

Beyond Lambert: Reconstructing Specular Surfaces Using Color

Satya P. Mallick[†] Todd E. Zickler[‡] David J. Kriegman[†] Peter N. Belhumeur[¶]
spmallick@vision.ucsd.edu zickler@eecs.harvard.edu kriegman@cs.ucsd.edu belhumeur@cs.columbia.edu

[†]Computer Science and Engineering, University of California at San Diego, La Jolla, CA 92093

[‡]Engineering and Applied Sciences, Harvard University, Cambridge, MA 02138

[¶]Computer Science, Columbia University, New York, NY 10027

Abstract

We present a photometric stereo method for non-diffuse materials that does not require an explicit reflectance model or reference object. By computing a data-dependent rotation of RGB color space, we show that the specular reflection effects can be separated from the much simpler, diffuse (approximately Lambertian) reflection effects for surfaces that can be modeled with dichromatic reflectance. Images in this transformed color space are used to obtain photometric reconstructions that are independent of the specular reflectance. In contrast to other methods for highlight removal based on dichromatic color separation (e.g., color histogram analysis and/or polarization), we do not explicitly recover the specular and diffuse components of an image. Instead, we simply find a transformation of color space that yields more direct access to shape information. The method is purely local and is able to handle surfaces with arbitrary texture.

1. Introduction

Scene analysis is complicated by the presence of highlights due to specular reflection, since these highlights are a complex function of viewpoint, illumination and surface microstructure. Typically, in order to obtain meaningful information about a scene from its images, complex reflection effects such as highlights are either roughly approximated or completely ignored by vision systems. For example, most stereo and structure-from-motion techniques ignore specular highlights altogether, instead modeling surfaces as Lambertian; and even state-of-the-art recognition systems are based on relatively simple (e.g., spatially-invariant, parametric) models of reflectance. (See, Ref. [3], for example.)

Even when the illumination of a scene can be controlled, the coupling of reflectance and shape makes the recovery of 3D shape a difficult task. This is the case for photometric stereo methods, for example, where images are acquired from a fixed viewpoint under multiple, known illumination conditions. Many photometric stereo techniques assume that surfaces are Lambertian [29], and others assume the




Figure 1. Removal of specular reflectance effects. A rotation of RGB color space (a) provides an image that is void of specular reflection effects (b) and provides a more direct measurement of scene shape.

reflectance is given *a priori* by a reference object [23], a linear basis of reference objects [9], or by an analytic BRDF model [10, 18, 24]. When these reflectance assumptions are not satisfied, the accuracy of the recovered shape can be compromised.

Recently, a small number of surface reconstruction techniques have overcome these limitations by effectively decoupling shape and reflectance in images. These techniques rely on the careful acquisition of images to exploit physical properties such as reflectance isotropy [14], reciprocity [15, 30] and the constancy of radiance in free space [12, 15]. By reducing or eliminating the restrictions on surface reflectance, these techniques provide accurate reconstructions for a much broader class of surfaces.

In this paper, we present a new method that is similar to these techniques in that it recovers estimates of surface shape that are independent of specular reflectance; but instead of requiring a complex acquisition system, the method is quite easy to implement. It is based on photometric stereo, and like existing photometric stereo techniques, it can recover shape from as few as three images captured under three (possibly unknown) illumination directions. We achieve independence from specular reflectance using a data-dependent rotation of RGB color space. This color space transformation allows us to distill the image, removing the specular effects, and leaving only the much

simpler—ideally Lambertian—diffuse effects. (See Fig. 1.) The technique can be applied to any *dichromatic surface* (i.e., surface whose reflectance can be represented using Shafer’s dichromatic model [22]) for which the body spectral reflectance varies significantly over the visible spectrum (i.e., not ‘white’.)

This paper begins with a brief review of the dichromatic reflectance model and previous work related to highlight removal and photometric stereo. Then, it introduces a transformation of color space as a means of removing specular reflection effects, and applies this color transformation to the problem of photometric stereo. We show that by eliminating specular effects, this transformation enables the application of well-known, Lambertian photometric stereo techniques to non-Lambertian surfaces.

2. Background and Related Work

At an appropriate scale, reflectance at a surface point is described by the bi-directional reflectance distribution function, or BRDF. Here, we consider it to be a five-dimensional function of wavelength and imaging geometry, and we write it as $f(\lambda, \boldsymbol{\theta})$, where $\boldsymbol{\theta} = (\theta_i, \phi_i, \theta_r, \phi_r)$ encodes the directions of the incident and reflected radiance in the local coordinate system.

