct correlation between the image color of the diffuse pixel $J_c(\mathbf{x})$ and the color of the diffuse component $I_c^{Diffuse}(\mathbf{x})$ of the specular pixel in image I. According to Eq. (8), for a pixel location $\mathbf{x}$ and $M_J^{Specular}(\mathbf{x}) = 0$, we have

$$\frac{I_c(\mathbf{x}) - M_I^{Specular}(\mathbf{x})\Gamma_c}{\sum_{u\in\{r,g,b\}}(I_u(\mathbf{x}) - M_I^{Specular}(\mathbf{x})\Gamma_u)} = \frac{I_c(\mathbf{x}) - M_I^{Specular}(\mathbf{x})\Gamma_c}{\sum_{u\in\{r,g,b\}}I_u(\mathbf{x}) - M_I^{Specular}(\mathbf{x})} = \frac{J_c(\mathbf{x})}{\sum_{u\in\{r,g,b\}}J_u(\mathbf{x})}, \quad (9)$$

thus

$$M_I^{Specular}(\mathbf{x}) = \frac{J_c(\mathbf{x})\sum_{u\in\{r,g,b\}}I_u(\mathbf{x}) - I_c(\mathbf{x})\sum_{u\in\{r,g,b\}}J_u(\mathbf{x})}{J_c(\mathbf{x}) - \Gamma_c\sum_{u\in\{r,g,b\}}J_u(\mathbf{x})} = \frac{J_c(\mathbf{x})I_{sum}(\mathbf{x}) - I_c(\mathbf{x})J_{sum}(\mathbf{x})}{J_c(\mathbf{x}) - \Gamma_c J_{sum}(\mathbf{x})}. \quad (10)$$

Note that c in Eq. (10) can be any color channel. In this paper, we choose $c = m$, where $m = \arg\min_{c\in\{r,g,b\}}J_c(\mathbf{x})$. According to Eq. (2), the diffuse component of image I at pixel location $\mathbf{x}$ is then:

$$I_c^{Diffuse}(\mathbf{x}) = I_c(\mathbf{x}) - \Gamma_c M_I^{Specular}(\mathbf{x}) = I_c(\mathbf{x}) - \Gamma_c\frac{J_m(\mathbf{x})I_{sum}(\mathbf{x}) - I_m(\mathbf{x})J_{sum}(\mathbf{x})}{J_m(\mathbf{x}) - \Gamma_m J_{sum}(\mathbf{x})}. \quad (11)$$

We place a conservative threshold on minimum intensity to prevent $J_{sum}(\mathbf{x}) = 0$. Note that even if the specular-free observability assumption is violated, the proposed approach will still reduce the specular effects (see Fig. 2). Also note that there is no requirement of lighting calibration.

2.4. Surface normal estimation from photometric stereo

The original PS algorithm of [33] gives us the ability to estimate local surface normal by using a number of photometric images of the same surface, taken from the same viewpoint but under

