

Color Photometric Stereo Using a Rainbow Light for Non-Lambertian Multicolored Surfaces

Sejuti Rahman¹, Antony Lam², Imari Sato³, and Antonio Robles-Kelly⁴

¹Carnegie Mellon University, USA, ²Saitama University, Japan, ³National Institute of Informatics, Japan, ⁴National ICT Australia (NICTA)

¹sejutir@andrew.cmu.edu, ²antonylam@cv.ics.saitama-u.ac.jp,
³imarik@nii.ac.jp, ⁴antonio.robles-kelly@nicta.com.au

Abstract. This paper presents a novel approach for recovering the shape of non-Lambertian, multicolored objects using two input images. We show that a ring light source with complementary colored lights has the potential to be effectively utilized for this purpose. Under this lighting, the brightness of an object surface varies with respect to different reflections. Therefore, analyzing how brightness is modulated by illumination color gives us distinct cues to recover shape. Moreover, the use of complementary colored illumination enables the color photometric stereo to be applicable to multicolored surfaces. Here, we propose a color correction method based on the addition principle of complementary colors to remove the effect of illumination from the observed color. This allows the inclusion of surfaces with any number of chromaticities. Therefore, our method offers significant advantages over previous methods, which often assume constant object albedo and Lambertian reflectance. To the best of our knowledge, this is the first attempt to employ complementary colors on a ring light source to compute shape while considering both non-Lambertian reflection and spatially varying albedo. To show the efficacy of our method, we present results on synthetic and real world images and compare against photometric stereo methods elsewhere in the literature.

1 Introduction

Reconstructing 3D shape from multiple images has attracted ample attention from the computer vision community. Along these lines, one of the most popular approaches is photometric stereo (PS) [1], which reconstructs the object shape by varying the light direction across multiple images that share the same viewpoint. By assuming Lambertian reflectance and using calibrated lights, PS can recover the object shape and spatially varying albedo making use of three images.

On the other hand, color PS [2–4] augments conventional PS with the idea of multiplexing in the spectral domain. Rather than using three grayscale images, the method uses a single color image of a Lambertian surface being illuminated by three different colored lights. This allows for its application to the reconstruction of deformable surfaces [5, 6]. However, in general, color PS approaches [2–4,

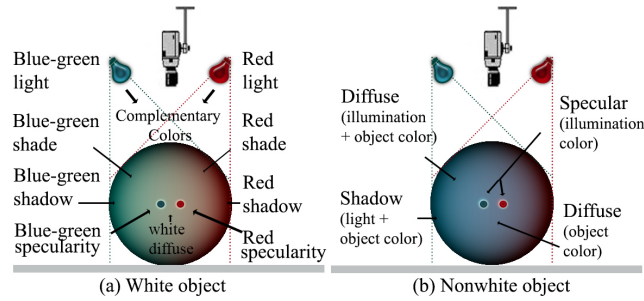


Fig. 1. Observed brightness under two complementary colored lights, (a) White colored object; (b) Non-white colored object.

7, 5] often assume constant object albedo and Lambertian reflectance. These assumptions can be overly restrictive in practice.

The aim of this paper is to extend color PS for non-Lambertian, multicolored surfaces. Here, we show that a ring light source with complementary colored lights has the potential to be effectively utilized for this purpose. The key observation is that, under a complementary colored light source, the observed color of an object surface varies with respect to different reflections (shown in Fig. 1). Therefore, analyzing how object color is modulated by illumination color gives us clues to estimate surface orientations. On top of that, the complementary colored illumination also allows the inclusion of multicolored surfaces in shape estimation. To deal with such surfaces, we propose a color correction method that exploits the addition principle of complementary colors to remove the effect of illumination from the observed color. This eases the estimation of object color and makes the method applicable to any number of chromaticities.

There have been a few approaches that consider scenes with varying chromaticity by either employing time multiplexing [8, 9], applying regularization to the normal field [10], using extra information provided by a depth camera [11], or a two camera stereo system [6]. However, our approach has significant advantages over the previous approaches as we can deal with shadows and specular highlights which are almost unavoidable phenomena in the real world.

It is worth noting in passing that our work has been partially motivated by the work in [12], where the same light configuration is utilized for depth edge extraction based on shadow cue. However, our approach differs from this work in that, we employ a colored ring light for 3D shape reconstruction, not for depth edge, and we utilize shading, shadow, and specular cues altogether. Our method is effective and requires only two input images taken under complementary colored lights. To the best of our knowledge, this is the first attempt that utilizes a complementary colored ring light for 3D reconstruction. The contributions of our paper are:

- We propose a new approach of color PS for non-Lambertian surfaces that utilizes the variation in observed surface color according to different reflections under complementary colored illumination.
- We present a color correction method based on the addition principle of complementary colors to render color PS applicable to multicolored surfaces.
- We show that the color difference between images under complementary illumination is useful for detecting diffuse, specular, and attached shadow reflections.

The rest of the paper is organized as follows. In Section 2, we summarize earlier research in PS for non-Lambertian, multicolored surfaces. In Section 3, we explain brightness variation under complementary illumination in different reflections, that we employ in Section 4 to develop theories to estimate shape. Experimental results and comparisons are presented in Section 5. Finally, in Section 6, we conclude on the developments presented here.

2 Related Work

Most of the surfaces around us are non-Lambertian. Conventional PS can not reliably reconstruct those, especially when deviation from the Lambertian assumption is large. Therefore, since the seminal work of Woodham [1], PS has received significant efforts to incorporate non-Lambertian phenomena. The existing approaches in this direction can be categorized into three broad classes.

