

A Class of Photometric Invariants: Separating Material from Shape and Illumination *

Srinivasa G. Narasimhan¹, Visvanathan Ramesh² and Shree K. Nayar¹

Department of Computer Science¹
Columbia University, New York 10027
{srinivas, nayar}@cs.columbia.edu

Real-time Vision and Modeling Department²
Siemens Corporate Research, Princeton 08540
rameshv@scr.siemens.com

Abstract

We derive a new class of photometric invariants that can be used for a variety of vision tasks including lighting invariant material segmentation, change detection and tracking, as well as material invariant shape recognition. The key idea is the formulation of a scene radiance model for the class of “separable” BRDFs, that can be decomposed into material related terms and object shape and lighting related terms. All the proposed invariants are simple rational functions of the appearance parameters (say, material or shape and lighting). The invariants in this class differ from one another in the number and type of image measurements they require. Most of the invariants in this class need changes in illumination or object position between image acquisitions. The invariants can handle large changes in lighting which pose problems for most existing vision algorithms. We demonstrate the power of these invariants using scenes with complex shapes, materials, textures, shadows and specularities.

1 Invariants in Vision

Appearances of scenes depend on a variety of factors such as lighting geometry and spectrum, scene structure and material properties, medium in which light travels, viewing geometry and sensor properties. Most often, these parameters combine non-linearly to yield an image. Recovering these factors from images is an important problem in vision. Direct estimation of these parameters from a set of images of a scene, however, is generally hard. Photometric invariants provide an intermediate solution to this problem.

Invariants usually transform images into a simpler feature space where more accurate algorithms can be developed for the task at hand. To be effective, invariants must satisfy two properties - (a) they must be invariant to certain appearance parameters (say, lighting) and (b) they must have good discriminability with respect to other parameters (say, material properties). Good invariants can be effective for common vision tasks such as illumination invariant recognition, material segmentation and lighting insensitive tracking.

There has been a lot of previous work in developing invariants and we will reference a few here. In several instances, the invariants are effective only in special situations. For ex-

ample, they may work only for Lambertian or matte surfaces (normalized RGB, [3, 27, 11, 1, 8], reflectance ratios [16], intrinsic images [26], color invariants [7, 6, 24]), or for specific surface geometry [22, 16], or may be only quasi-invariant [28, 2]. The dichromatic model [23, 10] has been used to separate diffuse and specular reflections (the diffuse component is invariant to specularities) [12, 19, 20, 21, 4]. However, most of the dichromatic model based works assume either (a) objects with homogeneous reflectance, or (b) specific models for diffuse and specular reflections, or (c) prior knowledge about the diffuse or specular colors, or (d) require at least six light sources [13]. In this work, we do not separate reflection components but are interested in separating material from lighting and shape. Compact representations of objects using images under a large number of lighting conditions have been proposed for lighting invariant recognition [1] and shadow removal [14, 5]. However, they often do not have a physical meaning (in terms of object material properties) and therefore hard to use for a problem like material segmentation.

In this paper, we present a class of photometric invariants that have the same computational framework for a large set of BRDFs, sensor types as well as the number of image measurements. We begin by deriving a simple image formation model that is valid for the class of “separable” BRDFs. We make no assumption on the exact models of BRDF (Lambertian, Torrance-Sparrow, etc.) in our image formation model. Then, we show how to decompose the observed image measurements into material properties of scene points and the lighting and scene geometry. This is the key idea for creating lighting and shape invariants as well as material invariants. Most of the invariants in this class need changes in illumination or object position between image acquisitions. The stabilities of the invariants increase when the changes in the scene lighting are large. Thus, our method proves effective in challenging illumination conditions that pose problems for most existing vision algorithms.

The invariants and their discrimination strengths are experimentally demonstrated using several representative scenes with complex shapes, textures, shadings, shadows and specularities. We show that the class of photometric invariants can have implications for vision; for instance, it can be used for lighting insensitive material segmentation, change detection and tracking under complex lighting conditions.

