

Color Subspaces as Photometric Invariants

Todd Zickler · Satya P. Mallick · David J. Kriegman ·
Peter N. Belhumeur

Received: 9 August 2006 / Accepted: 10 September 2007 / Published online: 30 October 2007
© Springer Science+Business Media, LLC 2007

Abstract Complex reflectance phenomena such as specular reflections confound many vision problems since they produce image ‘features’ that do not correspond directly to intrinsic surface properties such as shape and spectral reflectance. A common approach to mitigate these effects is to explore functions of an image that are invariant to these photometric events. In this paper we describe a class of such invariants that result from exploiting color information in images of dichromatic surfaces. These invariants are derived from illuminant-dependent ‘subspaces’ of RGB color space, and they enable the application of Lambertian-based vision techniques to a broad class of specular, non-Lambertian scenes. Using implementations of recent algorithms taken from the literature, we demonstrate the practical utility of these invariants for a wide variety of applications, including stereo, shape from shading, photometric stereo, material-based segmentation, and motion estimation.

Keywords Photometric invariants · Shape invariants · Color spaces · Dichromatic reflection · Multispectral imaging · Surface reconstruction · Photometric stereo · Shape from shading · Stereo · Color-based segmentation · Color-based optical flow

T. Zickler (✉)
Harvard University, 33 Oxford St., Cambridge, MA 02138, USA
e-mail: zickler@eecs.harvard.edu

S.P. Mallick · D.J. Kriegman
University of California, San Diego USA

P.N. Belhumeur
Columbia University, New York, NY 10027, USA

1 Introduction

An image is the product of the shape, reflectance and illumination in a scene. For many visual tasks, we require only a subset of this information, and we wish to extract it in a manner that is insensitive to variations in the remaining ‘confounding’ scene properties. For 3D reconstruction, for example, we seek accurate estimates of shape, and we design systems that are insensitive to variations in reflectance and illumination.

One practical approach to these problems is to compute a function of the input images that is invariant to confounding scene properties but is discriminative with respect to desired scene information. A number of these invariants are described in the literature, with the simplest example being a normalized-RGB image. For a Lambertian scene, the normalized RGB color vector at each pixel depends on the spectral reflectance of the corresponding surface patch but not its orientation with respect to a light source. It is a useful invariant for material-based segmentation.

Like normalized-RGB, most existing invariants seek to isolate information about the material properties in a scene and are therefore designed to be invariant to local illumination and viewing geometry. In contrast, this paper considers a class of invariants that deliberately preserve geometry information in a way that is invariant to specular reflections. The proposed invariants provide direct access to surface shape information through diffuse shading effects, and since diffuse shading is often well approximated by the Lambertian model, they satisfy the ‘constant-brightness assumption’ underlying most approaches to stereo reconstruction and structure-from-motion. In addition, these invariants provide access to surface normal information, which can be recovered using Lambertian-based photometric reconstruction methods.

The idea underlying the proposed invariants can be interpreted geometrically. When the illuminant color is known, and the reflectance of surfaces can be represented by the dichromatic model (Shafer 1985), we can linearly transform the space of RGB tristimulus vectors in a way that isolates specular reflection effects. Following the transformation, two sensor channels are free of these effects, and this two-dimensional “color subspace” constitutes a specular invariant. Since this operation is linear, the diffuse shading information is preserved by the transformation, and the invariant can be exploited photometrically. Also, the method places no restrictions on scene texture because the computation operates independently at each image point. Finally, it only requires knowledge about the spectral content of scene illumination and therefore makes no assumptions about the spatial distribution of light sources.

This paper begins with the case of RGB images and singly-colored illumination environments (Sect. 3), in which case the linear transformation can be interpreted as a transformation to an alternative, illuminant-dependent color space. We refer to this space as *SUV color space*. In addition to providing a specular invariant, we show that this color space leads naturally to a notion of generalized hue (Sect. 3.1). We are not limited to this case, however, and a similar procedure can be shown to handle mixed-illumination environments and hyper-spectral images (Sect. 4). To assess the utility of the proposed invariants they are applied in number of visual tasks, including binocular stereo, shape-from-shading, photometric stereo, optical flow estimation, and segmentation (Sect. 5). In each of these cases, when the source colors are known, significant improvements result from computing the invariants as a pre-process.

2 Background and Related Work

As mentioned in the introduction, many existing invariants seek to isolate information about material properties in a scene. One such property is surface reflectance, which is often described by the bi-directional reflectance distribution function, or BRDF. Here, we consider the BRDF to be a five-dimensional function of wavelength and imaging geometry, and we write it $f(\boldsymbol{\theta}, \lambda)$, 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. The simplest model of reflectance is the Lambertian model, according to which the BRDF is a constant function of the imaging geometry, so that $f(\boldsymbol{\theta}, \lambda) = f(\lambda)$.

A number of photometric invariants have been proposed for Lambertian scenes. Normalized-RGB, r - g chromaticity, and hue/saturation images are all examples of representations that are independent of *diffuse shading* (the geometric

relation between a surface normal and the illumination direction) and depend only on the spectral reflectance of the surface and the spectral power distribution (SPD) of the illuminant. Additional invariants to either local geometry or spectral reflectance can be computed from “reflectance ratios” when multiple images of a scene are available (e.g., Wolff and Angelopoulou 1994), or when the reflectance of the surface is spatially coherent (e.g., Nayar and Bolle 1996); an invariant to both local geometry and illuminant spectral power distribution (SPD) can be computed from a single image under appropriate imaging conditions (Hordley et al. 2002).

Invariants for more general scenes, including some scenes with specularities, can be derived from the Shafer’s dichromatic model of reflectance (Shafer 1985). According to this model, the BRDF of the surface can be decomposed into two additive components: the interface (specular) reflectance and the body (diffuse) reflectance. In theory, by separating an image according to these components, one can obtain invariants to either diffuse or specular reflection effects.

According to the dichromatic model (with the neutral interface assumption, Lee et al. 1990), the observation of a surface point can be written

$$e_k = \sigma_d d_k + \sigma_s s_k, \quad (1)$$

where σ_d and σ_s are geometric scale factors that depend on the material properties and the local view and illumination geometry ($\boldsymbol{\theta}$), and

$$d_k = \int E(\lambda) R(\lambda) C_k(\lambda) d\lambda, \quad (2)$$

$$s_k = \int E(\lambda) C_k(\lambda) d\lambda. \quad (3)$$

Here, $E(\lambda)$ is the SPD of the incident illumination, $R(\lambda)$ is the spectral reflectance of the surface, and $C_k(\lambda)$ is the spectral sensitivity of a linear sensor. A typical RGB camera yields three such observations, and in this case we write $\mathbf{e}_{RGB} = \{e_k\}_{k=R,G,B}$ and define $\mathbf{d} = \{d_k\}_{k=R,G,B}$ and $\mathbf{s} = \{s_k\}_{k=R,G,B}$ to be the *diffuse color* and *specular color*, respectively. These are conventionally assumed to be vectors of unit length.

There is practical utility in separating the diffuse and specular components in an image. Since diffuse reflections are typically well-represented by the Lambertian model, computing this separation as a pre-process allows the application of powerful Lambertian-based vision algorithms to a variety of non-Lambertian scenes. Materials that can be treated in this way include plant leaves, cloth, wood, and the skin of fruits (Lee et al. 1990; Tominga and Wandell 1989) in addition to a large number of dielectrics (Healey 1989). The dichromatic BRDF model has also proven useful for

a number of applications involving human skin (e.g., face recognition (Blanz and Vetter 2003), pigment-based analysis and synthesis (Tsumura et al. 2003)), even though the reflectance of human skin is more accurately described by the higher dimensional BSSRDF (Wann Jensen et al. 2001).

Despite its apparent utility, image analysis relying on explicit decomposition of the diffuse and specular components is rare because the separation problem is ill-posed. Classically, this separation problem is addressed using color histogram analysis. As made clear by (1), in the RGB cube, a collection of color vectors from a dichromatic material under multiple view and illumination configurations (i.e., different values of θ) lie in the *dichromatic plane*—the plane spanned by the specular and diffuse colors, \mathbf{s} and \mathbf{d} (Shafer 1985). 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 reflections (Gershon 1987; Klinker et al. 1988). When these limbs are sufficiently distinct, the diffuse and source vectors can be recovered, the two components can be separated, and the highlights can be removed (Klinker et al. 1988).

While this method works well for homogeneous, dichromatic surfaces in the noiseless case, there are two 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.) Overcoming these restrictions generally requires additional assumptions regarding spatial coherence on the surface (Nayar et al. 1997; Mallick et al. 2006; Tan and Ikeuchi 2003; Tan et al. 2003), specific parametric models for specular reflectance (Ragheb and Hancock 2001), or the use of multiple images that exploit additional cues such as polarization (Wolff and Boult 1991; Nayar et al. 1997).

When the source color is known and constant over a scene, one can compute specular invariants that are based on transformations of RGB color space and do not require explicit specular/diffuse separation. This is the approach taken in this paper, and it is related to the work of Tan and Ikeuchi (2003), who obtain such a specular invariant using a source-dependent *non-linear* transformation of the RGB values at a pixel. The transformation is computed independently at each point, and it yields a positive grayscale image that depends only on diffuse reflections (σ_d and \mathbf{d}) and is independent of specular effects (σ_s). Another non-linear transformation that provides a similar invariant under white illumination is proposed by Yoon and Kweon (2006a). As an alternative, Park (2003) defines a *linear* transformation that provides

two channels that, while not pure invariants, are highly insensitive to specular reflections. Following this transformation, the measurements in one channel correspond predominantly to specular reflectance information, while the other two are predominantly diffuse. Unlike these existing methods, we present true invariants that are computed linearly, and hence have the unique property of preserving diffuse shading (and geometry) information.