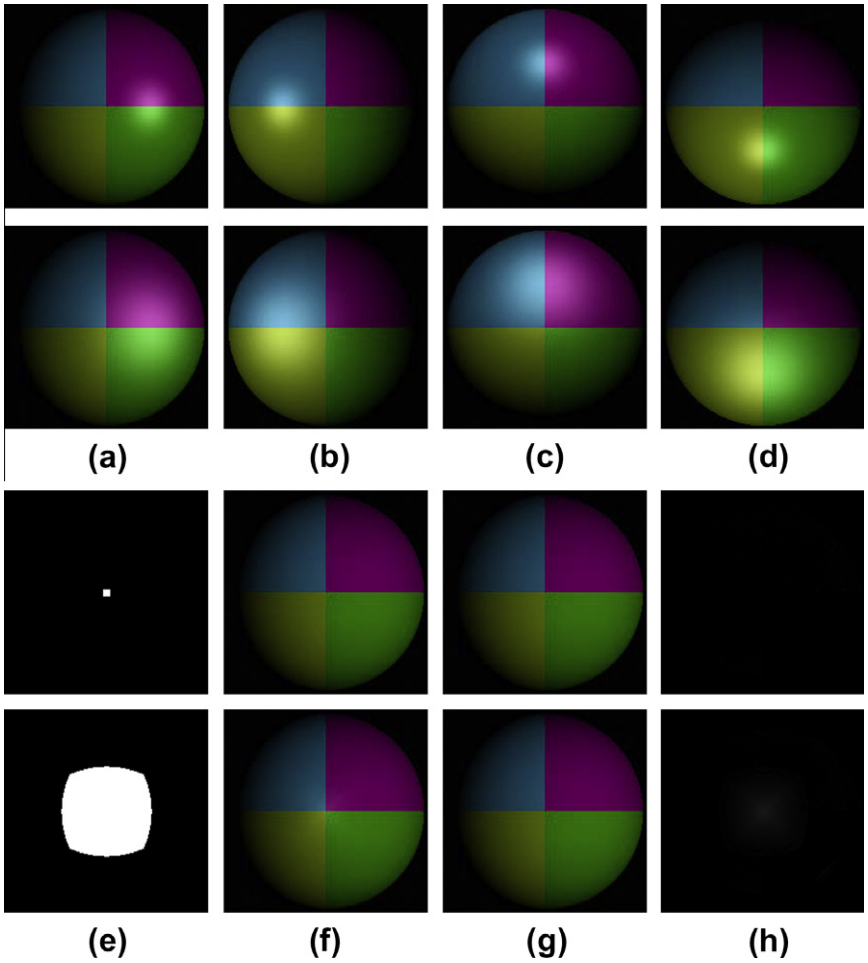

Fig. 1. Reflection components separation from four synthetic images. The first row shows the synthetic images generated by Phong reflection model using the shininess constant value of 40, and the second row sets the shininess constant to 10. (a–d) are the four input images; (e) shows the regions containing the pixels violating the specular-free observability assumption; (f) shows the recovered diffuse components for (a); (g) are the ground-truth diffuse components corresponding to (f); (h) provide the errors by comparing the extracted diffuse components with the ground-truth values. There is no noticeable error in the first row of (f), but some in the second row due to the violation of the specular-free observability assumption.