*This research was supported in parts by Siemens Corporate Research, Princeton, NJ 08540 and an NSF ITR Award (No. IIS-00-85864).

Model	K_1	G_1	K_2	G_2
Lambertian	ρ_d	$\frac{I}{\pi} \cos \theta_i$	-	-
Oren-Nayar	$\rho_d \left(1 - 0.5 \frac{\sigma^2}{\sigma^2 + 0.33}\right)$	$\frac{I}{\pi} \cos \theta_i$	$\rho_d \left(0.15 \frac{\sigma^2}{\sigma^2 + 0.09}\right)$	$\frac{I}{\pi} \cos \theta_i \cos(\phi_i - \phi_r) \sin \alpha \tan \beta$

Table 1: Examples of commonly used separable BRDF models for diffuse surfaces. Table shows material K , and geometry G , terms for diffuse reflectance. K_1 and G_1 terms represent material and geometry properties of smooth diffuse surfaces. The K_2 and G_2 terms represent material and geometry properties of rough surfaces. The subscripts i and r denote the incident and the reflected directions measured with respect to surface normal. Here, I is the source intensity, ρ_d is the diffuse albedo, σ is the surface roughness, $\alpha = \text{Max}[\theta_i, \theta_r]$ and $\beta = \text{Min}[\theta_i, \theta_r]$. The cosine terms ($\cos x$) should be replaced by $\text{Max}[0, \cos x]$ to prevent them from being negative (attached shadows). Multiple sources can be represented by just summing up (or integrating) the G terms for each source.

Model	K_3	G_3
NIR (Spike/Delta)	$\rho_s \mathcal{F}(\eta)$	$\delta(\theta_i - \theta_r) \delta(\phi_r)$
	K_4	G_4
Torrance-Sparrow	$\rho_s \mathcal{F}(\eta)$	$\frac{\mathcal{GAF}(\theta_i, \phi_i, \theta_r, \phi_r)}{\cos \theta_r} e^{-\frac{\theta_a^2}{2\sigma^2}}$

Surfaces	Smooth	Rough
Diffuse	1	2
Specular	1	1
Hybrid	2	3

Number of KG terms

Table 2: [Left] Examples of commonly used separable BRDF models for specular surfaces. Table shows material K , and geometry G , terms for specular reflection. The spike/delta model holds for smooth specular surfaces. The Torrance-Sparrow model well approximates the specular reflection from rough surfaces. Multiple sources can be represented by just summing up (or integrating) the G terms for each source. Here, \mathcal{GAF} is called the *Geometric Attenuation Factor*, and θ_a is the inclination of a micro-facet from the surface normal. The Fresnel coefficient $\mathcal{F}(\eta)$ is assumed to be constant with respect to incident angle. This assumption is valid for several materials when illuminated and viewed from non-grazing angles [17]. [Right] Typical number of KG terms needed for various real world surfaces.

2 Separable BRDFs

The radiance \mathcal{L} from a scene point can be written as the sum of diffuse and specular radiances: $\mathcal{L} = \mathcal{L}_d + \mathcal{L}_s$. In many cases, the individual radiances \mathcal{L}_d and \mathcal{L}_s can be further written as products of (a) material related terms and (b) lighting/viewing geometry and object shape related terms as:

$$\mathcal{L} = \sum_{j=1}^n K_j(\text{material}) G_j(\text{shape, lighting, viewpoint}) \quad (1)$$

where, the K terms depend only on intrinsic material properties such as diffuse and specular albedos, refractive index and surface roughness. On the other hand, the G terms depend only on the viewing/illumination geometry and object shape but not on the intrinsic material properties. We refer to the K terms simply as “material” terms and the G terms loosely as “geometry” terms. We now show that the above “separable BRDF model” is valid for a variety of real world surfaces. Tables 1 and 2 illustrate the separable forms for several commonly used BRDFs. This formulation is key to computing the class of photometric invariants presented in this paper.