The invariants presented in this paper assume knowledge of the scene illuminant \mathbf{s} . In controlled environments, or when the illuminants do not change significantly over time, the required source color vectors can be measured by imaging a calibration target. This is the approach taken in this paper. While not explored here, it may be possible to apply these invariants in more uncontrolled environments by combining them with existing image-based methods for illuminant estimation. For scenes with sufficient color diversity, for example, one can estimate the illuminant color using statistical knowledge of common sources and surfaces (Brainard and Freeman 1997; Finlayson 1996; Finlayson et al. 2001; Lehmann and Palm 2001; Rosenberg et al. 2001; Sapiro 1999; Tominga and Wandell 2002), and for glossy scenes with only a small population of diffuse colors, it can be estimated using methods based on the dichromatic model (Finlayson and Schaefer 2001; Lee 1986; Tan et al. 2004; Tominga and Wandell 1989). The accuracy of these methods depends on the materials and illuminants that are present in a particular scene, so in a generic setting, one would probably want to use some combination of them. For discussions, and for detailed evaluations of some of these algorithms, the reader is referred to (Barnard et al. 2002a, 2002b; Hordley and Finlayson 2006).

Invariants for scenes with more general reflectance functions are developed by Narasimhan et al. (2003). They describe a general model of reflectance consisting of a product of a ‘‘material’’ term (Lambertian albedo, Fresnel coefficient, etc.) and a ‘‘geometry’’ term that encodes the relationship between the surface normal, light-source, and viewing direction. Invariants to both of these terms can be computed from either multiple observations of a single point under variable view or illumination, or from one observation of a spatially-coherent scene. The geometry invariant is of particular interest, since it can be used directly for material-based segmentation.

3 A Source-Dependent Color Space

Suppose we treat RGB tristimulus values as points in \mathbb{R}^3 and linearly transform the RGB coordinate system by rotating the axes. Also, as shown in the left of Fig. 1, suppose this rotation is such that one of the axes (red, say) becomes




Fig. 1 Linear and non-linear transformations of the RGB cube. Three observations of the same material yield color vectors $\mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3$ in the dichromatic plane spanned by the source and diffuse colors \mathbf{s} and \mathbf{d} . *Left:* The SUV color space is defined by a rotation of the RGB coordinate vectors. One axis is aligned with the source color, and two of three resulting channels (UV) are invariant to specular reflections. Diffuse shading information is preserved in these channels and can be used to recover shape. Additionally, the ratio between the U and V channels represents *generalized hue* (ψ), which provides a second invariant depending only on spectral reflectance. *Right:* Unlike SUV space, central projection used to compute *r-g* chromaticity values and HSV-type color spaces does not preserve diffuse shading information

aligned with the direction of the effective RGB source vector \mathbf{s} . This transformation defines a new color space (see below), which we refer to as the *SUV color space*. It can be defined according to $\mathbf{e}_{SUV} = \mathbf{R}\mathbf{e}_{RGB}$ using any $\mathbf{R} \in SO(3)$ that satisfies $\mathbf{R}\mathbf{s} = (0, 0, 1)$. From (1) it follows that tristimulus vectors in the transformed space satisfy

$$\mathbf{e}_{SUV} = (\bar{\mathbf{d}}\sigma_d + \bar{\mathbf{s}}\sigma_s), \tag{4}$$

with

$$\bar{\mathbf{d}} = \mathbf{R}\mathbf{d}, \quad \text{and} \quad \bar{\mathbf{s}} = \mathbf{R}\mathbf{s} = (0, 0, 1).$$

Notice that according to our definition, the S channel is uniquely defined for a given \mathbf{s} (and thus a given illuminant SPD and sensor), while the U and V channels can be arbitrarily chosen from the family of orthonormal bases for the plane orthogonal to \mathbf{s} .

The SUV color space is a *source-dependent* color space because it depends on the effective source color vector in the image. It has two important properties. First, it separates the diffuse and specular reflection effects. 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 σ_s and are therefore specular invariants.

The second important property is that shading information is preserved by the linear transformation. This is clear from (4). If \mathbf{r}_i^\top denotes the i th row of \mathbf{R} , the values of the two diffuse channels satisfy

$$e_U = \mathbf{r}_1^\top \mathbf{d}\sigma_d \quad \text{and} \quad e_V = \mathbf{r}_2^\top \mathbf{d}\sigma_d. \tag{5}$$



Fig. 2 Input RGB image (*left*) and its corresponding specular invariant (*right*) computed pixel-wise according to (7) using the known illuminant color

Assuming Lambertian diffuse reflectance, σ_d is a constant function of the local view and illumination directions. In this case, the two-channel color vector

$$\mathbf{j} = (e_U, e_V) \tag{6}$$

and its monochromatic relative

$$j = \sqrt{e_U^2 + e_V^2} \tag{7}$$

provide direct information about the normal vector on the surface, with the terms $\mathbf{r}_2^\top \mathbf{d}$ and $\mathbf{r}_3^\top \mathbf{d}$ in (5) contributing to the effective Lambertian albedo values.

An example of the monochromatic specular invariant computed from SUV space is shown in Fig. 2. In this example, the invariant was computed using the source color determined by intersecting lines in chromaticity space (Lee 1986), and then transforming the image from RGB to SUV space on a pixel-by-pixel basis. (Here, we choose $\mathbf{R} = \mathbf{R}_G(-\theta_s)\mathbf{R}_B(\phi_s)$ where $\mathbf{R}_k(\theta)$ 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.) Comparing the result to the original image, we see that specular effects are largely removed. Note that the dichromatic model is violated when saturation occurs in the input images, and this causes errors at points of extreme brightness.

To see that SUV space is in fact a color space, recall that any linear color space can be defined by a linear transformation of the color matching functions of another. Such a transformation provides a mapping, say, between the ISO RGB color space (with an identified white point) and the CIE XYZ color space, and it induces a corresponding invertible linear mapping between the tristimulus vectors in the two spaces. The rotation matrix described above is a coordinate transformation, and it therefore defines a spectral space that is related to the original sensor space through a corresponding linear transformation of the sensor sensitivity functions. The important point is that following the transformation, the illuminant SPD $E(\lambda)$ integrates to black against two of the three transformed sensitivity functions. Thus, by

converting to SUV space, we are implicitly choosing a transformation of the sensor such that the transformed sensitivities \bar{C}_k satisfy

$$\int E(\lambda)\bar{C}_k(\lambda)d\lambda = 0, \quad k = 1, 2.$$

It is clear that the same invariant properties could be obtained using any transformation $\mathbf{T} \in GL(3)$ satisfying $\mathbf{T}\mathbf{s} \propto [0, 0, 1]^T$. The rotation matrix used in the definition above is simply one practical choice.

Figure 1 compares the linear, source-dependent SUV color space with conventional non-linear representations of color that also have invariant properties. Non-linear representations such as *r-g* chromaticity and hue-saturation-value (HSV) are computed by central projection. Each RGB vector in the RGB cube is intersected with the plane $R + G + B = c$ for some constant 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 representations such as these are useful for recognition, for example, because they remove Lambertian shading and shadow information. All positive scalar multiples of \mathbf{e}_{RGB} map to the same chromaticity coordinates and the same hue.

In contrast, the diffuse channels of SUV color space preserve diffuse reflection effects encoded in the geometric scale factor σ_d . Since diffuse reflectance is often well-approximated by the Lambertian model, this implies that the specular-invariant image often: (1) satisfies the ‘constant-brightness assumption’ underlying most stereo and structure-from-motion systems; and (2) provides access to surface normal information through Lambertian-based photometric reconstruction methods such as shape-from-shading and photometric stereo. As a result, by computing these invariants as a pre-processing step, we can successfully apply many Lambertian-based algorithms to a much broader class of specular, non-Lambertian surfaces. Applications are explored in Sect. 5.

3.1 Generalized Hue

An additional invariant is created by taking the ratio between the specular invariant channels of (6). The result,

$$e_U/e_V = \mathbf{r}_1^T \mathbf{d}/\mathbf{r}_2^T \mathbf{d},$$

is independent of both the diffuse and specular geometric scale factors σ_d and σ_s . As shown in Fig. 1, it is instructive to interpret this ratio as an angle and define

$$\psi = \tan^{-1}(e_U/e_V) = \tan^{-1}(\mathbf{r}_1^T \mathbf{d}/\mathbf{r}_2^T \mathbf{d}), \tag{8}$$

which we refer to as *generalized hue*. Notice that ψ reduces to the standard definition of hue when the source color \mathbf{s} is white.

Examples of generalized hue images are shown in Fig. 3 for a specular globe under two different source colors. In each case, the source vector is measured by imaging a Macbeth color checker, this vector is used to compute a two-channel subspace image as in (6), and the ratio between the two channels is used to compute ψ . Since it depends only on \mathbf{d} , the value of ψ within each country on the globe is constant and is invariant to both specular reflections and diffuse shading.

4 Color Subspaces

If we again think of RGB vectors as points in \mathbb{R}^3 , the invariants defined in the previous section are seen to derive from a projection onto the two-dimensional subspace orthogonal to the source vector \mathbf{s} . (See left of Fig. 1.) Based on this interpretation, the invariants defined in (6) and (7) can be generalized to environments with mixed illumination.