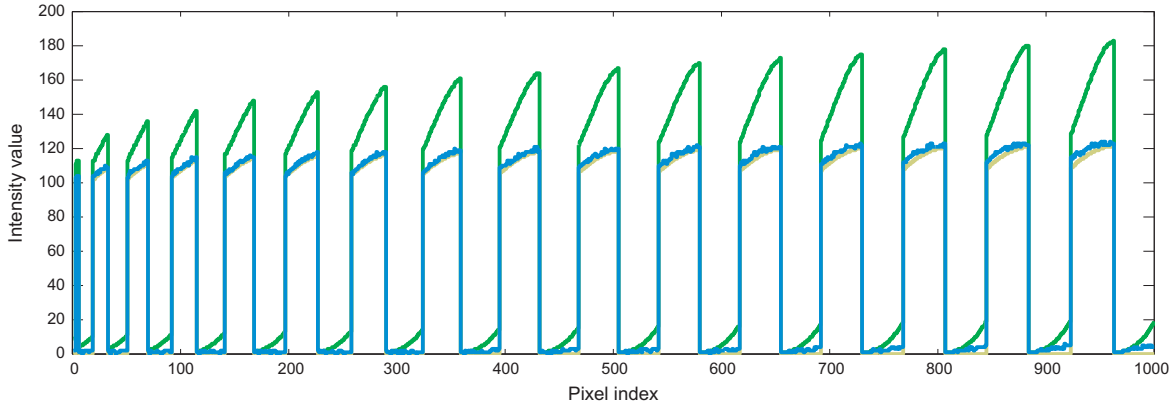

(a) Red channel

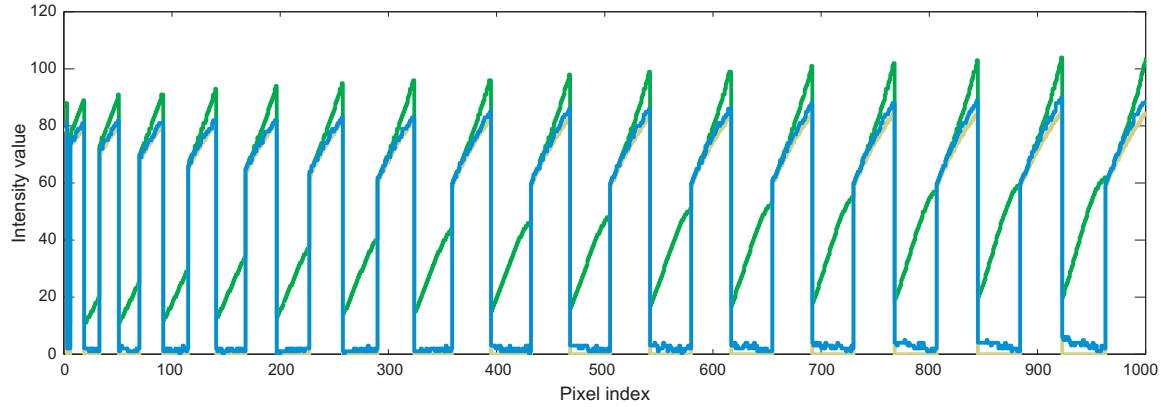

(b) Green channel

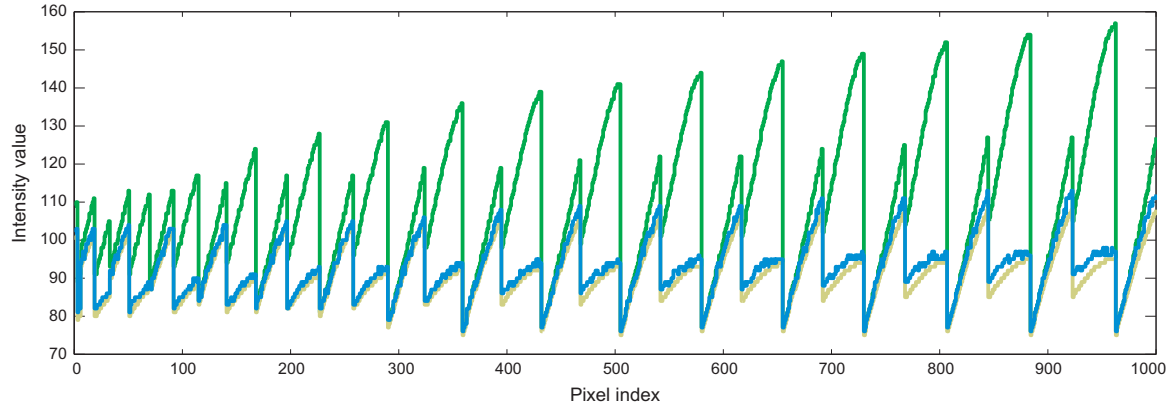

(c) Blue channel

Fig. 2. Robust reflection components separation. The green, blue and yellow curves are the intensity values of the synthetic images with the shininess constant set to 10, the estimated diffuse component, and the ground-truth diffuse component, respectively. Only the intensity values of part of the pixels violating the specular-free observability assumption are shown in the figure. Apparently, our method robustly reduces the specular effects even when the assumption is violated. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

Table 1
Evaluation of the estimated diffuse components against the ground-truth diffuse components. Apparently, error increases as the shininess constant decreases, because the number of pixels violating the specular-free observability assumption increases.

Shininess	Avg error	Min error	Max error	Std error
40	0.2036	0.0000	9.0000	0.5955
10	1.1739	0.0000	33.000	3.3954

illumination from different directions. Assuming Lambertian surfaces, the pixel intensity $I_c(=I_c^{Diffuse})$ in the images depends on local surface normal $\vec{n}=[n_x,n_y,n_z]^T$, albedo ρ_c, illumination direction $\vec{l}=[l_x,l_y,l_z]^T$ and strength L:

$$I_c = I_c^{Diffuse} = \rho_c L\vec{l}\cdot\vec{n}, \tag{12}$$

where $c\in\{R,G,B\}$. Assume L and $\vec{l}$ are given or can be calibrated, $\vec{n}$ and ρ_c can be recovered by solving the linear Eq. (12) with as few as three color images.

Using the dichromatic reflection model [23], we prove that a linear solution of Eq. (12) exists for specular surfaces by showing

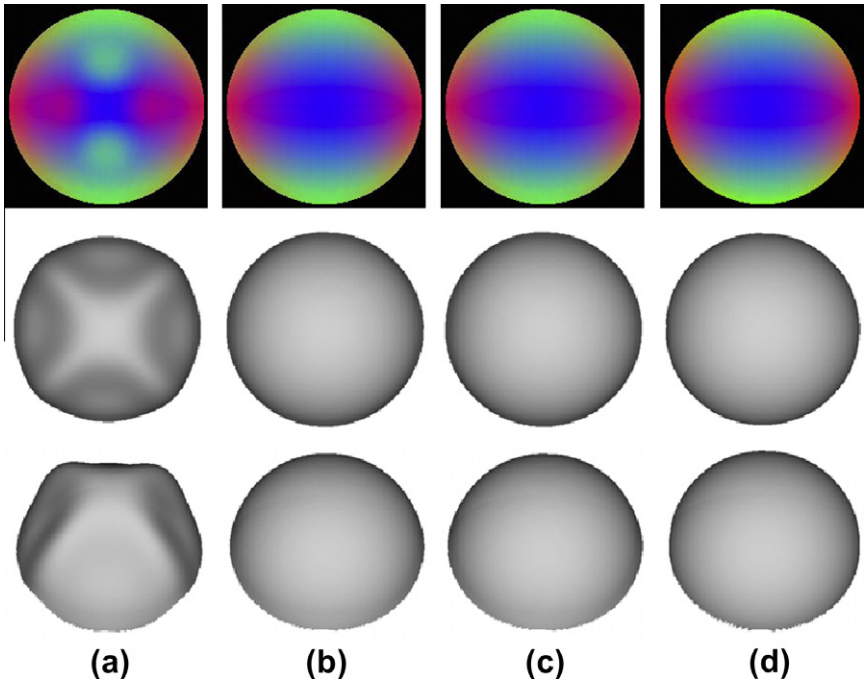

Fig. 3. Surface normal estimation from four synthetic images. (a) Experimental results using conventional PS approach with the captured images as the input. (b) Results using our approach. (c) Results using the conventional approach but with the estimated diffuse components as the input. (d) Ground-truth values. The first column shows the surface normal, where R, G, B channels represent the values of the x, y, z components (absolute values are displayed here). The last two columns show the corresponding shaded shapes.

Table 2
Evaluation against the ground-truth surface normal. The values in the table are the root-mean-square angular difference of the estimated surface normal and the ground truth in the specular regions. Conventional PS results are provided in the first row, the results of our method are in the second row, and those of the conventional method with the estimated diffuse components are presented in the last row. Obviously, the last two methods perform much better than the conventional method, and our method is the best for this synthetic data set. The conventional method using the estimated diffuse components is slightly worse because there are some errors in the estimated diffuse components as presented in Table 1.

Method	Avg error	Min error	Max error	Std error
Conv.	0.1720	0.0119	0.4157	0.1147
Ours	0.0070	0.0000	0.0238	0.0039
Conv. + RS	0.0075	0.0000	0.0261	0.0047

a linear relationship between the image color I_c and the surface normal $\vec{n}$.