Diffuse Surfaces : For matte surfaces, the Lambertian model is widely used. In terms of the separable model, we can write:

$$\mathcal{L} = \mathcal{L}_d = K_1 G_1, \quad (2)$$

where, K_1 denotes the diffuse albedo and G_1 is the dot product of source direction and intensity, with surface normal.

The Oren-Nayar model [18] is used for rough diffuse surfaces and we can write a two term separable formulation as:

$$\mathcal{L} = \mathcal{L}_d = K_1 G_1 + K_2 G_2, \quad (3)$$

where, the subscripts 1 and 2 denote the first and second terms of the Oren-Nayar diffuse reflection model. See table 1 for exact equations. Note that the normalized RGB or reflectance ratio [16] invariants used for Lambertian surfaces do not hold for rough surfaces.

Specular Surfaces : For smooth surfaces such as mirrors, the specular reflection is described using a delta function which can be rewritten using our notation as :

$$\mathcal{L} = \mathcal{L}_s = K_3 G_3, \quad (4)$$

where, K_3 denotes the specular albedo and Fresnel coefficient and G_3 is a double-delta function. For rough specular surfaces, we write the Torrance-Sparrow model [25, 17] in separable form as:

$$\mathcal{L} = \mathcal{L}_s = K_4 G_4. \quad (5)$$

See table 2 for exact equations¹. Note that most existing photometric invariants do not hold for specular surfaces.

¹A caveat for the G_4 term in the Torrance-Sparrow model is that it contains the surface roughness term σ . Surface roughness at a microscopic level can be considered as a geometry term (standard deviation of the Gaussian distributed orientations of the surface normals of the micro-facets) and hence is somewhat ambiguous.

Hybrid Surfaces : Hybrid surfaces can have both diffuse and specular reflectance properties. For smooth hybrid surfaces, we can write the radiance model as

$$\mathcal{L} = \mathcal{L}_d + \mathcal{L}_s = K_1 G_1 + K_3 G_3, \quad (6)$$

where, the subscript 1 denotes the diffuse product term and subscript 3 denotes the specular product term. Once again, refer to tables 1 and 2 for exact expressions. On the other hand, for rough surfaces that exhibit both diffuse and specular reflectances, 4 KG product terms are needed :

$$\mathcal{L} = \mathcal{L}_d + \mathcal{L}_s = \sum_{j=1}^4 K_j G_j. \quad (7)$$

At any instant only one of G_3 or G_4 is significant [17]. So, 3 KG products can approximate radiance. Table 2 summarizes the number of product terms needed for different surfaces.

We have shown through several examples that scene radiance can be expressed as an inner product of material terms K and geometry terms G . Note that the model (1) by itself does not explicitly take into account the exact expressions for the K or the G terms. The model is valid for any BRDF that is separable into material and geometry terms (*not just for the examples shown in tables 1 and 2*). In other words, the separable form masks the complexity of the material or geometric properties of the scene and the illumination. As a result, the invariants presented in this paper do not depend on the explicit expressions for BRDF.

3 The Class of Photometric Invariants

The invariants we present can be broadly classified into two categories: (a) *G-Invariants* that depend solely on the K terms and are invariant to illumination, viewing and scene geometry, and (b) *K-Invariants* that depend solely on the G terms and are invariant to material properties of the scene. Any invariant in each category depends on the number of KG products used as well as the number and type (say, RGB, gray scale) of image measurements.

For example, we consider a radiance model with only 2 KG product terms. Since image irradiance E is proportional to scene radiance \mathcal{L} , we write

$$E = g\mathcal{L} = g(K_d G_d + K_s G_s), \quad (8)$$

where g accounts for camera gain. In the remainder of the paper, we will combine g with source intensity (as a scale factor of the G terms) and drop g when we write the image formation model. The subscripts d and s do not have any particular significance. In one instance, the subscript d may denote the diffuse term and the subscript s may denote the specular term, or both the subscripts may denote only diffuse terms of rough diffuse objects².