The invariants of the previous section are based on (1), which in turn is premised on the assumption that the illuminant SPD is constant over the incident hemisphere of a surface point (i.e., that the illuminant ‘color’ is the same in all directions). Notationally, if $L(\boldsymbol{\omega}_i, \lambda)$ represents the incident radiance at a surface point, where $\boldsymbol{\omega}_i = (\theta_i, \phi_i) \in \Omega$ parameterizes the hemisphere of incident directions, the model requires that this input radiance field can be factored (with a slight abuse of notation) as $L(\boldsymbol{\omega})E(\lambda)$. To relate this to the terms in (1), recall that $f(\boldsymbol{\theta}, \lambda)$ with $\boldsymbol{\theta} = (\theta_r, \phi_r, \theta_s, \phi_s)$ denotes the BRDF of the surface, and write the image formation equation as

$$e_k = \int_{\lambda} \int_{\Omega} f(\boldsymbol{\theta}, \lambda)L(\boldsymbol{\omega}_i, \lambda)C_k(\lambda) \cos \theta_i \, d\boldsymbol{\omega}_i d\lambda. \tag{9}$$


According to the dichromatic model, the BRDF of the surface can be decomposed into additive diffuse and specular components, and each of these two components can be factored into a univariate function of wavelength and a multivariate function that depends on the imaging geometry. Finally, assuming a neutral interface, the index of refraction on the surface is constant over the visible spectrum, and the specular function of wavelength is constant. This leads to the common expression for the BRDF of a dichromatic surface,

$$f(\boldsymbol{\theta}, \lambda) = f_d(\boldsymbol{\theta})R(\lambda) + k_s f_s(\boldsymbol{\theta}), \tag{10}$$

where k_s is a constant. Substituting into (9) yields the expressions

$$\sigma_d = \int_{\Omega} f_d(\boldsymbol{\theta})L(\boldsymbol{\omega}_i) \cos \theta_i \, d\boldsymbol{\omega}_i,$$

Fig. 3 Pseudo-colored generalized hue images, each computed from a single RGB image of a globe under point source illumination having a distinct color. Generalized hue is invariant to both specularities and diffuse shading, and depends only on the spectral reflectance of the surface



$$\sigma_s = k_s \int_{\Omega} f_s(\theta) L(\omega_i) \cos \theta_i d\omega_i,$$

$$d_k = \int R(\lambda) E(\lambda) C_k(\lambda) d\lambda,$$

$$s_k = \int E(\lambda) C_k(\lambda) d\lambda.$$

To generalize the model, we consider a mixed-illumination environment whose spectral content can be written in terms of a finite linear basis:

$$L(\omega_i, \lambda) = \sum_{j=1}^N L_j(\omega_i) E_j(\lambda). \tag{11}$$

An example with $N = 2$ is an office environment where the illumination in every direction can be described as a mixture of daylight and fluorescent light. When the input radiance field can be decomposed in this manner, the BRDF decomposition of (10) yields

$$e_k = \sum_{j=1}^N \sigma_d^{(j)} d_k^{(j)} + \sigma_s^{(j)} s_k^{(j)}, \tag{12}$$

with

$$\sigma_d^{(j)} = \int_{\Omega} f_d(\theta) L_j(\omega_i) \cos \theta_i d\omega_i,$$

$$\sigma_s^{(j)} = k_s \int_{\Omega} f_s(\theta) L_j(\omega_i) \cos \theta_i d\omega_i,$$

$$d_k^{(j)} = \int R(\lambda) E_j(\lambda) C_k(\lambda) d\lambda,$$

$$s_k^{(j)} = \int E_j(\lambda) C_k(\lambda) d\lambda.$$

Equation (12) suggests the existence of a specular invariant that is analogous to the two-dimensional subspace defined of (6). In that section, the illuminant color is assumed constant over the input hemisphere (which corresponds to $N = 1$ in (12)) and the specular invariant subspace computed from a three-channel RGB image is two-dimensional. In general, given an M -channel (possibly hyper-spectral) image \mathbf{e} and an N -dimensional spectral basis $\{E_j(\lambda)\}_{j=1 \dots N}$ for the incident illumination, there exists a subspace of dimension $(M - N)$ that is independent of all $\sigma_s^{(j)}$ and there-




Fig. 4 (Color online) *Left:* For any mixture of two source SPDs, the specular invariant subspace is one-dimensional. By projecting RGB color vectors onto this line, a specular invariant can still be computed. *Right:* Two frames of an RGB video of a scene with mixed illumination and the corresponding specular invariants. A blue light on the *right* and a yellow light on the *left* induce complex specular effects. Projecting these images onto the one-dimensional subspace orthogonal to the source color vectors in RGB space yields an invariant to specular reflections that preserves diffuse shading information

fore invariant to specular reflections. Letting $\{\mathbf{r}_l\}_{l=1 \dots (M-N)}$ represent an orthonormal basis for this specular invariant subspace, the l th component (or ‘channel’) of the specular invariant image is given by

$$j_l = \mathbf{e}^T \mathbf{r}_l = \sum_{j=1}^N \sigma_d^{(j)} \mathbf{r}_l^T \mathbf{d}^{(j)}. \tag{13}$$

A specular invariant image with $(M - N)$ channels defined by this equation can be treated as an image, and as is the case for the U and V channels of Sect. 3, the channel values in this image can assume negative values. It is often more convenient to use the monochromatic specular invariant given by

$$j_{\text{inv}(M-N)} = \left(\sum_{l=1}^{M-N} j_l^2 \right)^{\frac{1}{2}}, \tag{14}$$

where the subscript $j_{\text{inv}(u)}$ is used to indicate that the grayscale invariant is derived from a u -dimensional specular invariant subspace. It is clear that (6) and (7) are specific examples of these invariants for the case $M = 3$ and $N = 1$.

Since the vast majority of cameras record three (RGB) channels, another interesting case to consider is $M = 3$, $N = 2$. An example is shown in Fig. 4, where light comes

from two sources with different SPDs. These SPDs induce two source color vectors $\mathbf{s}^{(1)}$ and $\mathbf{s}^{(2)}$ in RGB space (these are measured by imaging a calibration target), and by projecting the RGB color vectors of the input image onto the one-dimensional subspace orthogonal to these vectors, we create an image that is void of specular reflection effects.

4.1 Generalized Hue under Mixed Illumination

The concept of generalized hue (Sect. 3.1) can also be extended to handle hyper-spectral images and mixed illumination. In an M -channel image of a scene illuminated by a mixture of N illuminant SPDs, generalized hue can be defined as a scalar function defined on the surface of an $(M - N - 1)$ -dimensional unit sphere embedded in the $(M - N)$ -dimensional diffuse space. The sphere may be parameterized by a vector of angles Ψ .

As with RGB sensors and single illuminants, this expanded notion of generalized hue is independent of both shading and specularity, and it is consistent in that it reduces to the standard definition of hue for an RGB image acquired under white light.

4.2 Practical Considerations

The quality of the specular invariant signal depends on the spectral characteristics of the scene and the accuracy of the estimated source vectors. We discuss each separately in this section.

Spectral Characteristics When a surface is ‘white’, the spectral reflectance is a constant function of wavelength, so that $R(\lambda) = R$. In this case, since

$$d_k = R \int E(\lambda)C_k(\lambda)d\lambda = R s_k,$$

it follows that the observed color vector \mathbf{e} , the diffuse color vector \mathbf{d} and the source color vector \mathbf{s} are collinear. For these surfaces, the invariant images \mathbf{j} are zero; and as a result, they provide no information about the surface, regardless of the illuminant and sensors that are chosen. This is the same restriction noted by Klinker et al. (1988); when the diffuse and source colors are the same, there is no way to distinguish between the two reflection components.

More generally, the utility of the proposed invariants relies on the angular separation between the observed color vector \mathbf{e} and the source vectors \mathbf{s} . When this separation is small, the signal-to-noise ratio (SNR) in the invariant image can be prohibitively low. This is evident, for example, in the generalized hue image of the globe in the bottom-right of Fig. 3, where the hue variation within the People’s Republic of China is seen to be large.

Assuming independent, additive Gaussian noise with zero mean and variance σ^2 in each of the three channels of a




Fig. 5 The signal-to-noise ratio (SNR) of the two-channel diffuse image (\mathbf{j}) relative to that of the original image (\mathbf{e}) as a function of α , the angle between \mathbf{e} and the source color \mathbf{s} in the RGB cube

color vector \mathbf{e}_{RGB} , and assuming $\|\mathbf{e}_{RGB}\| \leq 1$, the signal-to-noise ratio (denoted $\text{SNR}(\mathbf{e}_{RGB})$) 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 $\|\mathbf{j}\| = \|\mathbf{e}_{RGB}\| \sin \alpha$, where α is the angle between the source color \mathbf{s} and color vector \mathbf{e}_{RGB} in color space. It follows that

$$\text{SNR}(\mathbf{j}) = \text{SNR}(\mathbf{e}_{RGB}) + 10 \log_{10}(\sin \alpha). \tag{15}$$

This relationship is shown in Fig. 5, 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.

In practice, this can be improved by increasing the SNR of the input images using multiple exposures (Grossberg and Nayar 2003). Additionally, since surface points with low SNR can be detected by monitoring the angle between the source colors \mathbf{s} and the input color vectors \mathbf{e}_{RGB} , this information can be easily incorporated into any robust vision algorithm (see, e.g., van de Weijer and Gevers 2004).

Source Color It is difficult to make general statements regarding the sensitivity of the invariants to errors in the source color estimates, because this sensitivity depends on the sensors as well as the spectral reflectances and illuminant SPDs of a particular scene. We can, however, gain some insight from the simple case of a homogeneous surface under a single illuminant. We present a qualitative description of this case here; related quantitative empirical results are presented in Sect. 5.4.2.

RGB observations of a homogeneous surface under a single illuminant lie in the dichromatic plane spanned the source and diffuse vectors \mathbf{s} and \mathbf{d} . Assuming the source vector is known, a two-channel invariant $\mathbf{j} = (j_1, j_2)$ is computed by projecting the vectors onto the subspace orthogonal to this source vector. When the estimate of the source color is inaccurate, the computed invariant also contains error. To describe sensitivity, we consider the square of the grayscale invariant $j_{\text{inv}(2)}^2 = j_1^2 + j_2^2$ and compute its derivatives with respect to angular deviations in \mathbf{s} .