Let $I_{sum} = \sum_{c\in\{R,G,B\}} I_c$, then from Eq. (2), we obtain $I_{sum} = (M^{Diffuse} + M^{Specular})$ and

$$I_{sum}\Gamma_c = (M^{Diffuse} + M^{Specular})\Gamma_c. \tag{13}$$

Substitute Eq. (13) into (2), we obtain

$$I_c = M^{Diffuse}\Lambda_c + \left(I_{sum} - M^{Diffuse}\right)\Gamma_c \Rightarrow$$

$$I_c = M^{Diffuse}(\Lambda_c - \Gamma_c) + I_{sum}\Gamma_c \Rightarrow$$

$$M^{Diffuse}(\Lambda_c - \Gamma_c) = I_c - I_{sum}\Gamma_c \Rightarrow$$

$$M^{Diffuse} = \frac{I_c - I_{sum}\Gamma_c}{\Lambda_c - \Gamma_c}. \tag{14}$$

According the definition of $M^{Diffuse}$ Eqs. (3) and (14) , we have

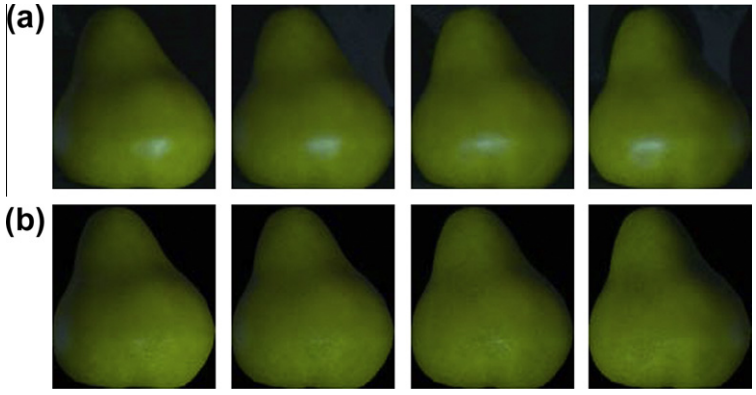

Fig. 4. The *pear* data set. (a) Four out of the 12 input color images. (b) The corresponding diffuse components recovered using our approach. *The images are gamma corrected for better illustration.*

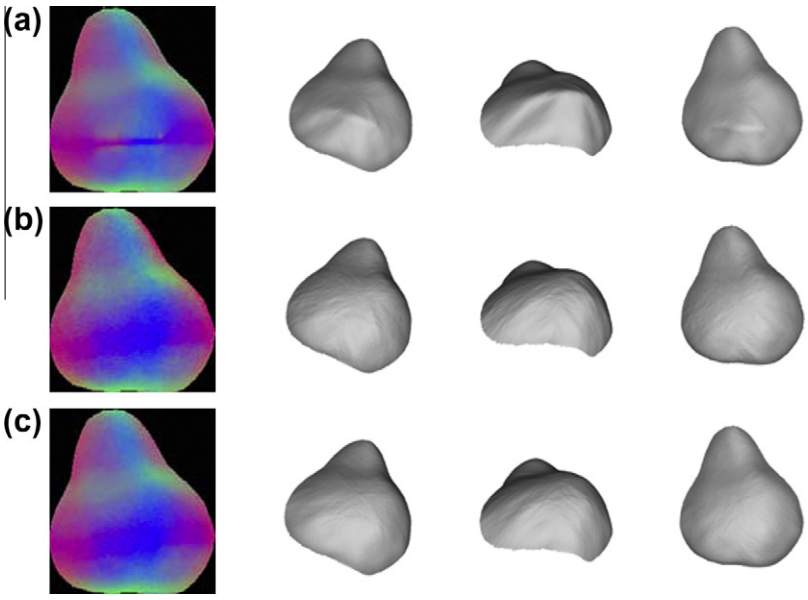

Fig. 5. Surface reconstruction. (a) Experimental results with conventional approach. (b) Results with the proposed approach. (c) Results with conventional approach using the estimated diffuse components as inputs. The first column shows the estimated surface normal (absolute values), and the last three columns show the corresponding shaded shapes from different viewing direction for visual comparison.

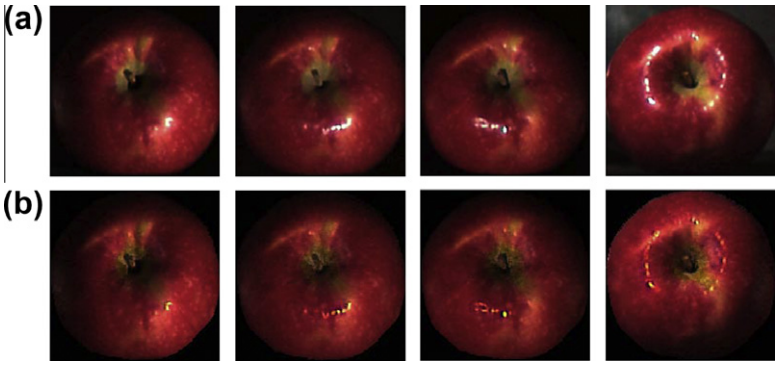

Fig. 6. The *apple* data set. (a) Four out of the 12 input color images. (b) The corresponding diffuse components recovered using our approach.