The first approach is to model non-Lambertian surfaces using parametric models that are more complex than the Lambertian, e.g. Georghiadis [13] uses a simplified Torrance-Sparrow (TS) model [14] to account for specular highlights but discards shadows as outliers; Goldman et al. [15] employ the Ward model [16] to include both shadow and highlights, but require more user inputs.

The second group [17, 18] assumes that the non-Lambertian phenomena, such as shadows and highlights, are restricted to small regions of an image. As a result, these can be treated as outliers and hence, removed from further consideration.

The third and more recent approach is to exploit general properties rather than assuming any specific reflectance model, e.g., radiance similarity [19], attached shadow codes [20], monotonicity and isotropy in diffuse reflectance [21], isotropy and reciprocity [22–24], reflectance monotonicity [25], and so on. However, approaches in the latter two categories use a large number of observations. Our method belongs to the first category as we utilize the TS model to include non-Lambertian surfaces. However, our method is a color PS approach and our novelty is that we exploit the potential of complementary colored illumination to extend the applicability of PS to non-Lambertian cases.

Apart from the non-Lambertian cases, we also develop a theory to incorporate multicolored surfaces, which is often ignored in color PS. It is known that given an image of a surface of uniform chromaticity, illuminated by three spectrally and spatially separated light sources, it is possible to estimate a surface normal at each pixel [2]. However, this is no longer possible for surfaces with multiple

chromaticities as a change in pixel color could be caused by either a change in surface orientation or a change in chromaticity. Therefore, most of the color PS algorithms [2–4, 7, 26, 5] are applicable to uniform colored surfaces. However, there are also a few works that consider scenes with varying chromaticity.

In [6, 11], Anderson et al. assume a scene to be comprised of multiple piecewise constant chromaticity. To render the problem tractable, they segment the input image into regions of constant chromaticity before applying PS to estimate shape. To perform the segmentation, they require low resolution geometry of the scene provided by either a two-camera stereo system [6] or a depth camera [11].

In another approach, Janko et al. [10] apply color PS to estimate the shape of a dynamic and multicolored surface. To separate the contribution of surface orientation from that of albedo in observed color, the method exploits a constraint based on the temporal constancy in surface albedo. However, their method has two major drawbacks. Firstly, it enforces spatial smoothness on the surface normal over neighboring pixels, which, in turn, can cause over smoothing of fine shape details. Secondly, to estimate albedo, their method requires the surface to be dynamic enough and the input sequence to be long enough so that all normals turn towards each light source at least once.

Decker et al. [8] and Kim et al. [9] resort to time multiplexing to relax the assumption of constant chromaticity and apply color PS for dynamic scenes. The method in [9] accounts for time varying surface orientation and requires three images to recover shape at each frame. On the other hand, the method in [8] assumes surface normal to be constant across the frames and requires at least 2 input images to estimate shape. The method in [8], like ours, also utilizes complementary colored light to estimate object albedo.

However, all these color PS methods, including that in [8], assume Lambertian reflectance. In a related development, Hernández et al. [26] considers shadows by applying an integrability constraint, but limit their attention to monochromatic surfaces and disregards specular highlights.

Like ours, there have been several other PS approaches [22, 23, 27, 24], that utilize the ring light setup. However, all these methods assume monochromatic lights and require a large number of input images. For example, the work in [22] requires 20-30 images to recover only partial reconstruction of isotropic surfaces and [23] requires much more (about 100). Zhou and Tan [27] recover Euclidean structure for multicolored, Lambertian surfaces using at least 5 images and 2 views. Based on a ring light setup and reflectance symmetry, Tan et al. [24] also recover Euclidean structure by extending the partial reconstruction developed in [22]. However, their method depends on the algorithm in [22] to compute the iso-depth contours and, therefore, requires a large number of images. This contrasts with our method, which only requires two images captured from a single view.

3 Brightness under Complementary Colors

This section describes our observations on brightness under complementary illumination, that we will utilize later to develop our theory for recovering shape.

Under a ring of colored lights, the observed color of an object includes not only surface color but also illumination color. Analyzing how object color is modulated by illumination color gives us clues to estimate surface orientations. Let us first consider the simple case depicted in the left-hand panel of Fig. 1, where the target object has a neutral (white) color and is illuminated by two lights with complementary hues, red and blue-green.

This complementary light configuration provides significant information for each of the diffuse, specular, and attached shadow reflections:

- Diffuse: Depending on the surface orientation, brightness in diffuse pixels can be of two types, non-white shaded and white shaded. As shown in Fig. 1(a), when the surface normal is tilted towards one of the light sources, we see the shade of that particular light. On the other hand, when the surface normal is not biased towards any specific light source, both lights contribute the same. As the mixture of complementary colors in an additive color space results in white, these pixels appear white.
- Specular: It is well known that specular reflection is the mirror-like reflection of light observed when the direction of the incoming light and reflected light have equal angles with respect to the surface normal. In the case of a color ring light, the color of specular reflection depends on the surface orientation. For instance, in Fig. 1(a), we see red specular reflection at points where its mirror reflection is oriented towards the red light source.
- Shadow: Attached shadow appears in a pixel when the angle between the surface normal at that point and a light source is more than 90° . In this case, the surface does not receive any light from this light source and the pixel is colorized according to the contribution from the other light source with which the incident angle is less than 90° .