²Note the model in the form of eq. 8, when the subscripts denote diffuse and specular terms, has been considered by several researchers [20, 21, 10]. However, this is just a special case of the general model (1).

Consider the following set of image measurements made using different values of the material terms K_d and K_s and the geometry terms G_d and G_s :

$$\begin{pmatrix} E_1^{(1)} & E_1^{(2)} & E_1^{(3)} \\ E_2^{(1)} & E_2^{(2)} & E_2^{(3)} \\ E_3^{(1)} & E_3^{(2)} & E_3^{(3)} \end{pmatrix} = \begin{pmatrix} K_d^{(1)} & K_s^{(1)} \\ K_d^{(2)} & K_s^{(2)} \\ K_d^{(3)} & K_s^{(3)} \end{pmatrix} \times \begin{pmatrix} G_d^{(1)} & G_d^{(2)} & G_d^{(3)} \\ G_s^{(1)} & G_s^{(2)} & G_s^{(3)} \end{pmatrix}. \quad (9)$$

The superscripts are written within brackets to avoid confusion with powers. For brevity, we write :

$$\mathbf{E} = \mathbf{K} \mathbf{G}. \quad (10)$$

The variation in the K terms could be due to measurements through different color filters $C = \{r, g, b\}$ in a camera. In this case, the superscripts (1), (2), (3) denote the red, green and blue material properties of a single scene point. In another case, the superscripts of the K terms could correspond to multiple scene points $P = \{p_1, p_2, p_3\}$ with different material properties. Similarly, the variation in the G terms could be due to changes in lighting $L = \{l, m, n\}$ and/or object positions $O = \{x, y, z\}$, between multiple image acquisitions, or could correspond to the G terms of different scene points.

3.1 Geometry (G) Invariants

To compute a function of the measurements $E_j^{(i)}$ that is invariant to all G terms, we consider the following sub-matrices from the first two columns of the matrix \mathbf{E} :

$$\begin{pmatrix} E_1^{(1)} & E_1^{(2)} \\ E_2^{(1)} & E_2^{(2)} \end{pmatrix} = \begin{pmatrix} K_d^{(1)} & K_s^{(1)} \\ K_d^{(2)} & K_s^{(2)} \end{pmatrix} \begin{pmatrix} G_d^{(1)} & G_d^{(2)} \\ G_s^{(1)} & G_s^{(2)} \end{pmatrix}$$

$$\begin{pmatrix} E_1^{(1)} & E_1^{(2)} \\ E_3^{(1)} & E_3^{(2)} \end{pmatrix} = \begin{pmatrix} K_d^{(1)} & K_s^{(1)} \\ K_d^{(3)} & K_s^{(3)} \end{pmatrix} \begin{pmatrix} G_d^{(1)} & G_d^{(2)} \\ G_s^{(1)} & G_s^{(2)} \end{pmatrix}$$

Taking the ratio of determinants of the above sub-matrices, we eliminate all the G terms to get a G-invariant:

$$\frac{E_1^{(1)} E_2^{(2)} - E_2^{(1)} E_1^{(2)}}{E_1^{(1)} E_3^{(2)} - E_3^{(1)} E_1^{(2)}} = \frac{K_d^{(1)} K_s^{(2)} - K_d^{(2)} K_s^{(1)}}{K_d^{(1)} K_s^{(3)} - K_d^{(3)} K_s^{(1)}}. \quad (11)$$

The above expression depends only on the material terms K and not the geometry (lighting, shape, viewpoint) terms G . Note that 3 different G-invariants can be obtained by cyclically changing the superscripts of the K terms in eq. 11.

3.2 Material (K) Invariants

Using a method similar to the derivation of the G-invariant, it is simple to show that one of 3 different selections of sub-matrix pairs in eq. 9 also yields a K-invariant. Computing the determinants of 2×2 sub-matrices selected from the first two

rows of the matrix \mathbf{E} and taking their ratios, we get:

$$\frac{E_1^{(1)} E_2^{(2)} - E_2^{(1)} E_1^{(2)}}{E_1^{(1)} E_2^{(3)} - E_2^{(1)} E_1^{(3)}} = \frac{G_d^{(1)} G_s^{(2)} - G_d^{(2)} G_s^{(1)}}{G_d^{(1)} G_s^{(3)} - G_d^{(3)} G_s^{(1)}}. \quad (12)$$

The above expression depends only on the geometry terms G and not the material terms K .