Let $\{\mathbf{r}_1, \mathbf{r}_2\}$ be an orthonormal basis for the subspace orthogonal to \mathbf{s} , and choose this basis such that \mathbf{r}_1 is in the dichromatic plane. Since the observed color vectors also lie in the dichromatic plane, any one vector will have coordinates of the form $\mathbf{e} = (e_1, 0, e_3)$ in the coordinate system defined by $\mathbf{r}_1, \mathbf{r}_2$, and $\mathbf{r}_1 \times \mathbf{r}_2 = \mathbf{s}$. Thus, the squared value of the grayscale invariant is simply e_1^2 .

To describe a perturbation of the source direction, we consider a small rotation about an axis (\mathbf{a} , say) orthogonal to it. The pencil of rotation axes orthogonal to \mathbf{s} can be parameterized by the angle from \mathbf{r}_1 (so that $\mathbf{a}(\varphi) = (\cos \varphi, \sin \varphi, 0)$), and the ‘noisy’ invariant ($j_{\text{inv}(2)}^2$) that results from a rotation by angle θ about any one of these axes is

$$j_{\text{inv}(2)}^2 = \langle \mathbf{R}_{\theta, \varphi} \mathbf{r}_1, \mathbf{e} \rangle^2 + \langle \mathbf{R}_{\theta, \varphi} \mathbf{r}_2, \mathbf{e} \rangle^2. \quad (16)$$

Here, the rotation matrix $\mathbf{R}_{\theta, \varphi}$ is obtained from the axis angle representation as usual,

$$\mathbf{R}_{\theta, \varphi} = \mathbf{I} + \sin \theta [\mathbf{a}(\varphi)]_{\times} + (1 - \cos \theta) [\mathbf{a}(\varphi)]_{\times}^2,$$

where $[\cdot]_{\times}$ is the skew symmetric matrix equivalent of the cross product.

A measure of the sensitivity of the invariant is obtained by taking the derivative of (16) with respect to θ and evaluating it at $\theta = 0$:

$$\left. \frac{\partial j_{\text{inv}(2)}^2}{\partial \theta} \right|_{\theta=0} = 2e_1 e_3 \sin \varphi. \quad (17)$$

This expression reveals that the sensitivity of the invariant is highly asymmetric. When $\varphi = 0$, the rotation axis lies in the dichromatic plane, and the source vector is perturbed in a direction orthogonal to that plane. In this case, the derivative is zero and the invariant is largely unaffected by small perturbations of the source estimate. In contrast, when the source color is perturbed *within* the dichromatic plane (i.e., $\varphi = \pi/2$), the magnitude of the derivative is maximal.

For any perturbation direction (for any φ), the sensitivity is proportional to the product of the two non-zero components of the color vector \mathbf{e} . Thus, if we consider vectors of equal norm, the sensitivity is largest when the angle between the observed color vector \mathbf{e} and the source vector \mathbf{s} is 45° .

Source Color and Interreflections In cases of significant interreflection, it is possible for one surface point (\mathbf{p} , say) to specularly reflect light that is first reflected at another point. When the first reflection is diffuse, the reflected spectral radiance is modulated by the spectral reflectance of the surface, and in general, it is not spectrally equivalent to the scene illuminant SPD. Thus, with respect to \mathbf{p} the first point behaves much like a light source having a distinct SPD, and the effective source vector at point \mathbf{p} is different from \mathbf{s} . In this case, the intensity observed at \mathbf{p} does not follow the image formation model of (1) (or (12) in the mixed case), and

the proposed invariants may be contaminated by specular effects.

One method for handling interreflection effects is to locally estimate the effective source colors, and to allow these source colors to vary from point to point. As shown by Nayar et al. (1997), this can be accomplished by capturing multiple exposures from a fixed viewpoint with polarization filters at different orientations. For the purposes of this paper, however, we assume interreflection effects to be negligible so that effective source vectors are the same at every point and a single image can be used as input.

5 Applications and Evaluation

This section investigates the utility of the subspace-based invariants for a number of visual tasks and compares the results to those obtained using standard grayscale images $e = (e_R + e_G + e_B)/3$. For RGB images, when the illumination is a mixture of two known SPDs, the two-channel specular invariant j_1 from (13) is grayscale and is equal to $j_{\text{inv}(1)}$ from (14). On the other hand, a single-SPD specular invariant computed from an RGB image includes two diffuse channels $\{j_1, j_2\}$, which can be combined into a grayscale invariant $j_{\text{inv}(2)}$ using (14). In this case, one can also compute generalized hue, which can be used to replace conventional hue as a material descriptor. The results in this section show that these invariants can have advantages over conventional grayscale and hue images in the presence of specular reflections.

For the experiments in this section, the source colors are measured by imaging a Munsell color checker in an offline calibration procedure, and we focus on cases in which the diffuse and source color vectors are distinct. A quantitative investigation of the sensitivity with respect to noise in the measured source colors is provided in Sect. 5.4.2.

5.1 Binocular Stereo

The vast majority of binocular stereo algorithms are based (either explicitly or implicitly) on the assumption that surfaces are Lambertian. Since specular reflections violate this assumption, stereo reconstructions of specular surfaces are often inaccurate. The most common approach to handle specular effects in binocular stereo is to treat them as outliers. These outliers can be either explicitly detected and removed (Brelstaff and Blake 1988) or handled implicitly using robust techniques (Yoon and Kweon 2006b). Instead of treating them as outliers, one may also reduce specular reflection effects by modifying the stereo matching function to permit a more general relationship between matched regions (Kim et al. 2003). Alternatively, one can use enhanced acquisition systems that allow the effects of specularities to




Fig. 6 Stereo reconstructions under a single-color illuminant. Both conventional grayscale images and specular invariant images (7) are computed from a rectified stereo pair (top) and these are used as input to existing binocular stereo algorithms. Middle row: disparity maps obtained from the grayscale (left) and specular invariant (right) images using the method of Birchfield and Tomasi (1998). Bottom row: those obtained using the method of Boykov et al. (1998)


be reduced or eliminated. Examples include multi-view acquisition schemes (Bhat and Nayar 1998; Li et al. 2002;

Jin et al. 2005), and binocular schemes with active illumination (Zickler et al. 2003; Tu and Mendonça 2003; Davis et al. 2005b).

More directly related to the present work are reconstructions systems that address specular reflection effects using color information. One approach is to solve the (ill-posed) problem of explicit specular/diffuse separation and use the diffuse images for stereo correspondence. This is explored by Lin et al. (2002), who show that the problem can be more manageable when additional viewpoints are available. Another approach is to use stereo matching based on specular invariants. For the case of monochromatic illumination, binocular stereo using a non-linear specular invariant (which does not preserve diffuse shading information) has been explored by Yoon and Kweon (2006a), and a method that exploits color information in a multi-view system with monochromatic illumination is presented by Yang et al. (2003).

Here we investigate the use of the proposed invariants, which are based on linear transformations and are applicable in both monochromatic and mixed illumination environments. In cases of significant specular reflections and complex illumination conditions, we can improve the accuracy of existing stereo algorithms by computing these specular invariants as a pre-process. Figure 6 compares the results of two binocular stereo algorithms (Birchfield and Tomasi 1998; Boykov et al. 1998) applied to grayscale e and single-illuminant invariant $j_{inv(2)}$ images derived from a rectified RGB stereo pair. There is a dramatic improvement in the quality of reconstruction when specular invariant images are used. This point is further emphasized in Fig. 7, which compares binocular stereo results obtained using conventional grayscale images, the single-illuminant (2D subspace) invariant $j_{inv(2)}$, and the two-color (1D subspace) invariant $j_{inv(1)}$. In this case, the original RGB image includes two specular highlights caused by blue and yellow illuminants.

Fig. 7 (Color online) Stereo reconstruction under mixed illumination. Top left: One image of an input stereo pair with blue and yellow illumination. Top center: Single-color invariant image $j_{inv(2)}$ from (13) and (14) with s in the direction of the blue source. Top right: Two-color invariant $j_{inv(1)}$ obtained by projecting to the 1D subspace orthogonal to both sources. Bottom row: depth map obtained using the stereo algorithm of Boykov et al. (1998) in each case



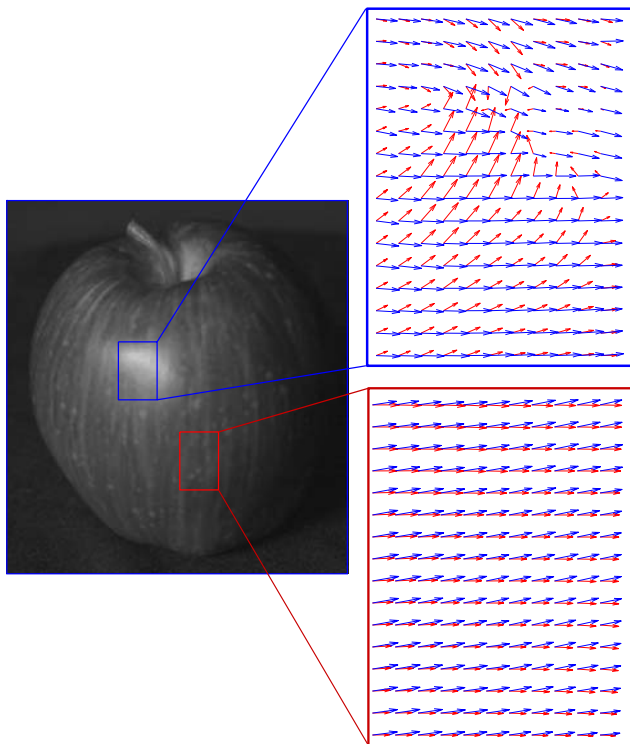


Fig. 8 (Color online) Optical flow under a single-color illuminant. An RGB image sequence is captured by a camera translating left relative to a specular apple. Both conventional grayscale and specular invariant images (7) are computed from this RGB sequence, and these are used as input to a robust optical flow algorithm (Black and Anandan 1993). *Left*: Single frame from the grayscale sequence. *Right*: flows obtained for regions that are highly specular and predominantly diffuse. *Red* flow is computed from the grayscale sequence and is severely corrupted by specular reflection. *Blue* flow is computed from the specular invariant sequence and is much closer to ground truth, which is horizontal and to the right

The blue highlight is largely eliminated in the single-color invariant $j_{inv(2)}$, while image $j_{inv(1)}$ is invariant to specular reflections of both colors. As expected, the results from the grayscale and single-color invariant images are poor in specular regions, and the depth map obtained using $j_{inv(1)}$ is significantly improved.