2.1. The Dichromatic Model

The dichromatic model of reflectance is a common special case of the BRDF model, and it was originally developed by Shafer [22] to model dielectrics. It assumes that the BRDF of the surface can be decomposed into two additive components: the interface (specular) reflectance and the body (diffuse) reflectance. Furthermore, it assumes that each of these two components can be factored into a univariate function of wavelength and a multivariate function that depends on the imaging geometry. That is,

$$f(\lambda, \boldsymbol{\theta}) = g_d(\lambda)f_d(\boldsymbol{\theta}) + g_s(\lambda)\tilde{f}_s(\boldsymbol{\theta}).$$

Finally, the model assumes that the index of refraction on the surface is constant over the visible spectrum—a valid assumption for many materials—so that $g_s(\lambda)$ is a constant function. This leads to the common expression for the BRDF of a dichromatic surface,

$$f(\lambda, \boldsymbol{\theta}) = g_d(\lambda)f_d(\boldsymbol{\theta}) + f_s(\boldsymbol{\theta}), \quad (1)$$

where $f_s(\boldsymbol{\theta}) = g_s\tilde{f}_s(\boldsymbol{\theta})$. The function $g_d(\lambda)$ is often referred to as the *spectral reflectance* and is an intrinsic property of the material.

Even though it was originally used to describe the reflectance of dielectrics [22], the dichromatic model has been used successfully as an approximation of the reflectance of many different materials. For example, although the reflectance of human skin is more accurately described

by a higher dimensional function (the 8-dimensional BSS-RDF [27]), the BRDF approximation—and more specifically, the dichromatic BRDF model—has proven useful for a number of applications, including face recognition [3, 6] and pigment-based image analysis and synthesis [26]. The validity of this approximation for skin reflectance is also supported the spectrophotometry measurements of Marszalec et al. [16].

In order to derive an expression for the image of a dichromatic surface, consider an observed surface point \mathbf{x} illuminated from direction $\hat{\mathbf{i}}$, and let $\boldsymbol{\theta} = (\theta_i, \phi_i, \theta_r, \phi_r)$ represent this direction and the viewing direction in the local coordinate system. We assume that the sensor is a linear device with sensitivity function $C_k(\lambda)$, and that the spectral power distribution (SPD) of the light source is given by $L(\lambda)$. Under these conditions, assuming that the BRDF at \mathbf{x} is given by Eq. 1, the sensor response is

$$I_k = (D_k f_d(\boldsymbol{\theta}) + S_k f_s(\boldsymbol{\theta})) \hat{\mathbf{n}} \cdot \hat{\mathbf{i}}, \quad (2)$$

with

$$D_k = \int C_k(\lambda)L(\lambda)g_d(\lambda)d\lambda$$

$$S_k = \int C_k(\lambda)L(\lambda)d\lambda.$$

An RGB color vector $\mathbf{I} = [I_R, I_G, I_B]^\top$ from a typical camera consists of three such measurements, each with a different sensor response with support in the visible spectrum.

Note that S_k represents the effective source strength as measured by the k^{th} sensor and is independent of the surface being observed. For notational simplicity, we define $\mathbf{S} = [S_R, S_G, S_B]^\top$ (with a corresponding definition for \mathbf{D}), and since scale can be absorbed by f_d and f_s , we assume $\|\mathbf{D}\| = \|\mathbf{S}\| = 1$.

2.2. Highlight Removal

As made clear by Eq. 2, in RGB color space a collection of color vectors from a dichromatic material under multiple view and illumination configurations (i.e., different values of $\boldsymbol{\theta}$) lie in the *dichromatic plane*—the plane spanned by the effective source and body colors, \mathbf{S} and \mathbf{D} [22]. In addition, it has been observed that these color vectors often cluster in the shape of a ‘skewed-T’ in this plane, where the two limbs of the skewed-T correspond to diffuse and specular reflection [7, 11]. When these limbs are sufficiently distinct, the diffuse and source colors can be recovered, the two components can be separated, and the highlights can be removed [11].

While this method works well for homogeneous, dichromatic surfaces in the noiseless case, there are three significant limitations that make it difficult to use in practice. First, many surfaces are textured and violate the homogeneous assumption. Even when an image *does* contain homogeneous

surfaces, a non-trivial segmentation process is required to identify them. Second, in order for the specular and diffuse limbs of the skewed-T to be distinct, the specular lobe must be sufficiently narrow (i.e., its angular support must be small relative to the curvature of the surface.) Finally, when the diffuse and specular colors are the same, there is no way to distinguish between the two components, and no color separation is possible.

Some of these restrictions can be overcome by using additional cues such as polarization [28]. Nayar et al. [17] use multiple exposures with varying polarization to estimate the source color independently at each point, and they exploit local spatial coherence to separate the diffuse and specular components. This method offers significant improvement over the use of only one of color or polarization information, but it still requires that the specular lobe is sufficiently narrow, and that the diffuse color is locally constant.

2.3. Photometric Stereo

Photometric stereo is the process of recovering 3D shape from a series of images captured from fixed viewpoint under multiple illuminations. Since they rely on the inversion of the image formation process, as described in Sect. 1, photometric stereo methods typically require significant knowledge about the reflectance of surfaces in the scene.