3.3 Discussion on the Invariants

Depending on the application at hand and the sensor used, only some of the matrix elements \mathbf{E} may be measurable. So, not all K -invariants or G -invariants can be computed from the measurements. However, we shall show that in several instances, these invariants can be measured and applied effectively to common vision tasks.

Handling Color : Observe a colored scene point under 2 different illumination conditions (as in figure 3). Both the camera and the scene remain stationary. The elements of the \mathbf{K} matrix in eq. 10 correspond to intrinsic colors $C = \{r, g, b\}$ of the observed scene point. Also, the elements in the \mathbf{G} matrix correspond to the changes in lighting $L = \{l, m\}$. Then the required G -invariant computed at every pixel independently is written from eq. (11):

$$\frac{E_r^{(l)} E_g^{(m)} - E_g^{(l)} E_r^{(m)}}{E_r^{(l)} E_b^{(m)} - E_b^{(l)} E_r^{(m)}} = \frac{K_d^{(r)} K_s^{(g)} - K_d^{(g)} K_s^{(r)}}{K_d^{(r)} K_s^{(b)} - K_d^{(b)} K_s^{(r)}}. \quad (13)$$

Note here that we require only one change in scene illumination (i.e., the third column in eq. (9) is not necessary). However, we need three color images under three different lighting conditions $L = \{l, m, n\}$ to compute a material invariant or K -invariant according to eq. 12:

$$\frac{E_r^{(l)} E_g^{(m)} - E_g^{(l)} E_r^{(m)}}{E_r^{(l)} E_g^{(n)} - E_g^{(l)} E_r^{(n)}} = \frac{G_d^{(l)} G_s^{(m)} - G_d^{(m)} G_s^{(l)}}{G_d^{(l)} G_s^{(n)} - G_d^{(n)} G_s^{(l)}}. \quad (14)$$

If lighting is fixed in the environment, but the objects are moving, then the lighting superscripts $L = \{l, m, n\}$ need to be just replaced by the object superscripts $O = \{x, y, z\}$.

Handling Gray-Scale : Using color makes it possible to compute invariants locally at every scene point. To compute invariants in gray-scale images we need multiple scene points. In other words, the superscripts 1, 2, 3 will correspond to different scene points $P = \{p_1, p_2, p_3\}$. In this case, invariants computed will be sparse. However, they will be valid for, say, scene points with same geometry terms (similar surface normals and source directions), and thus can be used in interesting grouping algorithms.

Handling more than 2 KG terms : All the invariants we show in this paper will be based on two KG terms (see 8). However, they can be easily generalized to include more KG terms. The key idea in computing the K -/ G -invariants is that the image measurements are arranged in the form of square

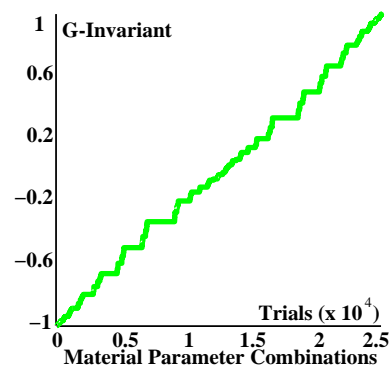


Figure 1: G -Invariant simulations (using Oren-Nayar model) illustrating the discriminability to about 25000 sets of values for material parameters (diffuse color albedos and roughness) of a scene point with fixed surface normal and viewing direction. Two lighting directions are used to compute the G -invariant. The absence of long flat regions in the plot shows good discriminability.

matrices and ratios of their determinants eliminate either the K terms or the G terms. As an example, consider 3 KG product terms in the radiance model. The measurement matrix is larger when compared to eq. 9 and is written as :

$$\mathbf{E}(4 \times 4) = \mathbf{K}(4 \times 3) \cdot \mathbf{G}(3 \times 4). \quad (15)$$

Then, the determinant ratio of two 3×3 sub-matrices of \mathbf{E} will yield the necessary K -/ G -Invariants.