The discussion above can be extended to a continuous ring light that is formed as a maximally saturated complementary hue circle with same brightness. Under this light, we experience similar observations. As for diffuse reflection, the white shade appears when the surface normal is toward the camera and thus receives an equal amount of light from all lights on the ring. On the other hand, a nonwhite shade appears when the surface normal is tilted towards a specific light on the ring. In such cases, it is colored according to a weighted average of all light colors, where that particular light weights more than the rest.¹

As for specular reflection, each specular pixel exhibits the color of a specific light on the ring when its mirror reflection is oriented towards that light. As already seen in the case of two lights, a shadow pixel is colorized according to the contributions from the visible light sources.

Note, however, that the scenario becomes more complex when the target object has a non-white surface color. As shown in the right-hand panel of Fig. 1, the shading and shadow are now a product of both, the illuminant and object

¹ This follows Newton’s geometrical weighting [28], which states that the additive mixture of any number of colors is determined as the weighted average of the positions of the original colors on the hue-saturation plane.

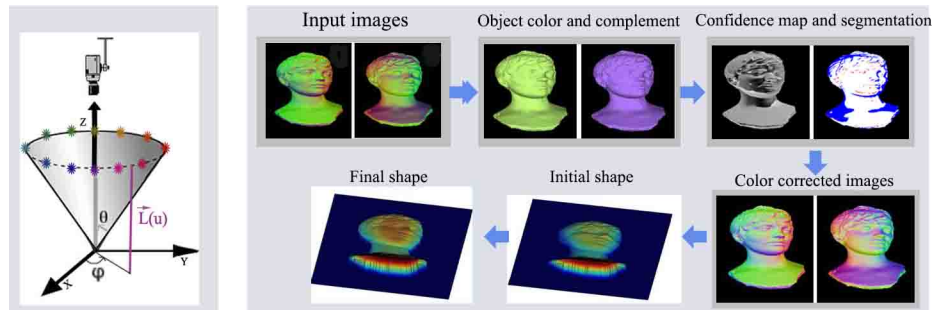


Fig. 2. Left-hand panel: illumination geometry for our colored ring light setup. Right-hand panel: a diagrammatic representation of our method.

color, whereas pure specular pixels show the illumination color. From the figure, we can understand that the separation of the illumination color from the object color becomes a challenging task. Later, in Section 4.1, we show how complementary illumination can be used to perform color correction. This delivers an image where the object color has been effectively canceled out.

4 Shape from Complementary Colors

We start from the assumption of directional illumination and orthographic projection. We consider a ring light setup where the observer is fixed at direction v . The global coordinate system is selected such that the Z -axis is aligned with the viewer direction, i.e., $v = (0, 0, 1)^T$. We assume that the scene is being illuminated by a ring light source, where the lights are arranged in equidistance on a cone centered at the viewing direction (see left-hand panel of Fig. 2). We further assume that our coordinate system is aligned with the HSV space and select colors on the ring such that the colors form a complementary hue circle. Our assumptions on scene properties are: no cast shadows in the scene, uniform roughness over the surface, and brightness distribution is Gaussian.

The right-hand panel of Fig. 2 shows a diagrammatic representation of our method, where our algorithm uses two input images and computes their sum to obtain the object color. Once the surface color is in hand, a pair of color-corrected images is recovered for purposes of color PS.

Next we consider the cases where objects are composed of multicolored diffuse surfaces in Section 4.1, specular surfaces in Section 4.2, and surfaces with both colored diffuse and specular reflections in Section 4.3.

4.1 Multicolored Diffuse Reflection

Let us assume that the scene is being illuminated by a ring-light source made up of k lights, i.e., $L = \{L_1, \dots, L_k\}$, each with a different hue from the set of hues, $H = \{H_1, \dots, H_k\}$. Now, if we capture an input image M_1 under H , the

color of each channel $c \in \{R, G, B\}$ for each diffuse pixel u , is in fact a linear combination of contribution from all the light sources, given as,

$$M_{1,c} = \sum_{i=1}^k (\mathbf{L}_i \cdot \mathbf{N}) \int \rho(\lambda) P_i(\lambda) Q_c(\lambda) d\lambda \quad (1)$$

where, the contribution of the i -th light source with power spectrum $P_i(\lambda)$ is integrated over the visible spectrum (380nm to 720nm). $Q_c(\lambda) = \{\bar{r}(\lambda), \bar{g}(\lambda), \bar{b}(\lambda)\}$ is the camera spectral sensitivity function. \mathbf{N} and $\rho(\lambda)$ denote surface normal and diffuse spectral reflectance at pixel u , respectively.

For the time being, let us assume that the surface is white and therefore has a constant spectral reflectance τ over all wavelengths. Following the scalar multiplication property and distributive law of dot products, we rewrite Eq. (1) as,

$$M_{1,c} = \tau \left(\sum_{i=1}^k (\mathbf{L}_i q_i(c)) \right) \cdot \mathbf{N} \quad (2)$$

where $q_i(c) = \int P_i(\lambda) Q_c(\lambda) d\lambda$. From Eq. (2), we can solve for surface normal as, $\mathbf{N} = A^{-1}b$, where $A = \sum_{i=1}^k \mathbf{L}_i q_i(c)$ and $b = [M_{1,c}]$. Thus our system of linear equations comprises of 3 equations to solve for 3 unknowns. Note that, we drop the term τ from our solution and assume that the surface normals are scaled by τ . This assumption is not exclusive to our method but rather common across other PS approaches [29].