Coleman and Jain [4] were perhaps the first to present a photometric technique for reconstructing non-Lambertian surfaces without an explicit reflectance model. In their method, the BRDF is assumed to be a linear combination of a Lambertian diffuse component and an undefined specular component with limited angular support. When four point-source illuminations are available, specular measurements can be treated as outliers and discarded, provided that the illumination directions are far from one another relative to the angular extent of the specular lobe. (This ensures that the specular reflectance component is zero for three of the four observations of each surface point.) Barsky and Petrou [2] refine this technique by using color information to improve the detection of specular measurements. Like the original work, however, specular measurements are treated as outliers, and the specular component is assumed to have limited angular support.

Another approach to photometric stereo for non-Lambertian surfaces is to assume dichromatic surfaces, and to remove highlights as a pre-processing step as discussed in Sect. 2.2. This is the approach taken by Schlüns and Wittig [21], who assume homogeneous dichromatic surfaces, and separate the diffuse and specular components using color histogram analysis techniques similar to Klinker et al. [11]. Sato and Ikeuchi [20] take a similar approach, but avoid the restriction to homogeneous surfaces by using a large number of light source directions to compute a distinct color histogram at each point. Because these meth-




Figure 2. Linear and non-linear transformations of RGB color space. Three observations of the same material yield color vectors $\mathbf{I}_1 \dots \mathbf{I}_3$ in the dichromatic plane spanned by the source and diffuse colors \mathbf{S} and \mathbf{D} . Left: The proposed SUV color space is a rotation of the RGB space. One axis is aligned with the source color, and two of three transformed channels (UV) are independent of specular reflectance. Diffuse shading information is preserved and can be used to recover shape. Right: Central projection used to compute r - g chromaticity values and HSV-type color spaces does not preserve shading information.

ods explicitly recover the diffuse and specular components, they have the additional benefit of providing an estimate of the diffuse color \mathbf{D} at each point in addition to recovering the surface shape. Since they are based on conventional dichromatic color separation, however, they are subject to the restrictions discussed in the previous section. Most importantly, they assume that the specular lobe is narrow relative to the surface curvature, an assumption similar to that underlying the four-source method of Coleman and Jain [4].

3. A Data-dependent Color Space

As described in the previous section, many photometric stereo methods assume the BRDF is Lambertian or known in some form. The methods that do not make this assumption require dichromatic surfaces without texture [21] or Lambertian-plus-specular surfaces with a narrow specular lobe [2, 4]. In this section we introduce a transformation of color space that seeks to overcome both of these limitations.

Suppose we linearly transform RGB color space by rotating the coordinate axes, and as shown in the left of Fig. 2, suppose this rotation is such that one of the axes (red, say) becomes aligned with the direction of the effective source color \mathbf{S} . This transformation defines a new color space, which we refer to as the SUV color space. It can be defined according to $\mathbf{I}_{SUV} = [R]\mathbf{I}$ using any $[R] \in SO(3)$ that satisfies $[R]\mathbf{S} = [1, 0, 0]^T$. Here, we choose $[R] = [R_G(-\theta_S)][R_B(\phi_S)]$ where $[R_k(\theta)]$ is a right-handed rotation about the k -axis by angle θ , and (θ_S, ϕ_S) are the elevation and azimuthal angles of the source vector \mathbf{S} in the RGB coordinate system. From Eq. 2 it follows that

$$\mathbf{I}_{SUV} = (\bar{\mathbf{D}}f_d(\theta) + \bar{\mathbf{S}}f_s(\theta)) \hat{\mathbf{n}} \cdot \hat{\mathbf{I}}, \quad (3)$$

where $\bar{\mathbf{D}} = [R]\mathbf{D}$ and $\bar{\mathbf{S}} = [R]\mathbf{S} = [1, 0, 0]^\top$.

This SUV space is a *data-dependent* color space because it depends on the effective source color in the image. It has two important properties. First, it separates the diffuse and specular reflection effects. The first channel (the S channel) encodes the entire specular component and an unknown fraction of the diffuse component, while the remaining two channels (U and V) are independent of $f_s(\boldsymbol{\theta})$ and are functions of only the diffuse reflectance.

The second important property is that, since the transformation is linear, shading information is preserved. This is clear from Eq. 3; if \mathbf{r}_i^\top denotes the i^{th} row of $[R]$, the two diffuse channels are

$$\begin{aligned} I_U &= \mathbf{r}_2^\top \mathbf{D} f_d(\boldsymbol{\theta}) \hat{\mathbf{n}} \cdot \hat{\mathbf{l}} \quad \text{and} \\ I_V &= \mathbf{r}_3^\top \mathbf{D} f_d(\boldsymbol{\theta}) \hat{\mathbf{n}} \cdot \hat{\mathbf{l}}. \end{aligned} \quad (4)$$

Assuming Lambertian diffuse reflectance, $f_d(\boldsymbol{\theta})$ is a constant function of $\boldsymbol{\theta}$. In this case, the two-channel color vector

$$\mathbf{J} = [I_U, I_V]^\top \quad (5)$$

provides direct information about the normal vector on the surface, with the coefficients $\rho_U = \mathbf{r}_2^\top \mathbf{D} f_d$ and $\rho_V = \mathbf{r}_3^\top \mathbf{D} f_d$ behaving as Lambertian albedos.