5.2 Optical Flow

Motion estimation through the computation of optical flow is another example of an application that can benefit from specular invariance. Recovering dense optical flow relies on the ‘constant-brightness assumption’, which is violated when an observer moves relative to a static, specular scene.

As is the case with stereo, existing work has shown that color information can be exploited to deal with violations of the constant-brightness assumption (see, Barron and Klette 2002, for a survey). Most existing algorithms exploit color by computing either a shading invariant (e.g., normalized RGB) or a white-illuminant specular invariant (e.g., hue) as

a pre-process, and studies have shown that these can provide improved estimates of the optical flow field.

We approach the problem of optical flow in a similar spirit using the invariants defined in Sect. 4, which have the advantages of handling non-white illuminants and mixed lighting environments. Figure 8 shows a comparison of optical flow estimation in the presence of specular reflections under a single-color illuminant. An RGB image sequence is captured by a camera translating horizontally relative to a static scene. The sequence is used to compute a conventional grayscale sequence $e(t)$ and a single-color invariant sequence $j_{inv(2)}(t)$, and these are used as input to a robust optical flow algorithm (Black and Anandan 1993). Since the camera undergoes pure translation, the ‘ground truth’ flow lies along parallel horizontal lines. As the figure shows, in regions that are predominantly diffuse, the flow obtained in both cases is close to the ground truth. In regions of specularity, however, there is a significant improvement in the quality of estimated flow when specular invariant images are used.

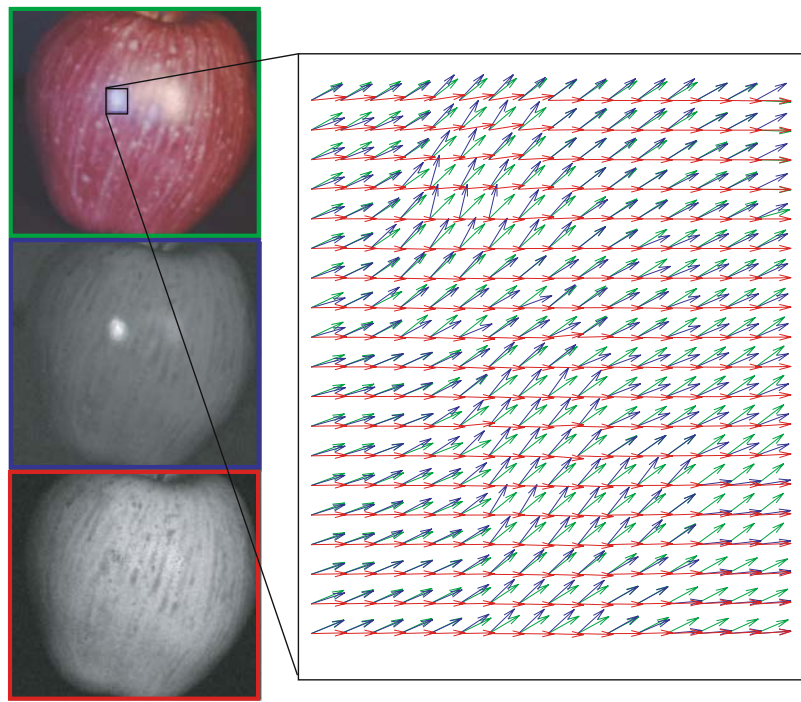
More interesting is the case of optical flow estimation under mixed illumination, which is shown in Fig. 9. A similar sequence is captured under illumination that is a mixture of two distinct SPDs, and the sequence is used to compute a conventional grayscale sequence $e(t)$, a single-color invariant sequence $j_{inv(2)}(t)$, and a two-color invariant sequence $j_{inv(1)}(t)$. These three videos are used as input to the same optical flow algorithm (Black and Anandan 1993). The left of Fig. 9 shows a single image from each sequence, and the right shows the recovered flows in the indicated window. The flow recovered using the conventional grayscale and single-color invariant sequences are severely corrupted by specular highlights. In contrast, the flow computed from the mixed-illuminant invariant (shown in red) is close to the ground truth and is largely unaffected by these non-Lambertian effects.

5.3 Shape from Shading

The previous two sections demonstrate the utility of the specular invariant for stereo matching and optical flow, both of which benefit from the fact that the specular invariant images do not change with viewpoint. The next three sections show that since they preserve diffuse (ideally Lambertian) shading information, these invariants can also be used to enhance photometric reconstruction methods.

In shape from shading, one seeks to recover surface shape from the photometric information available in a single image. The vast majority of the existing methods assume Lambertian reflectance, and even then the problem is a difficult one. Of the small number of methods that consider non-Lambertian effects, most assume reflectance to be of a specific parametric form—such as the Torrance-Sparrow or Oren-Nayar models—which must be known

Fig. 9 (Color online) Optical flow under mixed illumination. An RGB image sequence (top left) is captured by a camera translating left relative to a specular apple under yellow and blue illumination. Derived conventional grayscale $e(t)$, yellow-invariant $j_{inv(2)}(t)$ (left middle), and two-color invariant $j_{inv(1)}(t)$ (left bottom) sequences are computed and used as input to a robust optical flow algorithm (Black and Anandan 1993). Right: flows obtained in the three cases. Green and blue flows are from grayscale and yellow-invariant sequences, respectively, and both are corrupted by specular reflections. Red flow is computed from the two-color invariant and is much closer to ground truth, which is horizontal and to the right



a priori (Ahmed and Farag 2006; Bakshi and Yang 1994; Ragheb and Hancock 2001). The use of color in shape from shading is rare. One notable example is the work of Tian and Tsui (1997), which considers reflectance that is a linear combination of a Lambertian diffuse component and an ideal specular spike.

The invariants presented in Sect. 4 provide a means for considering a much broader class of surfaces. By combining these invariants with existing Lambertian-based methods for shape from shading, one can recover shape for surfaces having rather arbitrary specular components (i.e., general $f_s(\theta)$ in (10)) which need not be well-represented by any known parametric form. All that is required is that the surface conforms to the dichromatic model.

When illumination can be described as a single point source in direction \mathbf{l} (say) and the diffuse reflectance at a surface point is Lambertian, we can write $\sigma_d = f_d \mathbf{n}^\top \mathbf{l}$, where \mathbf{n} is the surface normal at the point and f_d is the albedo. When this is true, the specular invariant image of (14) reduces to

$$j_{inv(2)} = f_d ((\mathbf{r}_1^\top \mathbf{d})^2 + (\mathbf{r}_2^\top \mathbf{d})^2)^{\frac{1}{2}} \mathbf{n}^\top \mathbf{l}, \quad (18)$$

which is the image formation equation for a Lambertian surface with an effective albedo given by the first two terms. Thus, the specular invariant can be used directly as input to any Lambertian-based shape from shading algorithm.

The benefit of this approach is demonstrated in Fig. 10, where we assess the performance of a conventional shape from shading algorithm (Zheng and Chellappa 1991) for both a grayscale image e and a single-SPD invariant image $j_{inv(2)}$. The top of the figure shows grayscale and spec-

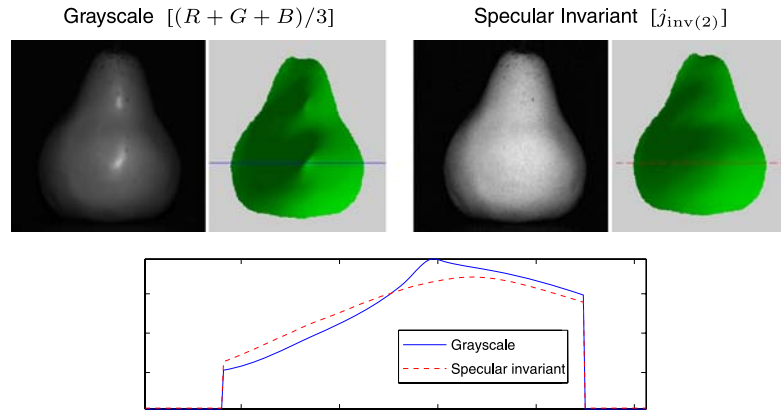
ular invariant images computed from an RGB image of a pear, and the middle row shows the surfaces that are recovered by applying the same algorithm in the two cases. The solid blue profile in the bottom graph shows that specular reflections cause severe artifacts when the algorithm is applied to the grayscale image. In contrast, as shown by the dashed red profile, one can obtain improved results using the same algorithm by computing the specular invariant as a pre-processing step.

5.4 Photometric Stereo

In photometric stereo, one seeks to recover shape from a set of images acquired from a fixed viewpoint under multiple illumination conditions. Like shape from shading, photometric stereo requires the inversion of the image formation process, and as a result, existing methods also require significant knowledge about the reflectance of surfaces in a scene. Many photometric stereo techniques assume that surfaces are Lambertian (Woodham 1978), and others assume the reflectance to be given *a priori* by a reference object (Silver 1980), a linear basis of reference objects (Hertzmann and Seitz 2003), or by a parametric BRDF model (Ikeuchi 1981; Nayar et al. 1990; Tagare and deFigueiredo 1991). When these reflectance assumptions are not satisfied, the accuracy of the recovered shape can be compromised.