To deal with nonwhite or multicolored surfaces, next we propose a color correction method that removes the contribution of object color to image brightness. For this, we need a second input image M_2 , that we capture under the complementary set of H , denoted by $\bar{H} = \{\bar{H}_1, \dots, \bar{H}_k\}$. The image brightness equation for M_2 is given by,

$$M_{2,c} = \sum_{i=1}^k (\mathbf{L}_i \cdot \mathbf{N}) \int \rho(\lambda) \bar{P}_i(\lambda) Q_c(\lambda) d\lambda \quad (3)$$

Note that, everything remains the same as in Eq. (1) except power spectrum of the i -th light $\bar{P}_i(\lambda)$ now corresponds to \bar{H}_i , i.e., the complementary hue of H_i .

Now if we add M_1 and M_2 , this results in an image showing only object color devoid of illumination color, as $(P_i(\lambda) + \bar{P}_i(\lambda))$ is equivalent to unity,²

$$M_{1,c} + M_{2,c} = \sum_{i=1}^k (\mathbf{L}_i \cdot \mathbf{N}) \left(\int \rho(\lambda) P_i(\lambda) Q_c(\lambda) d\lambda + \int \rho(\lambda) \bar{P}_i(\lambda) Q_c(\lambda) d\lambda \right)$$

² We note that in practice, $(P_i(\lambda) + \bar{P}_i(\lambda))$ do not always perfectly sum to 1 at each wavelength. Despite this, since we ultimately examine RGB values in our method and not individual wavelengths, we found the amount of error introduced did not adversely affect our algorithm.

$$\begin{aligned}
&= \sum_{i=1}^k (\mathbf{L}_i \cdot \mathbf{N}) \left(\int \rho(\lambda) (P_i(\lambda) + \bar{P}_i(\lambda)) Q_c(\lambda) d\lambda \right) \\
&= \sum_{i=1}^k (\mathbf{L}_i \cdot \mathbf{N}) \left(\int \rho(\lambda) Q_c(\lambda) d\lambda \right)
\end{aligned} \tag{4}$$

Now, if we assume that we have the complementary spectrum of $\rho(\lambda)$, we can generate a new image \bar{M}_1 whose equation in the RGB space is given by,

$$\bar{M}_{1,c} = \sum_{i=1}^k (\mathbf{L}_i \cdot \mathbf{N}) \int \bar{\rho}(\lambda) P_i(\lambda) Q_c(\lambda) d\lambda \tag{5}$$

where, in contrast to Eq. (1), ρ is now replaced by $\bar{\rho}$, which denotes the complementary spectrum of object color.

In practice, we use an RGB camera and do not have the complementary spectrum, so we utilize the HSV space to generate \bar{M}_1 . As the hue of $M_1 + M_2$ corresponds only to the object hue in the HSV space, we use this to compute the complementary hue of the object as, $\bar{obj}_{hue} = mod(obj_{hue} + 0.5, 1)$ where the object hue is normalized to $[0, 1]$. Then we generate an image $(\bar{obj})_{HSV}$, that shows only complementary object color, by changing the hue value of $(M_1 + M_2)_{HSV}$ from obj_{hue} to \bar{obj}_{hue} , keeping saturation and brightness unchanged. Now, adding $(M_1)_{RGB}$ to $(\bar{obj})_{RGB}$ results in only illumination color. Let's denote it by $(light)_{RGB}$. As \bar{M}_1 is composed of complementary object and illumination color (Eq. (5)), therefore, we get $(\bar{M}_1)_{RGB} = (\bar{obj})_{RGB} + (light)_{RGB}$.

Now adding \bar{M}_1 to M_1 , the term $(\rho(\lambda) + \bar{\rho}(\lambda))$ becomes unity, thus resulting in a brightness equation without the diffuse spectral reflectance term. This is given by,

$$\begin{aligned}
M'_{1,c} &= M_{1,c} + \bar{M}_{1,c} = \sum_{i=1}^k (\mathbf{L}_i \cdot \mathbf{N}) \left(\int (\rho(\lambda) + \bar{\rho}(\lambda)) P_i(\lambda) Q_c(\lambda) d\lambda \right) \\
&= \sum_{i=1}^k (\mathbf{L}_i \cdot \mathbf{N}) \left(\int P_i(\lambda) Q_c(\lambda) d\lambda \right)
\end{aligned} \tag{6}$$

We name the image M'_1 as the color corrected image of M_1 . Now with the color corrected image, we can use the solution presented earlier for white colored surfaces to multicolored surfaces without any loss of generality.

4.2 Specular Reflection

As described in Section 3, the observed color of a pure specular pixel corresponds to the illumination color. Therefore, we utilize the pixel's hue value to determine the illumination direction that causes the specular peak. As shown in the left-hand panel of Fig. 2, we define the illumination direction \mathbf{L} at pixel u using two angles: slant θ and tilt ϕ , where θ is the opening angle of the cone of the ring

light source. To compute tilt, we assume that our coordinate system is aligned with the HSV space, and use the hue value of u in HSV space to determine ϕ .

With the estimated illumination direction, we compute the surface normal using the half angle between the illumination and viewer direction. However, with the assumption of an idealized reflection direction, we miss off-specular reflections here. Therefore, we utilize the hue-based computations as initial shape estimates and update further using a coordinate descent [30] optimization scheme based on the simplified TS model [14]. According to this model, the image brightness equation for a specular pixel u can be expressed as,

$$M_{s,c} = \sum_{i=1}^k \left(\frac{1}{\cos \theta_s} \exp \left(-\frac{\theta_h^2}{2\sigma^2} \right) \right) \int P_i(\lambda) Q_c(\lambda) d\lambda \quad (7)$$

where, θ_s , θ_h , and σ denote reflection angle, half-angle, and surface roughness.