$$\sum_{c\in\{r,g,b\}} I_c^{Diffuse} = M^{Diffuse} = \frac{I_c - I_{sum}\Gamma_c}{\Lambda_c - \Gamma_c}. \tag{15}$$

From (12) and (15),

$$I_c - I_{sum}\Gamma_c = (\Lambda_c - \Gamma_c)\sum_{c\in\{r,g,b\}} \rho_c L\vec{l}\cdot\vec{n}. \tag{16}$$

According to the dichromatic reflection model, the surface chromaticity Λ_c and illumination chromaticity Γ_c are invariant to light directions and strength, thus assume $\rho'_c = (\Lambda_c - \Gamma_c)\sum_{c\in\{r,g,b\}}\rho_c$, ρ'_c will be invariant to the illumination too, and

$$I_c - I_{sum}\Gamma_c = \rho'_c L\vec{l}\cdot\vec{n}. \tag{17}$$

If we define

$$\mathbf{A} = \begin{bmatrix} l_x^1 & l_y^1 & l_z^1 \\ \vdots & \vdots & \vdots \\ l_x^m & l_y^m & l_z^m \end{bmatrix}$$

and

$$\mathbf{B} = \begin{bmatrix} (I_c - I_{sum}\Gamma_c)^1 \\ \vdots \\ (I_c - I_{sum}\Gamma_c)^m \end{bmatrix},$$

from a set of m color images, we obtain a set of m linear equations of the form

$$AX = B, \tag{18}$$

where $X = \rho'_c L\vec{n}$, $\vec{n}$ is the normalized vector of X to be estimated, ρ'_c is a constant for all the m views and L is the light strength that can be measured. For $m \geqslant 3$, Eq. (18) may be solved either exactly or by minimizing LSE given the correct illumination chromaticity Γ_c. The illumination chromaticity can be approximated by the image chromaticity of white reference taken by the camera or estimated from the current color images using a number of available methods (e.g. [29,11]).

3. Experiments

To demonstrate the effectiveness of our method, we conducted tests on both synthetic and real images. The synthetic images were

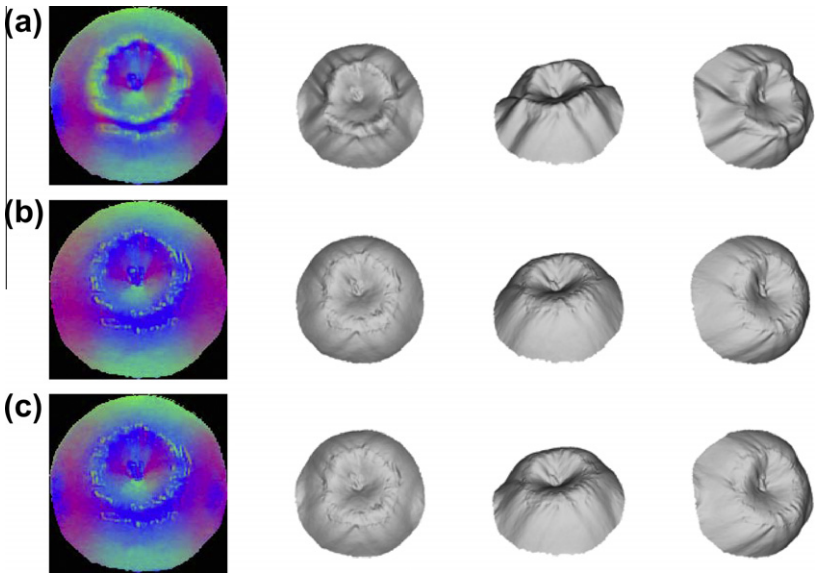

Fig. 7. Surface reconstruction. (a) Experimental results with conventional approach. (b) Results with the proposed approach. (c) Results with conventional approach using the estimated diffuse components as inputs. The first column shows the estimated surface normal (absolute values), and the last three columns show the corresponding shaded shapes from different viewing direction for visual comparison.

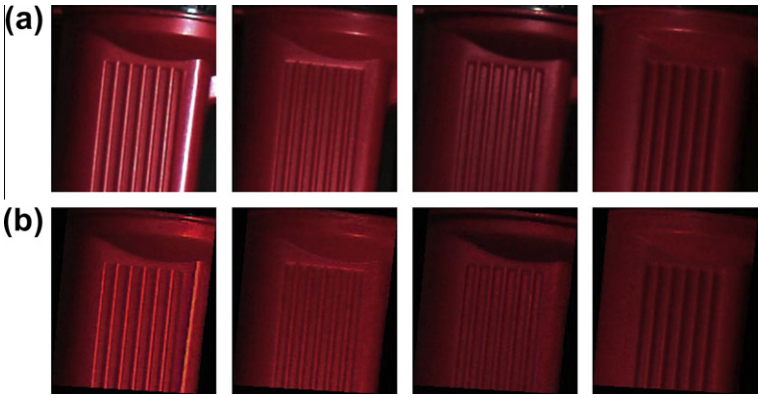

Fig. 8. The *flash light* data set. (a) Four out of the 12 input color images. (b) The corresponding diffuse components recovered using our approach.