4 Invariant Discriminability and Sensitivity

In theory, the value of the invariant expression may not be unique to a particular material. So, two or more materials could have the same determinant ratio. However, this ambiguity decreases by using invariants computed using cyclic combinations of superscripts in eq. 11. To empirically show that they indeed have good discriminative power, we simulated a fixed scene point and illumination/viewing geometry with about 25000 different combinations of material parameters (color diffuse albedos and surface roughness) using the Oren-Nayar model. The plot of the G -invariants vs. trial number is shown in figure 1. The absence of long flat regions in the curve shows that the invariant is discriminative. The numerical stability of the invariants depend on the amount of change in lighting/viewing geometry between image acquisitions. To study how the stability of G -invariant to lighting changes, we simulated several degrees of lighting changes (by changing the directions of 2 sources) on a scene point whose material and geometric properties are fixed. The plot of the computed invariant in the presence of noise for a range of lighting changes, shows that the invariants stability increases with increased change in lighting (see figure 2). We will leave the detailed statistical analysis for both the discriminability and the sensitivity issues for future work.




Figure 2: Lighting change sensitivity plot: Noise (STD = 1 in a scale of 0-256 gray levels) was added to the simulated radiances with gradual change in light source positions. The G-invariant then computed is plotted against the angular separation in degrees of the light sources between image acquisitions. For very small lighting changes, the invariants are unstable in the presence of noise and the stability increases with increased change in lighting.

5 Experimental Results : Finding Materials

We now describe various experiments that demonstrate the invariance to object shape, illumination intensity and direction and viewing geometry using eq. 13. The G-invariants are computed from multiple images of the scene under different lighting conditions. We have chosen the scenes so as to bring out the power of invariants in the presence of shadows, specularities and other complex shadings.

All our experiments are done using a 12-bits per pixel Kodak DCS 760 digital camera. Multiple images with different exposures are acquired, radiometrically calibrated and combined to obtain high dynamic range radiance maps of the scenes. Since our goal is to showcase the physics behind the invariants, we have chosen to use high quality data in our experiments. For low quality data, probabilistic analyses proposed by earlier works (eg., [15]) can be used to decide where in the image the computed invariants are unreliable.

5.1 Scene 1: Complex Shape and Self-Shadowing

The object of interest in the first experiment is a ceramic bust of David shown in figure 3. The bust has complex shape and creates specular effects as well as self shadows. The object was imaged under different lighting conditions using fluorescent (a) and halogen (b,c,d,e) light sources. The fluorescent source was ambient room illumination whereas the halogen sources were nearby area sources. The bust is made up of the same material irrespective of the complex shading. The G-invariant we compute should therefore produce a constant (or flat) image that does not have any shading or specularity.

The result of applying eq. 13 to image pairs is shown by G-invariant images (figs. 3(f)-(i)). Each invariant image is computed using a different pair of images [(a),(b)], [(a),(c)], etc). The accuracy of the invariant shown using the histograms of the invariant images, are narrow and at the same bin positions. The standard deviation of the G-invariant images were between 2.3 and 2.6 gray levels (out of 256).

Figure 4 shows the same bust against a more complex red velvet background. The cloth has creases and folds producing shadows and the bust itself casts shadows on the cloth. The G-invariant in this case should produce not only a flat image for the bust but also a flat image (of a different brightness) for the crumpled cloth. The G-invariant in fig. 4(c) is computed from 2 images (figs. 4(a) and (b)). The intensity histograms of the in (a, b) do not have distinctive peaks for different materials. However, the histogram in (c) has two separate peaks corresponding to the bust and the cloth. It is clear that a simple thresholding of the invariant image yields a near perfect segmentation in spite of the complex shadowing.