It is worth noting in passing that, in this section, we assume M_1 to be the observed specular-only reflection. Our goal is to fit this observed data to Eq. (7) using the least square minimization. We formulate our cost function in terms of two variables, surface normal and roughness.

At the first step, our algorithm utilizes the initial hue-based estimation of shape to estimate the surface roughness. We assume uniform surface roughness and employ the following cost function,

$$\sigma = \arg \min_{\sigma} \sum_{u \in M_1} \left[M_{1,c} - M_{s,c}(\mathbf{N}, \sigma) \right]^2 \quad (8)$$

We compute σ for each band separately and take the average if they are not the same. Once it estimates the optimal value for σ , the algorithm proceeds to the second step where the current estimate of σ is used to obtain surface normals. The cost function for \mathbf{N} is given as,

$$\mathbf{N} = \arg \min_{\mathbf{N}} \sum_c \left[M_{1,c} - M_{s,c}(\mathbf{N}, \sigma) \right]^2 \quad (9)$$

Thus the algorithm iterates between these two steps until none of the parameters change between two successive iterations or when a maximum number of iterations are completed.

4.3 A Unified Framework for Objects with both Diffuse and Specular Reflection

So far, we described our approach for diffuse-only and specular-only reflections. This section presents a unified approach that addresses surfaces with both diffuse and specular reflections.

We model the scene radiance as a linear combination of specular and diffuse reflectance based upon the TS reflectance model. Therefore, the image brightness at pixel u is given by,

$$M_{TS,c} = \sum_{i=1}^k \left[\max(0, (\mathbf{N} \cdot \mathbf{L}_i)) \int \rho(\lambda) P_i(\lambda) Q_c(\lambda) d\lambda \right. \\ \left. + K_s \frac{1}{\cos \theta_s} \exp\left(-\frac{\theta_h^2}{2\sigma^2}\right) \int P_i(\lambda) Q_c(\lambda) d\lambda \right] \quad (10)$$

Here, the first term on the right hand side denotes diffuse reflection where $\max(0, (\mathbf{N} \cdot \mathbf{L}_i))$ accounts for attached shadow, and the latter term denotes specular reflection with K_s as specular reflection coefficient at u .

From Eq. (10), we first show that our color correction process applies to surfaces with both diffuse and specular components. The image brightness equation for $M_1 + M_2$ is given as,

$$M_{1,c} + M_{2,c} = \sum_{i=1}^k \left[\max(0, (\mathbf{N} \cdot \mathbf{L}_i)) \int \rho(\lambda) (P_i(\lambda) + \bar{P}_i(\lambda)) Q_c(\lambda) d\lambda \right. \\ \left. + K_s \frac{1}{\cos \theta_s} \exp\left(-\frac{\theta_h^2}{2\sigma^2}\right) \int (P_i(\lambda) + \bar{P}_i(\lambda)) Q_c(\lambda) d\lambda \right] \\ = \sum_{i=1}^k \left[\max(0, (\mathbf{N} \cdot \mathbf{L}_i)) \int \rho(\lambda) Q_c(\lambda) d\lambda \right. \\ \left. + K_s \frac{1}{\cos \theta_s} \exp\left(-\frac{\theta_h^2}{2\sigma^2}\right) \int Q_c(\lambda) d\lambda \right] \quad (11)$$

As shown in Eq. (11), illumination colors are canceled for both diffuse and specular reflection. Therefore, all pixels, except the pure specular ones, now show only object color.

It is worth noting that the specular-only pixels become white in $M_1 + M_2$, which is an important observation as this relates to the accuracy in estimating object color from $M_1 + M_2$. In this context, we would like to mention that this causes noticeable error in shape estimation at specular pixels in the color PS method in [8], that also utilizes complementary colored illumination to estimate object albedo. However, this does not affect our method as we deal with object color using our color correction method.

From the brightness equation of $M_1 + \bar{M}_1$, we see that our color correction process is applicable to surfaces with both reflection components,

$$M'_{1,c} = M_{1,c} + \bar{M}_{1,c} = \sum_{i=1}^k \left[\max(0, (\mathbf{N} \cdot \mathbf{L}_i)) \int (\rho(\lambda) + \bar{\rho}(\lambda)) P_i(\lambda) Q_c(\lambda) d\lambda \right. \\ \left. + 2K_s \frac{1}{\cos \theta_s} \exp\left(-\frac{\theta_h^2}{2\sigma^2}\right) \int P_i(\lambda) Q_c(\lambda) d\lambda \right] \quad (12)$$

As the albedo term is exclusive to the diffuse part, adding \bar{M}_1 to M_1 , removes the object color from the diffuse part, while leaving the illumination color unchanged in the specular part.