Coleman and Jain (1982) 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

Fig. 10 (Color online) Shape from shading comparison. *Top*: An RGB image of a pear is used to compute conventional grayscale (*left*) and specular invariant (*right*) images, and these are input to a shape from shading algorithm (Zheng and Chellappa 1991) yielding the surfaces shown in *green*. *Bottom row*: cross-sections of the recovered surfaces along the indicated *horizontal lines*



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 (2003) 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 explicitly separate the diffuse and specular components as a pre-processing step. This is the approach taken by Schlüns and Wittig (1993), who assume homogeneous dichromatic surfaces, and separate the diffuse and specular components using color histogram analysis techniques similar to Klinker et al. (1988). Sato and Ikeuchi (1994) 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 methods 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 explicit specular/diffuse separation, however, they are subject to the restrictions discussed in Sect. 2. 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 (1982).

By using the invariants from Sect. 4 in conjunction with existing Lambertian-based methods for photometric stereo, many of these limitations can be overcome. In fact, this provides a reconstruction method that operates completely independent of specular reflections (i.e., independent of

$f_s(\theta)$ in (10)) and therefore requires no additional assumptions regarding the specular behavior of a surface. In this sense, this approach to photometric stereo is related to other recent reconstruction methods that exploit physical properties such as reflectance isotropy (Lu and Little 1999), reciprocity (Magda et al. 2001; Zickler et al. 2002), the constancy of radiance in free space (Magda et al. 2001; Koudelka et al. 2001), and light transport constancy (Davis et al. 2005a) to enable accurate reconstructions of very broad classes of surfaces. An important difference, however, is that the photometric stereo method described here requires a simple acquisition system and is quite easy to implement.

To use the proposed invariants for photometric stereo, we assume directional monochromatic illumination as in the previous section. Let $\mathbf{j}^1, \mathbf{j}^2, \mathbf{j}^3$ be three two-channel color vectors produced by observing a single point under three different lighting directions $\mathbf{l}^1, \mathbf{l}^2, \mathbf{l}^3$, and computing specular invariants according to (13). Assuming Lambertian diffuse reflectance, we see that

$$\mathbf{j}^k = [j_1^k, j_2^k]^T = (\mathbf{n}^T \mathbf{l}^k) \boldsymbol{\rho}, \tag{19}$$

with

$$\boldsymbol{\rho} = [\rho_1, \rho_2]^T = f_d[\mathbf{r}_1^T \mathbf{d}, \mathbf{r}_2^T \mathbf{d}]^T$$

being an effective two-channel albedo, and it follows that these specular invariant images can be used as input to a Lambertian photometric stereo algorithm. In what follows, we adapt the algorithm of Barsky and Petrou (2001) that was originally designed to handle RGB images of Lambertian scenes.

Similar to Barsky and Petrou (2001), a *shading vector* is defined as $\mathbf{h} = [h^1, h^2, h^3]^T = [\mathbf{l}^1 \ \mathbf{l}^2 \ \mathbf{l}^3]^T \mathbf{n}$, and the invariant images resulting from the three lighting directions are combined in an *intensity matrix* satisfying

$$\mathbf{J} = \begin{bmatrix} j_1^1 & j_2^1 \\ j_1^2 & j_2^2 \\ j_1^3 & j_2^3 \end{bmatrix} = \begin{bmatrix} h^1 \rho_1 & h^1 \rho_2 \\ h^2 \rho_1 & h^2 \rho_2 \\ h^3 \rho_1 & h^3 \rho_2 \end{bmatrix} = \mathbf{h} \boldsymbol{\rho}^T. \tag{20}$$

Fig. 11 Photometric stereo procedure. Three or more RGB images are acquired under known illumination conditions, and specular invariants \mathbf{j} are computed according to (13). The invariants 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


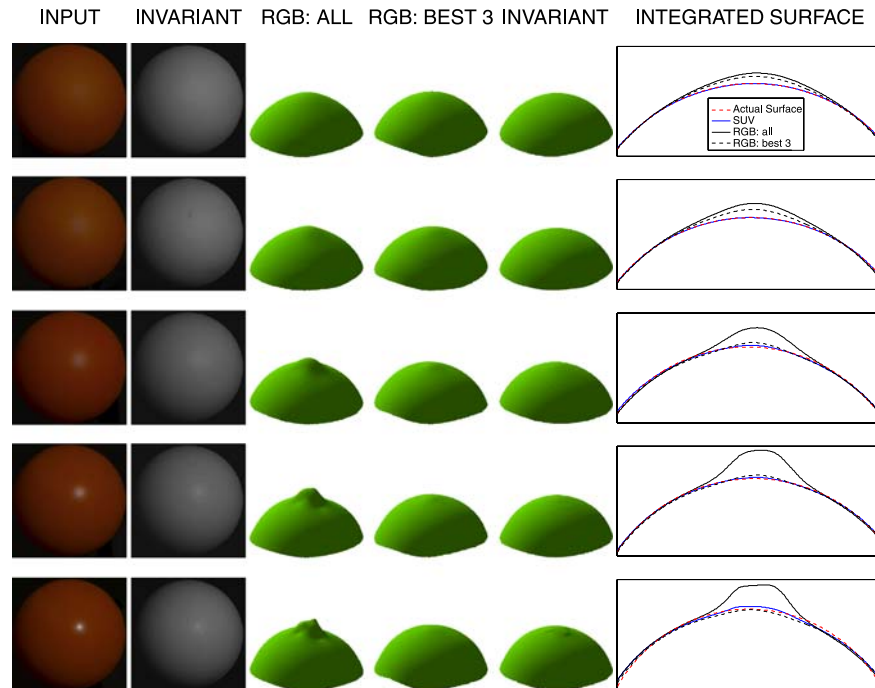


Fig. 12 (Color online) Comparison of photometric stereo methods. Five red spheres with increasing specular reflectance are each observed under four illumination directions, and these images are used to recover the surface. From left to right, each row shows: (i) an input RGB image, (ii) the corresponding specular invariant \mathbf{j} from (13), (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



The least-squares estimate of the shading vector \mathbf{h} is computed from the intensity matrix; it is the principal eigenvector of $\mathbf{J}\mathbf{J}^T$. Once the shading vector is determined, the surface normal is found by solving the matrix equation $\mathbf{h} = [\mathbf{l}^1 \ \mathbf{l}^2 \ \mathbf{l}^3]^T \mathbf{n}$. This reconstruction procedure is outlined in Fig. 11, and it can be applied without change to any number of images larger than two.

5.4.1 Experimental Results

Photometric stereo provides a convenient means for quantitative analysis of the proposed invariant, since we can directly measure the accuracy of reconstructed shapes having different material properties. To perform such an analysis, we painted five identical spheres, shown in Fig. 12, with standard latex paints that were mixed to have approximately the same color pigment and five different levels of glossy finish: flat, eggshell, satin, semi-gloss, and high-gloss. The

observed incident-plane BRDFs of these spheres are shown in Fig. 13.

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 used to compute the specular invariants \mathbf{j} and $j_{inv(2)}$ according to (13) and (14). The second column of Fig. 12 confirms that the specular invariant $j_{inv(2)}$ depends largely on the diffuse reflectance.

Using the two-channel specular invariant images, the surface normals of each sphere were estimated using the photometric stereo method described above. As a means of comparison, we implemented two alternative RGB-based photometric techniques. The first method uses all four RGB images and assumes Lambertian reflectance (Barsky and Petrou 2001). The second method assumes Lambertian+specular reflectance and reconstructs the surface by choosing the three ‘least specular’ RGB measurements

Fig. 13 Comparison of photometric stereo methods. *Left:* Relative BRDFs (in decibels) of the five red spheres of Fig. 12 as a function of half-angle. *Right:* Mean-square angular error in the recovered surface normals as a function of increasing specularity using both the proposed specular invariants and existing RGB methods





Fig. 14 Invariant-based photometric stereo applied to natural surfaces. *Left:* Input RGB images show significant specular reflectance and texture. By computing the specular invariant, the specular effects are removed, enabling accurate recovery of shape. *Middle,* *Right:* Surfaces recovered by integrating the estimated surface normals



at each pixel (Barsky and Petrou 2003; Coleman and Jain 1982). The results are shown in Figs. 12 and 13. The recovered surfaces, including cross-sections overlaid on the true shape, are displayed in Fig. 12. Quantitative results are shown in Fig. 13, with the right of that figure displaying the angular difference between the true and estimated surface normals as a function of increasing specularity. These results demonstrate that the invariant-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 (Barsky and Petrou 2001) assumes Lambertian reflectance and its performance degrades monotonically as gloss increases; and while the three-image RGB method (Barsky and Petrou 2003; Coleman and Jain 1982) 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 14 shows the results of applying the invariant-based photometric stereo method to two natural objects (a pear and a pumpkin). Since the computation of the specular invariant is purely local, the method requires no spatial coherence in the image, and it performs well for surfaces with arbitrary texture. This is not true for alterna-

tive photometric stereo techniques that rely on explicit diffuse/specular separation (e.g., Schlüns and Wittig 1993), since these methods generally require some form of spatial coherence in the spectral reflectance of a surface.

5.4.2 Sensitivity to Illuminant Color

Photometric stereo also provides an opportunity to quantitatively evaluate the sensitivity of the proposed invariants to perturbations in the measured illuminant color. This complements the qualitative analysis presented in Sect. 4.2. To measure sensitivity, we repeated the photometric reconstruction procedure in Fig. 13 using invariants computed with perturbed source vectors. When the source vector is perturbed from its true value, the specular invariant images are contaminated by specular effects, and the reconstruction error in the Lambertian-based photometric stereo result is expected to increase.

Figure 15 shows the result of this experiment using the red sphere from the second row of Fig. 12. Depicted is the angular mean-square reconstruction error (in degrees) resulting from perturbations of the unit source vector. Since source vectors are of unit length, the domain of the error function is the unit sphere, and the figure shows the stereographic projection of this error function centered at the true

source color (indicated by +). Concentric circles in Fig. 15 correspond to angular source perturbations of 5° , 10° and 15° , and the diagonal black line is the projection of the dichromatic plane, which is the plane spanned by the diffuse vector of the homogeneous surface and the true source vector.