Figure 2 compares this linear, data-dependent color space with conventional non-linear color spaces. Non-linear color spaces such as *r-g* chromaticity and hue-saturation-value (HSV) are computed by central projection. Each RGB pixel corresponds to a vector in the RGB cube, and these are intersected with the plane $R + G + B = c$ for some c . For example, hue and saturation correspond to the distance and polar angle of these intersection points relative to the cube diagonal, and chromaticity coordinates are derived from the intersection of these color vectors with the plane $R + G + B = 1$. Non-linear color spaces such as these are useful for recognition, for example, since they remove Lambertian shading and shadow information. (All positive scalar multiples of \mathbf{I} map to the same chromaticity coordinates and the same hue.) Since they do not preserve photometric information, however, they are generally unsuitable for recovering shape.

In contrast, the SUV color space *does* preserve shading information, and by providing two channels that are independent of the complex, specular component of reflectance, it can be a useful tool for scene analysis. One important application is photometric stereo, and Sect. 4 provides some examples in which these images are used for that purpose.

It is worth noting that Park [19] uses a transformation of color space that is similar to SUV space for the purposes of segmentation. His transformation is different from ours, however, since it is not a pure rotation and thus does not exactly isolate diffuse reflectance effects.

3.1. Practical Considerations

For shape recovery we are generally interested in the two-channel diffuse signal \mathbf{J} in Eq. 5. The quality of this signal depends on the spectral characteristics of the surface in addition to sensor noise.

Spectral Reflectance

When the surface is ‘white’, the spectral reflectance is a constant function of wavelength. In this case, $g_d(\lambda) = g_d$, and since

$$D_k = g_d \int C_k(\lambda) L(\lambda) d\lambda = g_d S_k,$$

it follows that the color vector \mathbf{I} and the source color \mathbf{S} are collinear in color space. For these surfaces, the UV channels are zero; and as a result, they provide no information about the surface, regardless of the illuminant and sensors that we choose. (This is the same restriction experienced by Klinker et al. [11]; when the diffuse and source colors are the same, there is no way to distinguish between the two reflection components.) In order to recover surface shape for these materials without restricting the BRDF, alternative (and more complex) reconstruction methods can be used. See, for example, Refs. [12, 30].

Sensor Noise

Assuming independent, additive Gaussian noise with zero mean and variance σ^2 in each of the three channels of color vector \mathbf{I} , and assuming $\|\mathbf{I}\| \leq 1$, the signal-to-noise ratio (denoted $\text{SNR}(\mathbf{I})$) is $10 \log_{10}(1/\sigma)$ dB. The magnitude of the diffuse color vector \mathbf{J} is related to that of the original color vector by $\|\mathbf{J}\| = \|\mathbf{I}\| \sin \alpha$, and since the noise is the same in both cases, it follows that

$$\text{SNR}(\mathbf{J}) = \text{SNR}(\mathbf{I}) + 10 \log_{10}(\sin \alpha). \quad (6)$$

This relationship is shown in Fig. 3, and it suggests that when the angle between the image and the source color is less than 10° , the two-channel diffuse signal suffers severe degradation. The effects of this degradation can be mitigated by using multiple exposures to collect high dynamic range (HDR) images. When exposures are chosen judiciously [8], we can increase the SNR of the original image and therefore that of the two-channel diffuse image as well.

3.2. Computing Diffuse Images

In order to compute the two-channel diffuse image, it is necessary to know the effective source color, \mathbf{S} . Under controlled conditions, the source color can be calibrated, and in an uncontrolled setting it can be estimated using a number of established techniques.

Tominga and Wandell [25] demonstrate a method for estimating the source color based on the fact that color vectors from a homogeneous surface span the dichromatic plane.




Figure 3. The signal-to-noise ratio (SNR) of the two-channel diffuse image (\mathbf{J}) relative to that of the original image (\mathbf{I}) as a function of α , the angle between \mathbf{I} and the source color \mathbf{S} in RGB color space. (See Fig. 2.)

They determine the source color by intersecting multiple dichromatic planes (from different dichromatic surfaces) in color space. Similarly, Lee [13] finds the source color by intersecting lines in 2D chromaticity space. More recently, Finlayson and Schaefer [5] use the fact that the chromaticity of the majority of illuminants lie along a known curve in chromaticity space and recover the source color from an image of a single homogeneous dichromatic surface.

Figure 1 shows a diffuse image that was computed using the source color determined by intersecting lines in chromaticity space [13]. (Three homogeneous surfaces were manually segmented for this purpose.) The monochrome diffuse image was created by transforming the input image to the SUV color space and computing $\|\mathbf{J}\|$ at each pixel. Comparing this to the original image, we see that the specular effects have been removed as expected. Note that the dichromatic model is violated when saturation occurs in the input images, and this causes errors at points of extreme specularity. Additional examples computed with calibrated sources are shown in Sect. 4.1.