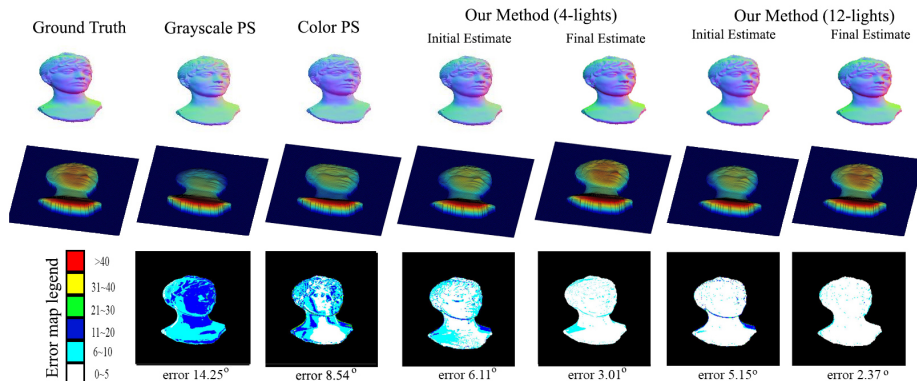


Fig. 3. Comparative results on 3D reconstruction for a synthetic shape. We present the RGB-encoded normal map and depth map for the ground truth data and that yielded by grayscale PS [1], color PS [8], our method with 4-lights and 12-lights, respectively. The bottom row shows error maps and average per-pixel angular error measures (in degrees) for the surface normals yielded by each of the methods under consideration as compared to the ground truth.

Next, we pose the shape estimation problem as an iterative minimization and fit the observed data to Eq. (10) to recover the parameters \mathbf{N} , σ , and K_s .

Initialization For initialization, we first segment image pixels into diffuse, specular, or shadows on the assumption that either they are purely diffuse, specular, or shadows based on the color difference between M_1 and M_2 . Then we utilize the approaches in Section 4.1 and 4.2 to initialize surface normals in diffuse-only and specular-only pixels.

For shadow pixels, we employ a hue-based computation. The surface orientation at a shadow pixel u is defined using two angles: slant θ and tilt ϕ , where θ comes from the opening angle of the cone of the ring light source and ϕ is set according to the hue of u in the HSV space.

For segmentation, we compute the Euclidean distance between the images in the hue-saturation plane and name this the shadow-specularity confidence map. The distance has higher values in specular and shadow pixels and lower values in diffuse. This conforms to our observations in Section 3, where we see that the color difference between two input images are higher in these reflections than the diffuse. Moreover, the distribution of this distance is bimodal and, hence, we can apply Otsu’s adaptive threshold method [31] to select D_{th} , a threshold on this distance, that can be employed to segment the specular and shadow pixels from diffuse pixels. Finally, we apply a threshold on brightness, V_{th} to separate shadows from specular highlights. Since we have considered the distribution of brightness to be Gaussian and specular pixels to occupy a minor portion of the scene, V_{th} should be on the right bottom rim of the bell curve for the distribution

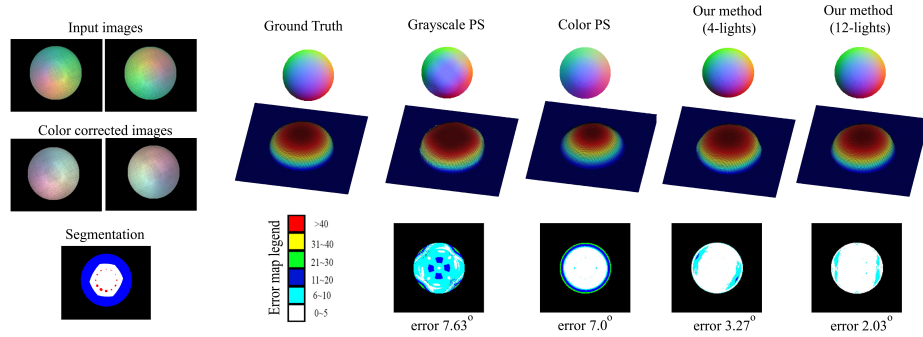


Fig. 4. The left-hand panel shows input images, color corrected images, and segmentation for a real-world sphere. The right-hand panel shows comparative results for 3D reconstruction. The top row presents the RGB-encoded normal maps and the middle row presents depth maps for the ground truth data (left) and that yielded, in turn, by grayscale PS [1], color PS [8], our method with 4 lights, and our method with 12 lights. The bottom row shows error maps and average per-pixel angular error (in degrees) for the surface normals yielded by each of the methods (compared to the ground truth).

and can be computed as, $V_{th} = mean(V) + 1.5 * std(V)$, where $mean(\cdot)$ and $std(\cdot)$ denote the mean and standard deviation, respectively.

Iterative Updates The coordinate descent approach comprises of three interleaved minimization steps. We write the cost functions using the TS model given in Eq. (10). Note that Eq. (10) shows the general form of the TS model, whereas the model we fit in assumes all $\rho(\lambda) = 1$, since our method allows for the ρ terms in Eq. (12) to sum to unity. Also, Eq. (12) shows $2K_s$ when Eq. (10) only shows K_s . In Eq. (13-15), this scalar difference is accounted for during optimization.

At the first step, the algorithm solves for surface roughness over all pixels by using the initial estimates of shape as,

$$\sigma = \arg \min_{\sigma} \sum_{u \in M'_1} \left[M'_{1,c} - M_{TS,c}(K_s, \mathbf{N}, \sigma) \right]^2 \quad (13)$$

At the second step, with the current estimates of σ and \mathbf{N} , it minimizes the following cost function to obtain K_s at each pixel,

$$K_s = \arg \min_{K_s} \sum_c \left[M'_{1,c} - M_{TS,c}(K_s, \mathbf{N}, \sigma) \right]^2 \quad (14)$$

At the third step, it solves for \mathbf{N} using the updated values of K_s and σ , by minimizing the following cost function,

$$\mathbf{N} = \arg \min_{\mathbf{N}} \sum_c \left[M'_{1,c} - M_{TS,c}(K_s, \mathbf{N}, \sigma) \right]^2 \quad (15)$$

Thus once the former variables are at hand, we obtain optimal values for the latter ones. The algorithm iterates between these three steps, until convergence.