generated with the Phong reflectance model [20] in which the specular term can be represented as

$$I^S = K^S(R \cdot V)^{\alpha}, \tag{19}$$

where K^S is the ratio of reflection of the specular term of incoming light, R is the direction that a perfectly reflected ray of light, V is the direction towards the viewer, and α is a shininess constant for the surface material, which decides how "evenly" light is reflected from a shiny spot. Note that the ground-truth values of the diffuse/specular component and the surface normal of these synthetic images are available for quantitative evaluation. The real images were captured under a single LED light source by a Sony DFW-X700 camera with gamma correction off. Using a mirror sphere and a diffuse sphere [5] placed next to the target object, we estimated the light source direction and intensity. Note that for the proposed reflection components separation method, lighting calibration is not required.

3.1. Synthetic data set

We show the effectiveness of our method with synthetic images by comparing our results with the ground truth. Fig. 1 displays results for a four-colored sphere. There is no visible error when the shininess constant given in Eq. (19) is set to 40. However, when the shininess constant decreases to 10, the area of the regions containing the pixels violating the specular-free observability assumption greatly increases, and errors become noticeable. But our method still successfully decreases the specular effect for those pixels. Fig. 2 depicts this improvement explicitly and Table 1. provides the numerical evaluation.

Fig. 3 shows the results of shape estimation from four synthetic images. The improvement is obvious, and the main differences of the recovered shape and the ground-truth appear around the edges of the sphere where the color images exhibit low intensity. The quantitative evaluation is provided in Table 2. The values in the table are the root-mean-square angular difference of the estimated surface normal and the ground truth in the specular regions. In our experiment, we treated a pixel to be specular if the root-mean-square difference of its color [0,255] and the ground-truth values of the diffuse component [0,255] is larger than 3. The synthetic images used in this experiment were generated with the Phong reflectance model by setting the shininess constant to be 40. To recover the shape, we chose an arbitrary height at an arbitrary point, and used the gradients (computed from the estimated

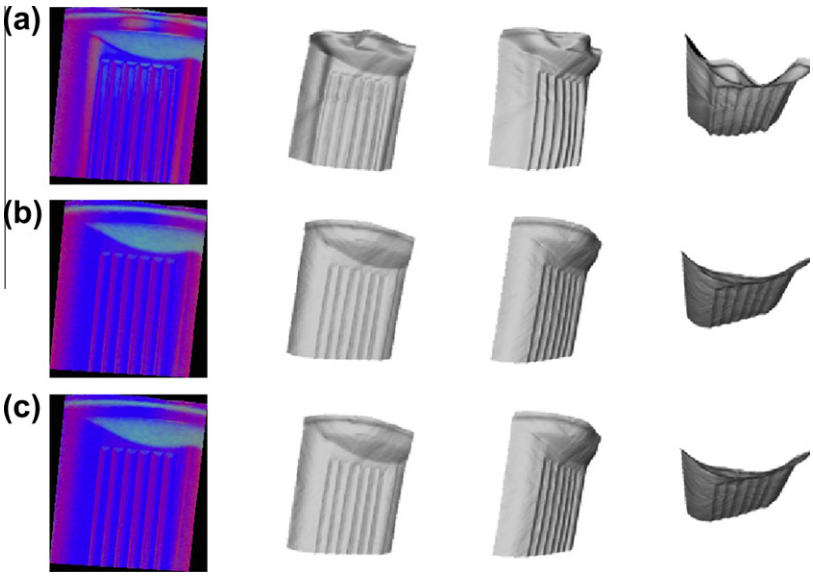

Fig. 9. Surface reconstruction. (a) Experimental results with conventional approach. (b) Results with the proposed approach. (c) Results with conventional approach using the recovered diffuse components. The first column shows the estimated surface normal (absolute values), and the last three columns show the corresponding shaded shapes from different viewing direction for visual comparison.

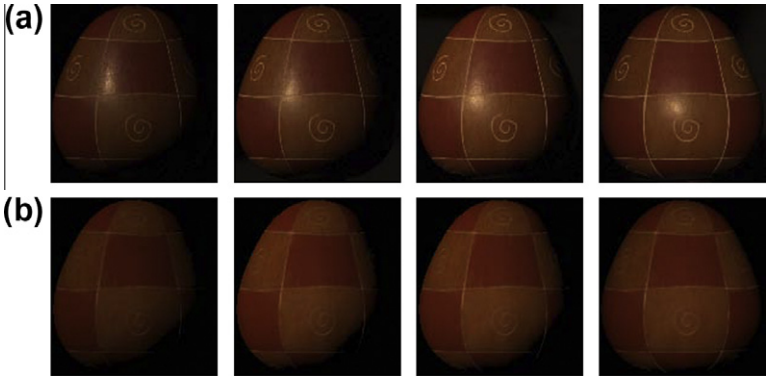

Fig. 10. The *gourd1* data set. (a) Four out of the 12 input color images. (b) The corresponding diffuse components recovered using our approach.

normal) to compute the height at a nearby point. We chose the simplest shape reconstruction method because our paper is focused on how to estimate the accurate surfaces normal for the specular surfaces, not accurate 3D surface reconstruction from the normal. The reconstructed shapes presented in the paper are mainly for visually evaluation of the estimated normal.