4. Application to Photometric Stereo

The two-channel diffuse image of the previous section is derived without making any assumptions about the nature of specular reflectance, and yet it is independent of specular reflection effects for dichromatic surfaces. By assuming that the diffuse component of the reflectance is Lambertian, we can use these images along with established techniques for photometric stereo to obtain estimates of surface shape.

Assuming the light source color is known, the color vector \mathbf{I} at each pixel can be transformed to SUV space using Eq. 3. The two diffuse UV channels preserve shading information, and the two-channel image \mathbf{J} is independent of the specular component. Assuming that the diffuse reflection component is Lambertian, standard photometric stereo techniques can be used for surface reconstruction. Here, we discuss a modified version¹ of the color photometric stereo method of Barsky and Petrou [1].

Let \mathbf{J}^1 , \mathbf{J}^2 and \mathbf{J}^3 be three 2-channel color vectors produced by observing a single point under three different light

¹The method was originally designed for color images with three channels. Here we use the same method for a two-channel images.

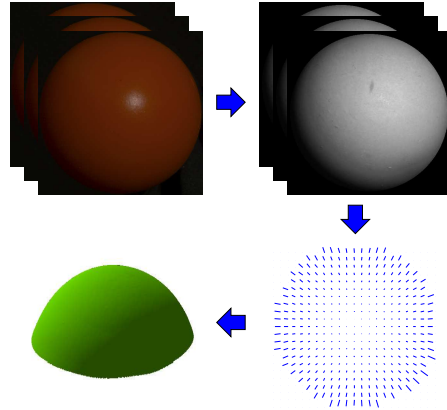


Figure 4. Photometric stereo using SUV color space. Three or more RGB images are acquired under known illumination conditions, and these are transformed to SUV space using the known source color. The UV channels represent diffuse images of the object, and these are used with standard photometric stereo techniques to estimate the surface normal at each pixel. The normals are integrated to recover the surface.

source directions $\hat{\mathbf{l}}^1$, $\hat{\mathbf{l}}^2$ and $\hat{\mathbf{l}}^3$. As in Eq. 4, it follows that

$$\mathbf{J}^k = [I_U^k, I_V^k]^\top = (\hat{\mathbf{n}} \cdot \hat{\mathbf{l}}^k) \boldsymbol{\rho}, \quad (7)$$

where $\boldsymbol{\rho}$ is a 2-channel UV albedo. A *shading vector* is defined according to $\mathbf{F} = [f^1, f^2, f^3]^\top = [\hat{\mathbf{l}}^1 \hat{\mathbf{l}}^2 \hat{\mathbf{l}}^3]^\top \hat{\mathbf{n}}$, and the shading information from all three observations is combined in an *intensity matrix* that satisfies

$$[\mathbf{J}] = \begin{bmatrix} J_1^1 & J_2^1 \\ J_1^2 & J_2^2 \\ J_1^3 & J_2^3 \end{bmatrix} = \begin{bmatrix} f^1 \rho_U & f^1 \rho_V \\ f^2 \rho_U & f^2 \rho_V \\ f^3 \rho_U & f^3 \rho_V \end{bmatrix} = \mathbf{F} \boldsymbol{\rho}^\top. \quad (8)$$

The least squares estimate of the shading vector \mathbf{F} can easily be computed from the intensity matrix; it is the principal eigenvector of $[\mathbf{J}][\mathbf{J}]^\top$. Once the shading vector is known, the surface normal is found by solving the matrix equation $\mathbf{F} = [\hat{\mathbf{l}}^1 \hat{\mathbf{l}}^2 \hat{\mathbf{l}}^3]^\top \hat{\mathbf{n}}$.

This reconstruction procedure is outlined in Fig. 4, and it can be applied without change to any number of images larger than three.

4.1. Experimental Results

To validate the method quantitatively, we used objects of known shape with varying material properties. The objects are shown in Fig. 7, and they consist of a set of spheres with increasing specularity. (The incident-plane BRDFs are shown in Fig. 5.)