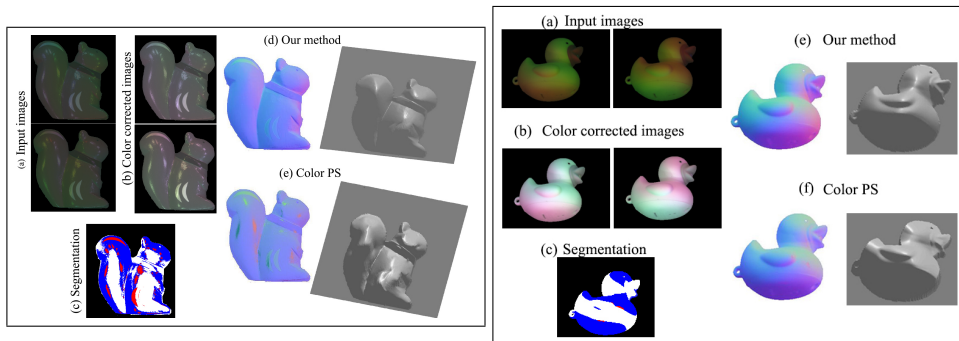


Fig. 5. The left-hand and right-hand panels show 3D reconstruction results for two real world objects, a squirrel and a duck, respectively. In each panel: (a) input images, (b) color corrected images, (c) segmentation, (d) and (e) present RGB-encoded normal maps and depth maps delivered by our method and color PS [8], respectively.

5 Experimental Results

We now illustrate the accuracy of our method for the purposes of 3D reconstruction using synthetic and real world data. To evaluate our method, we compare our results against ground truth data and two other PS approaches. As our method requires only two input images, we compare our results with those obtained by the methods that require a small number of images. The first of these is the PS method in [1], for which we use four images. To obtain these images, we turn on each of the 4 evenly distributed lights on the ring one by one.

The other of our alternatives is a closely related work, i.e., the color PS method in [8], that uses complementary illumination for estimating surface albedo to relax the assumption of constant object chromaticity. This method does not work for non-Lambertian surfaces and requires $(k+2)/3$ images over k lights. This contrasts with our method, which can be applied to non-Lambertian objects and to any number of light sources greater than or equal to two.

In fact, the accuracy of our estimated shape increases with the number of light sources on the ring. We illustrate this behavior in Fig. 3, where we show results for a synthetic shape for four and twelve lights. For both $k=4$ and $k=12$ cases, we evenly distribute the colored lights on the ring. The human face in Fig. 3 has been rendered (shown in Fig. 2) using the reflectance model in Eq. (10) and the albedo of a Aloe vera leaf obtained in house using a spectrometer. For the microfacet slope values, we have used a normal distribution with a mean of 2.0 and a standard deviation of 0.01.

In the figure, we also show the RGB encoded normal map and depth map delivered by our method and the two alternatives, i.e., the grayscale PS [1] and the color PS [8]. We include initial and final estimates of the surface shape to provide a better understanding of our optimization scheme. The bottom row shows error maps and average per-pixel angular error measures (in degrees) for the surface normals yielded by each method compared against the ground truth.

Note that the depth map is recovered by applying the Frankott and Chellappa integration method [32] on the surface normals yielded by each of the methods. We can see that our surface normals are in good accordance with the shape of the face even at initial value. Further, after optimization, improvements are noticeable specifically in specular and shadowed regions. This is confirmed by our error rates, which are typically less or around 3° .

For our real world imagery, we have used three non-Lambertian objects. These are a glossy wooden sphere (Fig. 4), a squirrel figurine with a very specular surface (Fig. 5), and a plastic duck (Fig. 5). All the objects are multicoloured and include specular highlights and shadows.

For each object, we show input images, color corrected images, segmentation results, and reconstructed shape. In Fig. 4, we also show the ground truth data and provide error maps for surface normals estimated by our method and that by two other alternative methods. Our method outperforms the alternatives, which is evident in the qualitative results, i.e., normals maps and depth maps, and also in quantitative error measures.

Next we present shape reconstruction results for the squirrel and the duck in Fig. 5. As mentioned earlier, our method can take input images acquired using any number of illuminants. However, to make our method comparable to that in [8], in the following we only employ two input images with 4 light sources. For each object, we compare the RGB encoded normal map and depth map yielded by our method with that delivered by the color PS [8]. From the figures we can see that the color PS fails to estimate the correct shape in specular and shadowed pixels. This contrasts with our method, which delivers better surface detail, and, in turn, more accurate shapes.

6 Conclusion

In this paper, we have presented a novel method for recovering the 3D shape of a non-Lambertian, multicolored object by utilizing a ring light source. The underlying theory for our method is based on the properties of complementary colors and the brightness variations in diffuse, specular, and shadow reflections. This allows for the computation of a pair of color corrected images which can be used for purposes of color PS making use of an iterative optimization scheme. As far as we know, this is the first attempt to utilize colored illuminants on a ring to recover object shape. Moreover, our method can naturally process non-Lambertian, multicolored surfaces with unknown reflectance properties using only two input images. A future research direction would be to extend the theory for deformable surfaces and uncalibrated PS.

Acknowledgement. This research was supported in part by the Ministry of Education, Science, Sports and Culture Grant-in-Aid for Scientific Research on Innovative Areas.