The qualitative analysis from Sect. 4.2 reveals that the specular invariant is more sensitive to source perturbations within the dichromatic plane than it is to perturbations away from the plane. This effect is also observed in Fig. 15, where a 10° perturbation within the plane causes the error to increase by nearly a factor of two, while the same angular perturbation in the orthogonal direction induces only a 25% increase.

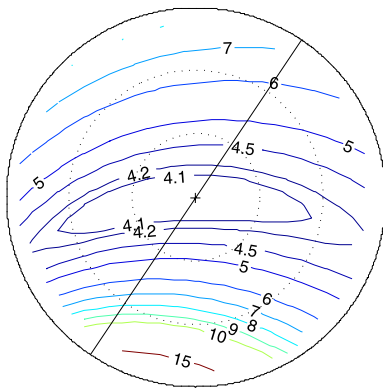


Fig. 15 Sensitivity of a photometric stereo reconstruction with respect to errors in estimated source color. The field of surface normals of a sphere from Fig. 12 is recovered using invariant-based photometric stereo with source vectors \mathbf{s} perturbed from truth, and the angular MSE in the normal field is recorded. Shown is a contour plot of the stereographic projection of this error (in degrees) as a function of the angular perturbation to the source vector. Concentric circles are cones of source vectors displaced by 5° , 10° and 15° from the true vector (+), and the diagonal line is the projection of the dichromatic plane for this homogeneous surface. The angular MSE for the true source vector is 3.98° . The reconstruction is more sensitive to source perturbations within the dichromatic plane than those orthogonal to it

While this experiment provides some insight into the sensitivity of the proposed invariants, one must be cautious about the conclusions one draws. Since photometric stereo is an active illumination technique, one typically has the opportunity to directly measure the source color. When this is the case, the noise in the source estimate will be much smaller than the 15° error considered here. For other applications (stereo, optical flow, etc.) in which the source color is difficult to measure or is time-varying, one would need to rely on existing image-based methods for illuminant estimation as discussed in Sect. 2. Angular source errors may be larger in this case—some empirical studies suggest that errors of 10° are not uncommon (Barnard et al. 2002a, 2002b)—and these errors will depend very strongly on the particular materials and illuminants that are present in the scene, the sensors being used, and the illuminant estimation algorithm(s) being employed.

5.5 Photometric/Geometric Reconstruction

In addition to the applications presented thus far, the specular invariant can be used to improve the performance of a broad class of Lambertian-based reconstruction systems in the presence of specular, non-Lambertian surfaces. This includes methods that combine both geometric and photometric constraints to obtain accurate surface shape (Jin et al. 2004; Lim et al. 2005; Zhang et al. 2003). To provide an example, we use the passive photometric stereo algorithm described by Lim et al. (2005). This method begins with an approximate, piece-wise planar reconstruction obtained by tracking a small number of features across a video sequence under (possibly varying) directional illumination. Then, an iterative method based on uncalibrated Lambertian photometric stereo simultaneously refines the reconstruction and estimates the unknown illumination directions.

Figure 16 compares the results obtained from an image sequence that consists of a moderately specular cylin-

Fig. 16 Comparison of shape from combined photometric and geometric constraints. *Left*: three RGB frames of a specular cylinder moving under fixed view and illumination. *Right frame*: result of simultaneous tracking and photometric reconstruction (Lim et al. 2005) using both the conventional grayscale (*left*) and specular invariant (*right*) sequences





Fig. 17 Generalized hue for material-based segmentation. Each panel shows a pseudo-colored representation that is computed from the RGB image on the *top-left*. The generalized hue image on the *bottom-right* is useful for segmentation because it depends only on the spectral reflectance of the surfaces. The same is not true for a conventional hue image (*bottom-left*) unless the illuminant is white



der moving under fixed illumination and viewpoint. The shape is estimated by applying the same algorithm to both the conventional grayscale sequence ($e(t)$) and the specular invariant sequence ($j_{\text{inv}(2)}$) computed from the same RGB data. The right-most surface in Fig. 16 shows that the reconstruction obtained using the specular invariant is nearly cylindrical, while that computed from the conventional grayscale sequence is severely corrupted by specular reflections.

5.6 Material-Based Segmentation

Sections 5.1–5.5 demonstrate the utility of the proposed specular invariants for a variety of visual tasks. This section demonstrates an applications of the second invariant, generalized hue, which is independent of both the specular reflections and diffuse shading in an image. We consider its application to the problem of material-based segmentation, although other potential applications include lighting-insensitive tracking and recognition.

Figure 17 shows an RGB image of a dichromatic scene under uniform source color ($N = 1$) along with a series of pseudo-colored representations related to the invariants presented in Sect. 4. The top row shows conventional grayscale and specular invariant images, and in the latter, the specular effects (most notably on the green apples, the pumpkin, and the red pepper) are largely eliminated. The bottom-right of Fig. 17 shows the generalized hue image given by (8), which is invariant to diffuse shading in addition to specular reflections, and therefore depends only on the spectral reflectance. The fact that the generalized hue within each region is relatively constant suggests that it is a useful representation for segmentation. The same is not true for the

conventional hue image (shown on the bottom-left) because the illuminant is not white.

6 Conclusion

This paper presents photometric invariants that are derived from color subspaces. They can be efficiently computed from a single image of a dichromatic scene and can be applied in cases of both monochromatic and mixed illumination environments. Two important features of these invariants are that: (1) they are free of specular reflectance effects; and (2) they preserve the diffuse shading information in an image. The latter means that they can be used directly for Lambertian-based photometric analysis including shape from shading and photometric stereo.

The invariants are computed point-wise and therefore place no restriction on scene texture. Additionally, while they require knowledge of the effective source color(s), they place no restrictions on the angular distribution of incident light.

The utility of these invariants is demonstrated by their ability to improve the performance of a wide variety of vision algorithms, including those for binocular stereo, motion estimation, and photometric reconstruction. They are directly applicable in cases where the source color is measured or known, and in these cases, they are shown to allow many Lambertian-based algorithms to be applied more successfully to a much broader class of surfaces.

An important next step is to explore applications in uncontrolled environments, where illumination spectra cannot be measured or are time varying. By combining the proposed invariants with existing methods for illuminant estimation and robust Lambertian-based vision algorithms, they may prove to be useful in these cases as well.