For each sphere, a set of four high dynamic range (HDR) images were captured from a fixed viewpoint and four known illumination directions. The source color was calibrated by imaging a Macbeth color checker, and it was

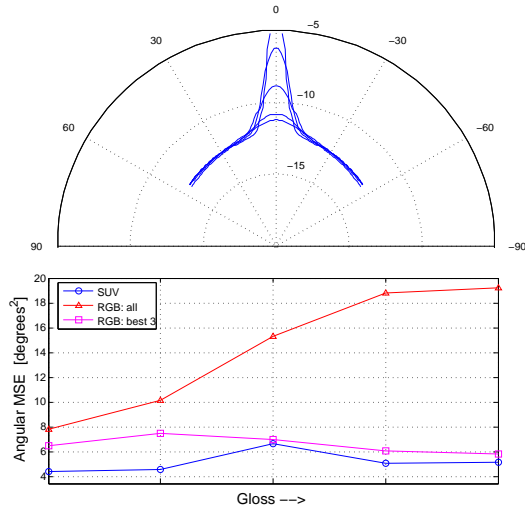


Figure 5. Comparison of photometric stereo using two different color spaces. Top: Relative BRDFs (in decibels) of the five red spheres of Fig. 7 as a function of half-angle. Bottom: Mean-square angular error in the recovered surface normals as a function of increasing specularity using both the SUV color space and existing RGB methods.

used to compute the SUV images as described in Sect. 3. The second column of Fig. 7 confirms that the UV channels of these images depends largely on the diffuse reflectance. (These monochrome diffuse images show the magnitude of the two diffuse channels, computed using $\|\mathbf{J}\| = (I_U^2 + I_V^2)^{\frac{1}{2}}$ at each pixel.)

Using the diffuse UV images, the surface normals of each sphere were estimated using the photometric stereo method described in Sect. 4. As a means of comparison, we implemented two different RGB-based photometric techniques. The first method uses all four RGB images and assumes Lambertian reflectance [1]. The second method assumes Lambertian+specular reflectance and reconstructs the surface by choosing the three ‘least specular’ RGB measurements at each pixel [2, 4].

The results are shown in Figs. 5 and 7. The recovered surfaces, including cross-sections overlaid on the true shape, are displayed in Fig. 7. More quantitative results are shown in Fig. 5, with the bottom of that figure displaying the angular difference between the true and estimated surface normals as a function of increasing specularity. (This plot shows the mean-square error computed over a window encompassing the visible specularities in the four images.) These results demonstrate that the SUV-based reconstruction is largely independent of the specular reflectance, whereas both the four-image and three-image RGB methods are affected by it. The four-image method assumes Lambertian reflectance and its performance degrades monotonically as gloss increases; and while the three-image RGB method performs well for the high-gloss (narrow specular

lobe) spheres, it performs less well when the angular support of the specular lobe is large relative to the separation of the light source directions.

Figure 6 shows the results of applying our SUV-based photometric stereo method to two natural objects (a pear and a pumpkin.) For each object, four HDR images were captured and transformed to SUV color space, and the two-channel diffuse images were used to estimate the surface normals. For each object, we show an input image and the surfaces recovered by integrating these surface normals. Since the transformation to SUV space is purely local, the method requires no spatial coherence in the image, and it performs well for surfaces with arbitrary texture as shown by the pear example. This is not true for alternative photometric stereo techniques that use color-based highlight removal [11,21], since these methods generally require spatial coherence.

5. Conclusion

This paper presents a data-dependent rotation of RGB color space that separates the diffuse reflection effects from the more complex, specular reflection effects in images of dichromatic surfaces. Since it is linear, this transformation preserves diffuse shading information, and images in this transformed space can be used to obtain photometric reconstructions that are independent of the specular reflectance. The result is a photometric stereo method for non-diffuse materials that does not require an explicit BRDF model or reference object. The method is evaluated both qualitatively and quantitatively, and it is shown to perform well for both painted surfaces with varying specular reflectance and natural surfaces with and without texture.

We currently assume that the color of the illuminant is constant over the surface, although one can imagine enhancing this method with polarization, for example, to estimate the source color locally. Additionally, although it is not explored in this paper, it is straight-forward to adapt the method to the ‘uncalibrated’ case in which the lighting directions are unknown. Finally, while this paper explores the application of this color transformation to photometric stereo, it can be useful for any visual task (e.g., shape-from-shading, stereo, structure-from-motion) that is simplified by the absence of specular reflectance. Exploring these applications is an interesting direction for future work.

Acknowledgments

The authors would like to thank Sameer Agarwal for helpful suggestions regarding this manuscript. This work was supported in part by the National Science Foundation. D. J. Kriegman and S. P. Mallick were supported under grants IIS-0308185 and EIA-0224431, and P. N. Belhumeur was supported under grants IIS-00-85864 and IIS-03-08185.




Figure 6. SUV-based photometric stereo applied to natural surfaces. Left: Input RGB images show significant specular reflectance and texture. By transforming the images to SUV space, the specular effects are removed, enabling accurate recovery of shape. Middle, Right: The surfaces recovered by integrating the estimated surface normals.