References

1. Woodham, R.J.: Photometric method for determining surface orientation from multiple images. *Optical Engineering* **19** (1980) 139–144
2. Drew, M., Kontsevich, L.: Closed-form attitude determination under spectrally varying illumination. In: *IEEE Conference on Computer Vision and Pattern Recognition*. (1994)
3. Kontsevich, L., Petrov, A., Vergelskaya, I.: Reconstruction of shape from shading in color images. *Journal of Optical Society of America A* **11** (1994) 1047–1052
4. Woodham, R.J.: Gradient and curvature from photometric stereo including local confidence estimation. *Journal of Optical Society of America A* **11** (1994) 3050–3068
5. Brostow, G., Hernández, C., Vogiatzis, G., Stenger, B., Cipolla, R.: Video normals from colored lights. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **33** (2011) 2104–2114
6. Anderson, R., Stenger, B., Cipolla, R.: Color photometric stereo for multicolored surfaces. In: *IEEE International Conference on Computer Vision*. (2011)
7. Hernández, C., Vogiatzis, G., Brostow, G.J., Stenger, B., Cipolla, R.: Non-rigid photometric stereo with colored lights. In: *IEEE International Conference on Computer Vision*. (2007)
8. Decker, B.D., Kautz, J., Mertens, T., Bekaert, P.: Capturing multiple illumination conditions using time and color multiplexing. In: *IEEE Conference on Computer Vision and Pattern Recognition*. (2009)
9. Kim, H., Wilburn, B., Ben-Ezra, M.: Photometric stereo for dynamic surface orientations. In: *European Conference on Computer Vision*. (2010)
10. Janko, Z., Delaunoy, A., Prados, E.: Colour dynamic photometric stereo for textured surfaces. In: *Asian Conference on Computer Vision*. (2010)
11. Anderson, R., Stenger, B., Cipolla, R.: Augmenting depth camera output using photometric stereo. In: *MVA*. (2011)
12. Taguchi, Y.: Rainbow flash camera: Depth edge extraction using complementary colors. In: *European Conference on Computer Vision*. (2012)
13. Georghiades, A.: Incorporating the Torrance and Sparrow model of reflectance in uncalibrated photometric stereo. In: *IEEE International Conference on Computer Vision*. (2003)
14. Torrance, K., Sparrow, E.: Theory for off-specular reflection from roughened surfaces. *Journal of Optical Society of America A* **57** (1967) 1105–1114
15. Goldman, D., Curless, B., Hertzmann, A., Seitz, S.: Shape and spatially-varying BRDFs from photometric stereo. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **32** (2010) 1060–1071
16. Ward, G.: Measuring and modeling anisotropic reflection. *Computer Graphics* **26** (1992) 265–272
17. Barsky, S., Petrou, M.: The 4-source photometric stereo technique for three dimensional surfaces in the presence of highlights and shadows. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **25** (2003) 1239–1252
18. Wu, L., Ganesh, A., Shi, B., Matsushita, Y., Wang, Y., Ma, Y.: Robust photometric stereo via low-rank matrix completion and recovery. In: *Asian Conference on Computer Vision*. (2011)
19. Sato, I., Okabe, T., Yu, Q., Sato, Y.: Shape reconstruction based on similarity in radiance changes under varying illumination. In: *IEEE International Conference on Computer Vision*. (2007)

20. Okabe, T., Sato, I., Sato, Y.: Attached shadow coding: Estimating surface normals from shadows under unknown reflectance and lighting conditions. In: IEEE International Conference on Computer Vision. (2009)
21. Higo, T., Matsushita, Y., Ikeuchi, K.: Consensus photometric stereo. In: IEEE Conference on Computer Vision and Pattern Recognition. (2010)
22. Alldrin, N., Kriegman, D.: Toward reconstructing surfaces with arbitrary isotropic reflectance : A stratified photometric stereo approach. In: IEEE International Conference on Computer Vision. (2007)
23. Alldrin, N., Zickler, T., Kriegman, D.: Photometric stereo with non-parametric and spatially-varying reflectance. In: IEEE Conference on Computer Vision and Pattern Recognition. (2008)
24. Tan, P., Quan, L., Zickler, T.: The geometry of reflectance symmetries. IEEE Transactions on Pattern Analysis and Machine Intelligence **33** (2011) 2506–2520
25. Shi, B., Tan, P., Matsushita, Y., Ikeuchi, K.: Elevation angle from reflectance monotonicity: Photometric stereo for general isotropic reflectances. In: European Conference on Computer Vision. (2012)
26. Hernández, C., Vogiatzis, G., Cipolla, R.: Shadows in three-source photometric stereo. In: European Conference on Computer Vision. (2008)
27. Zhou, Z., Tan, P.: Ring-light photometric stereo. In: European Conference on Computer Vision. (2010)
28. MacEvoy, B.: Color vision. <http://www.handprint.com/LS/CVS/color.html> (2008)
29. Higo, T., Matsushita, Y., Joshi, N., Ikeuchi, K.: A hand-held photometric stereo camera for 3-d modeling. In: IEEE International Conference on Computer Vision. (2009)
30. Friedman, J., Hastie, T., Hfling, H., Tibshirani, R.: Pathwise coordinate optimization. The Annals of Applied Statistics **1** (2007) 302–332
31. Otsu, N.: A threshold selection method from gray-level histogram. IEEE Trans. Systems Man, and Cybernetics **9** (1979) 62–66
32. Frankot, R., Chellappa, R.: A method for enforcing integrability in shape from shading algorithms. IEEE Transactions on Pattern Analysis and Machine Intelligence **10** (1988) 439–451