References

- Ahmed, A., & Farag, A. (2006). A new formulation for shape from shading for non-lambertian surfaces. In *Proceedings of IEEE conference on computer vision and pattern recognition* (pp. 1817–1824).
- Bakshi, S., & Yang, Y.-H. (1994). Shape from shading for non-lambertian surfaces. In *Proceedings of IEEE international conference on image processing* (Vol. 2, pp. 130–134).
- Barnard, K., Cardei, V., & Funt, B. (2002a). A comparison of computational color constancy algorithms. I: Methodology and experiments with synthesized data. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 11(9).
- Barnard, K., Martin, L., Coath, A., & Funt, B. (2002b). A comparison of computational color constancy algorithms. II: Experiments with image data. *IEEE Transactions on Image Processing*, 11(9), 985–996.
- Barron, J. L., & Klette, R. (2002). Quantitative color optical flow. In *Proceedings international conference on pattern recognition* (Vol. 4, pp. 251–255). Washington: IEEE Computer Society.
- Barsky, S., & Petrou, M. (2001). Colour photometric stereo: Simultaneous reconstruction of local gradient and colour of rough textured surfaces. In *Proceedings of IEEE international conference on computer vision* (pp. 600–605).
- Barsky, S., & Petrou, M. (2003). 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(10), 1239–1252.
- Bhat, D., & Nayar, S. (1998). Stereo and specular reflection. *International Journal of Computer Vision*, 26(2), 91–106.
- Birchfield, S., & Tomasi, C. (1998). Depth discontinuities by pixel-to-pixel stereo. In *Proceedings of IEEE international conference on computer vision* (pp. 1073–1080).
- Black, M. J., & Anandan, P. (1993). A framework for the robust estimation of optical flow. In *Proceedings of IEEE international conference on computer vision* (pp. 231–236).
- Blanz, V., & Vetter, T. (2003). Face recognition based on fitting a 3D morphable model. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 25(9).
- Boykov, Y., Veksler, O., & Zabih, R. (1998). Markov random fields with efficient approximations. In *Proceedings of IEEE conference on computer vision and pattern recognition* (pp. 648–655).
- Brainard, D. H., & Freeman, W. T. (1997). Bayesian color constancy. *Journal of Optical Society of America A*, 14, 1393–1411.
- Brelstaff, G., & Blake, A. (1988). Detecting specular reflection using lambertian constraints. In *Proceedings of IEEE international conference on computer vision* (pp. 297–302).
- Coleman, E., & Jain, R. (1982). Obtaining 3-dimensional shape of textured and specular surfaces using four-source photometry. *Computer Vision, Graphics and Image Processing*, 18(4), 309–328.
- Davis, J., Yang, R., & Wang, L. (2005a). BRDF invariant stereo using light transport constancy. In *Proceedings of IEEE international conference on computer vision* (Vol. 1).
- Davis, J. E., Yang, R., & Wang, L. (2005b). BRDF invariant stereo using light transport constancy. In *ICCV '05: Proceedings of the tenth IEEE international conference on computer vision (ICCV'05)* (Vol. 1, pp. 436–443). Washington: IEEE Computer Society.
- Finlayson, G. D. (1996). Color in perspective. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 18, 1034–1036.
- Finlayson, G., & Schaefer, G. (2001). Constrained dichromatic colour constancy. In *Proceedings of European conference on computer vision* (Vol. 1, pp. 342–358).
- Finlayson, G. D., Hordley, S. D., & Hubel, P. M. (2001). Color by correlation: A simple, unifying framework for color constancy. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 23, 1209–1221.
- Gershon, R. (1987). *The use of color in computational vision*. Ph.D. thesis, University of Toronto.
- Grossberg, M. D., & Nayar, S. K. (2003). High dynamic range from multiple images: Which exposures to combine? In *Proceedings of IEEE workshop on color and photometric methods in computer vision (CPMVC)*.
- Healey, G. (1989). Using color for geometry-insensitive segmentation. *Journal of Optical Society of America A*, 6(6), 920–937.
- Hertzmann, A., & Seitz, S. (2003). Shape and material by example: a photometric stereo approach. In *Proceedings of IEEE conference on computer vision and pattern recognition*.
- Hordley, S., & Finlayson, G. (2006). Reevaluation of color constancy algorithm performance. *Journal of Optical Society of America A*, 23(5), 1008–1020.
- Hordley, S. D., Finlayson, G. D., & Drew, M. S. (2002). Removing shadows from images. In *Proceedings of European conference on computer vision* (pp. 823–836).
- Ikeuchi, K. (1981). Determining surface orientations of specular surfaces by using the photometric stereo method. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 3(6), 661–669.
- Jin, H., Cremers, D., Yezzi, A., & Soatto, S. (2004). Shedding light on stereoscopic segmentation. In *Proceedings of IEEE conference on computer vision and pattern recognition*.
- Jin, H., Soatto, S., & Yezzi, A. J. (2005). Multi-view stereo reconstruction of dense shape and complex appearance. *International Journal of Computer Vision*, 63(3), 175–189.
- Kim, J., Kolmogorov, V., & Zabih, R. (2003). Visual correspondence using energy minimization and mutual information.
- Klinker, G., Shafer, S., & Kanade, T. (1988). The measurement of highlights in color images. *International Journal of Computer Vision*, 2(1), 7–32.
- Koudelka, M., Magda, S., Belhumeur, P., & Kriegman, D. (2001). Image-based modeling and rendering of surfaces with arbitrary BRDFs. In *Proceedings of IEEE conference on computer vision and pattern recognition* (pp. 568–575).
- Lee, H.-S. (1986). Method for computing the scene-illuminant chromaticity from specular highlights. *Journal of Optical Society of America A*, 3(10), 1694–1699.
- Lee, H. C., Breneman, E. J., & Schulte, C. P. (1990). Modeling light reflection for computer color vision. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 12(4), 402–409.
- Lehmann, T. M., & Palm, C. (2001). Color line search for illuminant estimation in real-world scenes. *Journal of Optical Society of America A*, 18(11), 2679–2691.
- Li, Y., Lin, S., Lu, H., Kang, S. B., & Shum, H.-Y. (2002). Multi-baseline stereo in the presence of specular reflections. In *ICPR '02: Proceedings of the 16th international conference on pattern recognition (ICPR'02)* (Vol. 3, p. 30573). Washington: IEEE Computer Society.
- Lim, J., Ho, J., Yang, M.-H., & Kriegman, D. (2005). Passive photometric stereo from motion. In *Proceedings of IEEE international conference on computer vision*.
- Lin, S., Li, Y., Kang, S. B., Tong, X., & Shum, H.-Y. (2002). Diffuse-specular separation and depth recovery from image sequences. In *ECCV '02: proceedings of the 7th European conference on computer vision—Part III* (pp. 210–224). London: Springer.
- Lu, J., & Little, J. (1999). Reflectance and shape from images using a collinear light source. *International Journal of Computer Vision*, 32(3), 1–28.
- Magda, S., Kriegman, D., Zickler, T., & Belhumeur, P. (2001). Beyond Lambert: Reconstructing surfaces with arbitrary BRDFs. In *Proceedings of IEEE international conference on computer vision* (pp. 391–398).
- Mallick, S. P., Zickler, T. E., Belhumeur, P. N., & Kriegman, D. J. (2006). Specularity removal in images and videos: A PDE approach. In *Proceedings of European conference on computer vision*.

- Narasimhan, S. G., Ramesh, V., & Nayar, S. K. (2003). A class of photometric invariants: Separating material from shape and illumination. In *Proceedings of IEEE international conference on computer vision* (Vol. 2, pp. 1387–1394).
- Nayar, S. K., & Bolle, M. (1996). Reflectance based object recognition. *International Journal of Computer Vision*, 17(3), 219–240.
- Nayar, S., Ikeuchi, K., & Kanade, T. (1990). Determining shape and reflectance of hybrid surfaces by photometric sampling. *IEEE Journal of Robotics and Automation*, 6(4), 418–431.
- Nayar, S., Fang, X., & Boulton, T. (1997). Separation of reflection components using color and polarization. *International Journal of Computer Vision*, 21(3), 163–186.
- Park, J. B. (2003). Efficient color representation for image segmentation under nonwhite illumination. In *SPIE* (Vol. 5267, pp. 163–174).
- Ragheb, H., & Hancock, E. (2001). Separating lambertian and specular reflectance components using iterated conditional modes. In *Proceedings of British machine vision conference* (pp. 541–522).
- Rosenberg, C., Hebert, M., & Thrun, S. (2001). Color constancy using KL-divergence. In *Proceedings of IEEE international conference on computer vision* (pp. 239–247).
- Sapiro, G. (1999). Color and illumination voting. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 21, 1210–1215.
- Sato, Y., & Ikeuchi, K. (1994). Temporal-color space analysis of reflection. *Journal of Optical Society of America A*, 11(11), 2990–3002.
- Schlüns, K., & Wittig, O. (1993). Photometric stereo for non-Lambertian surfaces using color information. In *Proceedings of international conference on image analysis and processing* (pp. 505–512).
- Shafer, S. (1985). Using color to separate reflection components. *COLOR Research and Applications*, 10(4), 210–218.
- Silver, W. (1980). *Determining shape and reflectance using multiple images*. Master's thesis, MIT.
- Tagare, H., & deFigueiredo, R. (1991). A theory of photometric stereo for a class of diffuse non-lambertian surfaces. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 13(2), 133–152.
- Tan, P., Lin, S., & Quan, L. (2006). Separation of highlight reflections on textured surfaces. In *Proceedings of IEEE conference on computer vision and pattern recognition*.
- Tan, R. T., & Ikeuchi, K. (2003). Separating reflection components of textured surface using a single image. In *Proceedings of IEEE international conference on computer vision* (pp. 870–877).
- Tan, R. T., Nishino, K., & Ikeuchi, K. (2004). Color constancy through inverse-intensity chromaticity space. *Journal of Optical Society of America A*, 21(3), 321–334.
- Tian, Y., & Tsui, H. (1997). Shape recovery from a color image for non-lambertian surfaces. *Journal of Optical Society of America A*, 14(2), 397–404.
- Tominga, S., & Wandell, B. (1989). Standard surface-reflectance model and illuminant estimation. *Journal of Optical Society of America A*, 6(4), 576–584.
- Tominga, S., & Wandell, B. A. (2002). Natural scene-illuminant estimation using sensor correlation. *Proceedings of IEEE*, 90, 42–56.
- Tsumura, N., Ojima, N., Sato, K., Shiraishi, M., Shimizu, H., Nabeshima, H., Akazaki, S., Hori, K., & Miyake, Y. (2003). Image-based skin color and texture analysis/synthesis by extracting hemoglobin and melanin information in the skin. In *Proceedings of ACM SIGGRAPH* (pp. 770–779).
- Tu, P., & Mendonça, P. (2003). Surface reconstruction via Helmholtz reciprocity with a single image pair. In *Proceedings of IEEE conference on computer vision and pattern recognition* (Vol. 1, pp. 541–547).
- van de Weijer, J., & Gevers, T. (2004). Robust optical flow from photometric invariants. In *Proceedings of IEEE international conference on image processing* (pp. 1835–1838).
- Wann Jensen, H., Marschner, S., Levoy, M., & Hanrahan, P. (2001). A practical model for subsurface light transport. In *Proceedings of ACM SIGGRAPH* (pp. 511–518).
- Wolff, L., & Angelopoulos, E. (1994). Three-dimensional stereo by photometric ratios. *Journal of Optical Society of America A*, 11, 3069–3078.
- Wolff, L. B., & Boulton, T. E. (1991). Constraining object features using a polarization reflectance model. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 13(7), 635–657.
- Woodham, R. (1978). Photometric stereo: A reflectance map technique for determining surface orientation from image intensity. In *Proceedings of SPIE* (Vol. 155, pp. 136–143).
- Yang, R., Pollefeys, M., & Welch, G. (2003). Dealing with textureless regions and specular highlights—a progressive space carving scheme using a novel photo-consistency measure. In *ICCV '03: Proceedings of the ninth IEEE international conference on computer vision* (p. 576). Washington: IEEE Computer Society.
- Yoon, K., & Kweon, I. (2006a). Correspondence search in the presence of specular highlights using specular-free two-band images. In *Proceedings of Asian conference on computer vision* (pp. 761–770).
- Yoon, K.-J., & Kweon, I. S. (2006b). Adaptive support-weight approach for correspondence search. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 28(4), 650–656.
- Zhang, L., Curless, B., Hertzmann, A., & Seitz, S. M. (2003). Shape and motion under varying illumination: Unifying structure from motion, photometric stereo, and multi-view stereo. In *Proceedings of IEEE international conference on computer vision* (pp. 618–625).
- Zheng, Q., & Chellappa, R. (1991). Estimation of illuminant direction, albedo, and shape from shading. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 13(7), 680–702.
- Zickler, T., Belhumeur, P., & Kriegman, D. (2002). Helmholtz stereopsis: Exploiting reciprocity for surface reconstruction. In *Proceedings of European conference on computer vision* (pp. 869–884).
- Zickler, T. E., Ho, J., Kriegman, D. J., Ponce, J., & Belhumeur, P. N. (2003). Binocular helmholtz stereopsis. In *Proceedings of IEEE international conference on computer vision* (p. 1411). Washington: IEEE Computer Society.