References

- [1] S. Barsky and M. Petrou. Colour photometric stereo: Simultaneous reconstruction of local gradient and colour of rough textured surfaces. In *Proc. Int. Conf. on Computer Vision*, pages 600–605, 2001.
- [2] S. Barsky and M. Petrou. The 4-source photometric stereo technique for three-dimensional surfaces in the presence of highlights and shadows. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 25(10):1239–1252, October 2003.
- [3] V. Blanz and T. Vetter. Face recognition based on fitting a 3D morphable model. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 25(9), september 2003.
- [4] E. Coleman and R. Jain. Obtaining 3-dimensional shape of textured and specular surfaces using four-source photometry. *Computer Vision, Graphics and Image Processing*, 18(4):309–328, April 1982.
- [5] G. Finlayson and G. Schaefer. Constrained dichromatic colour constancy. In *Proc. European Conf. on Computer Vision*, volume 1, pages 342–358, 2001.
- [6] A. Georghiadis. Incorporating the Torrance and Sparrow model of reflectance in uncalibrated photometric stereo. In *Proc. Int. Conf. on Computer Vision*, volume 1, pages 816–823, 2003.
- [7] R. Gershon. *The use of color in computational vision*. PhD thesis, University of Toronto, 1987.
- [8] M. D. Grossberg and S. K. Nayar. High dynamic range from multiple images: Which exposures to combine? In *Proc. IEEE Workshop on Color and Photometric Methods in Computer Vision (CPMCV)*, October 2003.
- [9] A. Hertzmann and S. Seitz. Shape and material by example: a photometric stereo approach. In *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, 2003.
- [10] K. Ikeuchi. Determining surface orientations of specular surfaces by using the photometric stereo method. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 3(6):661–669, 1981.
- [11] G. Klinker, S. Shafer, and T. Kanade. The measurement of highlights in color images. *Int. Journal of Computer Vision*, 2(1):7–32, 1988.
- [12] M. Koudelka, S. Magda, P. Belhumeur, and D. Kriegman. Image-based modeling and rendering of surfaces with arbitrary BRDFs. In *Proc. IEEE Conf. Computer Vision and Pattern Recognition*, pages 568–575, 2001.
- [13] H.-S. Lee. Method for computing the scene-illuminant chromaticity from specular highlights. *J. Optical Society of America A*, 3(10):1694–1699, October 1986.
- [14] J. Lu and J. Little. Reflectance and shape from images using a collinear light source. *Int. Journal of Computer Vision*, 32(3):1–28, 1999.
- [15] S. Magda, D. Kriegman, T. Zickler, and P. Belhumeur. Beyond Lambert: Reconstructing surfaces with arbitrary BRDFs. In *Proc. Int. Conf. on Computer Vision*, pages II: 391–398, 2001.
- [16] E. Marszalec, B. Martinkauppi, M. Soriano, and M. Pietikäinen. Physics-based database for color research. *Journal of Electronic Imaging*, 9(1):32–38, January 2000.
- [17] S. Nayar, X. Fang, and T. Boulton. Separation of reflection components using color and polarization. *Int. Journal of Computer Vision*, 21(3):163–186, 1997.
- [18] S. Nayar, K. Ikeuchi, and T. Kanade. Determining shape and reflectance of hybrid surfaces by photometric sampling. *IEEE J. of Robotics and Automation*, 6(4):418–431, 1990.
- [19] J. Park. Efficient color representation for image segmentation under non-white illumination. In *Intelligent Robots and*