3.2. Real data sets

In this section, experimental results are presented for four real objects: the side face of a pear, the top of an apple and the side face of a flash light and side face of a gourd. 12 color images of each object captured under different illumination directions are used for PS.

Fig. 4 shows the pear images and the recovered diffuse components using the proposed approach. The visual improvement is obvious. Fig. 5 shows three sets of the shape estimation results of: conventional approach, proposed approach, and the conventional approach with the estimated diffuse components as input. Apparently, the main difference between the conventional approach and the proposed approach appears around the specular areas. The shape estimated with the conventional approach indicates that there is a sharp slope at the bottom of the pear. However, our approach does not exhibit this artifact as shown in Fig. 5b. Fig. 5c shows the results using the recovered diffuse components, which are similar to the results of our approach.

The experimental results on the top of an apple are presented in Figs. 6 and 7. The existence of the saturated pixels decreases the performance of our method. However, the improvements are still obvious. Using the conventional approach, the reconstructed shape is greatly distorted. In contrast, our approach suppresses the distortion except at the saturated pixels. Again, the results of our method agree with those of the conventional method with the recovered diffuse components. Figs. 8 and 9 present results for a plastic flash light. Similar to the *apple* data set, saturated pixels decreases the performance.

Lastly, we show results obtaining using high dynamic range (HDR) images in Figs. 10 and 11. The HDR images avoids saturation, thus correctly recover the diffuse reflections as demonstrated in Fig. 10. The first column in Fig. 11 presents the estimated surface normals obtained from conventional approach and our method, and the others are the shaded shapes under different light directions. Apparently, conventional approach fails around specular regions and introduces errors. Our method successfully corrects these errors.

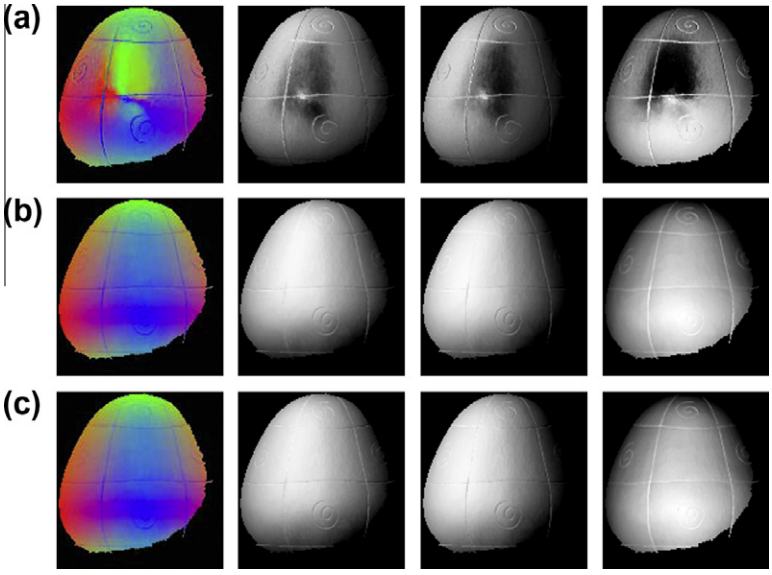

Fig. 11. Surface reconstruction. (a) Experimental results with conventional approach. (b) Results with the proposed approach. (c) Results with conventional approach using the recovered diffuse components. The first column shows the estimated surface normal (absolute values), and the last three columns show the corresponding shaded shapes under different light direction for visual comparison.

Table 3
Quantitative comparison of the surface normal of the proposed method and the conventional method using the estimated diffuse components. The numbers in the table are the root-mean-square angular difference of the reconstructed surface normal. If both the estimated surface normal and diffuse components are correct, the conventional method using the estimated diffuse components should produce the same surface normal as the proposed shape estimation method. This table demonstrates this claim.

	Avg diff.	Min diff.	Max diff.	Std diff.
Pear	0.0025	0.0000	0.4892	0.0112
Apple	0.0029	0.0000	0.4744	0.0169
Flash light	0.0029	0.0000	0.4486	0.0129
Gourd	0.0017	0.0000	0.3876	0.0049

Since no ground-truth values are available for real scenes, we only show the visual improvements. The experimental results of our method are consistent with those of the conventional shape estimation approach with the recovered diffuse components, which shows that both the shape estimation and the reflection component separation are reasonable. The quantitative comparison of the two methods is provided in Table 3.

4. Conclusions

In this paper, a new photometric stereo method of simultaneously estimating specular reflection and surface normal is presented. The proposed method can be used for all dielectric inhomogeneous surfaces. Compared with the conventional photometric stereo method, our extra requirement is illumination chromaticity which can be estimated by a white reference image or a number of existing methods. For reflection components separation, we make an additional assumption, namely that for every pixel, at lease one of the images is free of specularities. However, the proposed highlight removal method does not require lighting calibration. We demonstrate the robustness of our separation method by showing that even when the assumption is violated, our method will still reduce the specular highlights effectively.