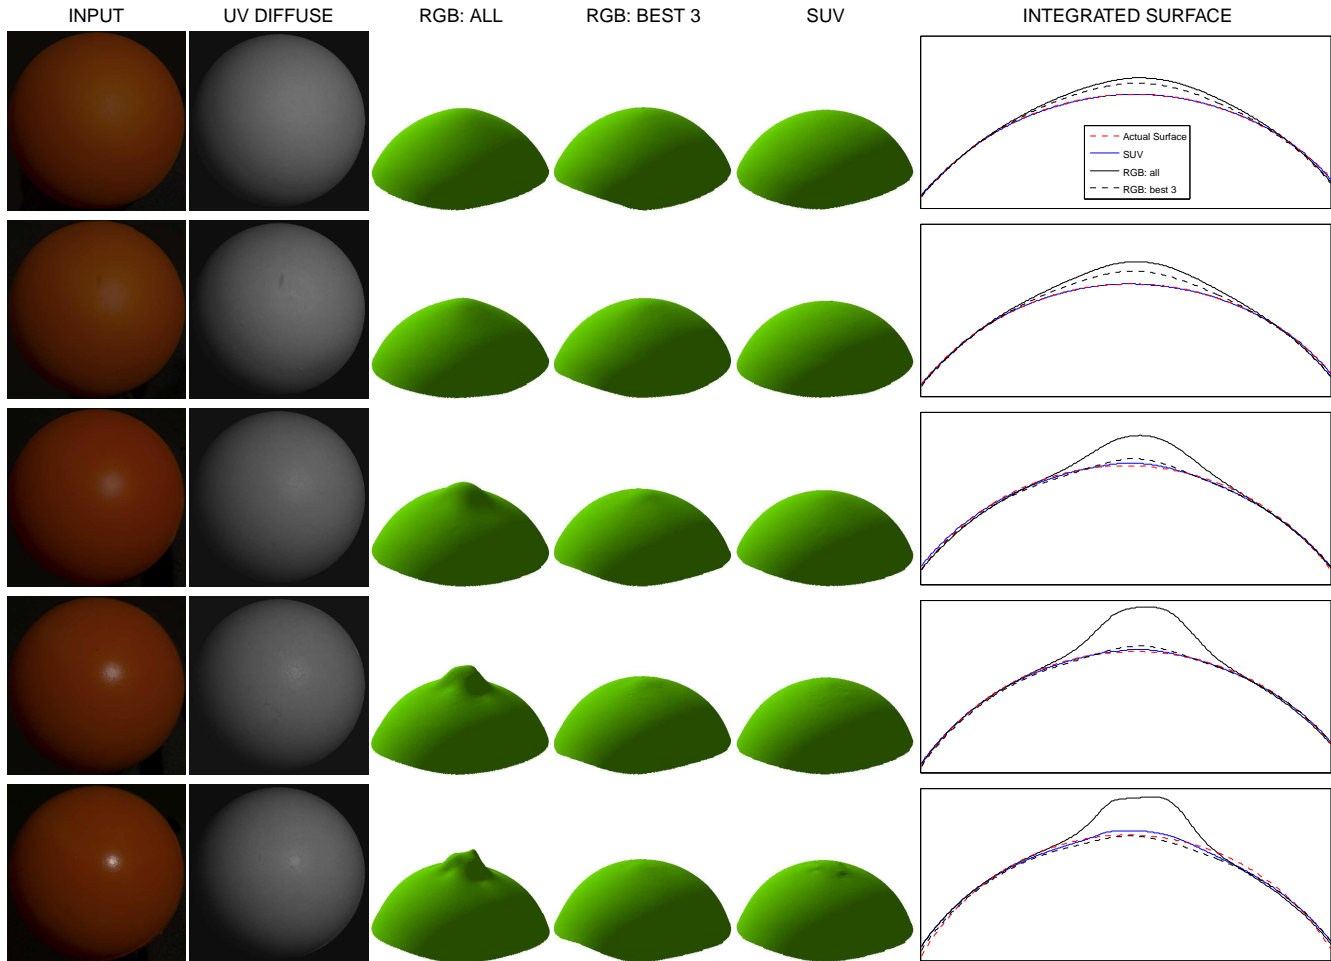


Figure 7. Comparison of photometric stereo methods using two different color spaces. Five red spheres with increasing specular reflectance are each observed under four illumination directions, and these images are used to recover the surface. From left to right, each row shows: i) an input RGB image, ii) the corresponding diffuse image given by the linear transformation to SUV space, iii) surfaces integrated from the surface normals estimated by three photometric stereo methods, and iv) cross-sections of the surfaces overlaid on the true shape.

- Computer Vision XXI: Algorithms, Techniques, and Active Vision (Proc. SPIE, Vol. 5267)*, pages 163–174, 2003.
- [20] Y. Sato and K. Ikeutchi. Temporal-color space analysis of reflection. *J. Optical Society of America A*, 11(11):2990–3002, 1994.
- [21] K. Schlüns and O. Wittig. Photometric stereo for non-Lambertian surfaces using color information. In *Proc. Int. Conf. on Image Analysis and Processing*, pages 505–512, September 1993.
- [22] S. Shafer. Using color to separate reflection components. *COLOR research and applications*, 10(4):210–218, 1985.
- [23] W. Silver. Determining shape and reflectance using multiple images. Master’s thesis, MIT, 1980.
- [24] H. Tagare and R. deFigueiredo. A theory of photometric stereo for a class of diffuse non-lambertian surfaces. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 13(2):133–152, February 1991.
- [25] S. Tominga and B. Wandell. Standard surface-reflectance model and illuminant estimation. *J. Optical Society of America A*, 6(4):576–584, April 1989.
- [26] N. Tsumura, N. Ojima, K. Sato, M. Shiraishi, H. Shimizu, H. Nabeshima, S. Akazaki, K. Hori, and Y. Miyake. Image-based skin color and texture analysis/synthesis by extracting hemoglobin and melanin information in the skin. In *Proc. SIGGRAPH*, pages 770–779, 2003.
- [27] H. Wann Jensen, S. Marschner, M. Levoy, and P. Hanrahan. A practical model for subsurface light transport. In *Proc. SIGGRAPH*, pages 511–518, 2001.
- [28] L. B. Wolff and T. E. Boult. Constraining object features using a polarization reflectance model. *IEEE Trans. Pattern Analysis and Machine Intelligence*, 13(7):635–657, 1991.
- [29] R. Woodham. Photometric stereo: A reflectance map technique for determining surface orientation from image intensity. In *Proc. SPIE*, volume 155, pages 136–143, 1978.
- [30] T. Zickler, P. Belhumeur, and D. Kriegman. Helmholtz stereopsis: Exploiting reciprocity for surface reconstruction. In *Proc. European Conf. on Computer Vision*, pages III: 869–884, May 2002.