References

[1] M. Ashikhmin, P. Shirley, An anisotropic Phong BRDF model, J. Graph. Tools 5 (2) (2000) 25–32.
[2] S. Barsky, M. Petrou, The 4-source photometric stereo technique for three-dimensional surfaces in the presence of highlights and shadows, PAMI 25 (10) (2003) 1239–1252.
[3] E.N. Coleman, R. Jain, Obtaining 3-dimensional shape of textured and specular surfaces using four-source photometry, Shape Recov. (1982) 180–199.
[4] R. Cook, K. Torrance, A reflectance model for computer graphics, Siggraph 15 (3) (1981) 307–316.
[5] D.B. Goldman, B. Curless, A. Hertzmann, S.M. Seitz, Shape and spatially-varying BRDFS from photometric stereo, in: ICCV, 2005, pp. 341–348.
[6] X. He, K. Torrance, F. Sillion, D. Greenberg, A comprehensive physical model for light reflection, Siggraph (1991) 175–186.
[7] A. Hertzmann, S.M. Seitz, Example-based photometric stereo: shape reconstruction with general, varying BRDFS, PAMI 27 (8) (2005) 1254–1264.
[8] K. Ikeuchi, Determining surface orientations of specular surfaces by using the photometric stereo method, Shape Recov. (1992) 268–276.
[9] G. Klinker, S. Shafer, T. Kanade, The measurement of highlights in color images, IJCV 2 (1) (1988) 7–32.
[10] E. Lafortune, S.-C. Foo, K.E. Torrance, D.P. Greenberg, Non-linear approximation of reflectance functions, Siggraph (1997) 117–126.
[11] H.-C. Lee, Method for computing the scene-illuminant chromaticity from specular highlights, Color (1992) 303–308.
[12] S. Lin, Y. Li, S.B. Kang, X. Tong, H.-Y. Shum, Diffuse-specular separation and depth recovery from image sequences, in: ECCV, 2002.
[13] S. Lin, and H.-Y. Shum, Separation of diffuse and specular reflection in color images, in: CVPR, 2001, pp. 341–346.
[14] S.P. Mallick, T.E. Zickler, D.J. Kriegman, P.N. Belhumeur, Beyond lambert: reconstructing specular surfaces using color, in: CVPR, 2005, pp. 619–626.
[15] S.K. Nayar, K. Ikeuchi, T. Kanade, Determining shape and reflectance of hybrid surfaces by photometric sampling, Shape Recov. (1992) 200–212.
[16] S.K. Nayar, X.S. Fang, T. Boult, Separation of reflection components using color and polarization, IJCV 21 (3) (1996).
[17] A. Ngan, F. Durand, W. Matusik, Experimental validation of analytical BRDF models, in: ACM SIGGRAPH 2004 Sketches, SIGGRAPH '04, 2004, p. 90.
[18] F.E. Nicodemus, Directional reflectance and emissivity of an opaque surface, Appl. Optics 7 (7) (1965) 767–773.
[19] B.T. Phong, Illumination for computer generated pictures, Commun. ACM 18 (6) (1975) 311–317.
[20] B.T. Phong, Illumination for computer generated pictures, Commun. ACM 18 (1975) 11–317.
[21] Y. Sato, K. Ikeuchi, Temporal-color space analysis of reflection, JOSA 11 (11) (1994) 2990–3002.
[22] Karsten Schlüns, Photometric stereo for non-lambertian surfaces using color information, in: CAIP, 1993, pp. 444–451.
[23] S.A. Shafer, Using color to separate reflection components, Color Res. App. 10 (4) (1985) 210–218.
[24] W. Silver, Determining Shape and Reflectance using Multiple Images, Master Thesis, MIT, 1980.
[25] H.D. Tagare, R. deFigueiredo, A theory of photometric stereo for a class of diffuse non-Lambertian surfaces, PAMI 13 (2) (1991) 133–152.
[26] P. Tan, S. Lin, L. Quan, H.-Y. Shum, Highlight removal by illumination-constrained inpainting, in: ICCV, 2003, p. 164.
[27] P. Tan, L. Quan, S. Lin, Separation of highlight reflections on textured surfaces, in: CVPR, 2006, pp. 1855–1860.

[28] Tan, R.T., Ikeuchi, K., Reflection components decomposition of textured surfaces using linear basis functions, in: CVPR, 2005, pp. 125–131.
[29] R.T. Tan, K. Nishino, K. Ikeuchi, Illumination chromaticity estimation using inverse-intensity chromaticity space, in: CVPR, 2003, pp. 673–680.
[30] R.T. Tan, K. Ikeuchi, Separating reflection components of textured surfaces using a single image, PAMI 27 (2) (2005) 178–193.
[31] K. Torrance, E. Sparrow, Theory for off-specular reflection from roughened surfaces, J. Opt. Soc. Am. 57 (9) (1967) 1105–1114.
[32] G.J. Ward, Measuring and modeling anisotropic reflection, Siggraph 26 (4) (1992) 265–272.
[33] Robert J. Woodham, Photometric Method for Determining Surface Orientation from Multiple Images, Shape from Shading, 1980, pp. 513–531.
[34] Robert J. Woodham, Gradient and curvature from the photometric stereo method, including local confidence estimation, JOSA 11 (1994) 3050–3068.