5.2 Scene 2: Textured Objects

The previous subsection described experiments with an object made of homogeneous material. We now describe an experiment performed with a textured doll made of cotton shown in figure 6. The G-invariant computed should be a flat 2D texture of the doll without any shadowing or other complex shading due to the 3D shape of the doll.

As before, several images under different illumination conditions are captured (figures 6(a,b,d,f)). Three different pairs of images are used to compute the three G-invariant images shown in figs. 6(c,e,g). The G-invariants were computed using the same expression (eq. 13) used before. Notice that the shading on the object as well as shadows present in the original images are completely eliminated yielding a flat textured appearance of the doll. This facilitates simple 2D template matching type of algorithms for recognition.

5.3 Scene 3: Complex Specularities

In this example, we will demonstrate the G-invariant in the presence of strong specularities due to nearby sources. Note that this is a hard case for most existing techniques. The object is a green plastic pear shown in figure 7. Three images of the pear are acquired under different lighting conditions and are shown in figures 7(a,b,d). Notice the strong specularities on the images present at different positions. The G-invariant images computed using two pairs of images are shown in figures 7(c) and (e). Figure 7(f) shows a K-invariant computed using all the three original images. The expression used to compute the K-invariant is the same as eq. 12 with the superscripts of the K and E terms denoting the different color channels r, g, and b. Note that all the specularities are visible in the K-invariant image. This K-invariant is a function of lighting and shape but not material properties (diffuse and specular albedos) of the pear.

6 Implications for Vision

In this section, we discuss the various applications for the invariant class we developed in this paper. The invariants, especially the G-invariants, are very useful in transforming images taken under complex lighting conditions into simpler material maps. Existing vision algorithms for segmentation,

Input Images (Two images are used to compute a G-Invariant)



(a) Image 1 (b) Image 2 (c) Image 3 (d) Image 4 (e) Image 5

G-Invariant Images (each image computed from a pair of input images)



(f) From 1 and 2 (g) From 1 and 3 (h) From 1 and 4 (i) From 1 and 5

Histograms of G-Invariant Images





Figure 3: Separating Material from Illumination, Shape and Viewing direction using the G-invariant. [(a) - (e)] Scene 1: A ceramic bust of David is imaged under different lighting conditions. [(f) - (i)] Computed G-invariant Images. Each image is computed (eq. 13) using a different pair of images. Histograms of the invariant images are narrow (the invariant images are indeed flat) showing that robustness of the G-invariant to complex shape and shadows.

recognition and tracking can then be more accurately applied to the invariant images.


We presented two examples of material segmentation in figures 4 and 5. Note that lighting insensitive recognition of objects can be made more robust if applied to G-invariant images rather than the raw input images. For instance, the recognition of the textured doll in figure 6 could be done by applying 2D template matching to the G-invariant image that is devoid of shadows or shading. Recall that most existing invariants for lighting insensitive change detection and tracking rely on the lambertian reflectance model. With the invariants developed in this paper, we believe that more efficient change detection and tracking can be performed since our invariants can handle both complex material and lighting.

7 Summary


We presented a class of photometric invariants that can be used for a variety of vision tasks such as recognition, material classification and tracking. The invariants are valid for the class of separable BRDFs which can be written as dot products of material and geometry terms. All the invariants in this class are determinant ratios of measurement sub-matrices. They are effective for complex shapes in complex lighting environments. The geometry invariants are invariant to specularities and to even cast, attached and self shadows. The code required for our algorithm is about 5 lines in Matlab (to compute two matrix determinants and their ratios). Our goal in this paper is to show the capability of the invariant functions on complex materials and lighting conditions. For this,



(a) Input image 1 and its intensity histogram



(b) Input image 2 and its intensity histogram



(c) G-Invariant image using (a) and (b) and its histogram.




(d) Material Segmentation

Figure 4: Finding materials in a scene with complex shapes and shadows. The scene consists of the bust of David with a crumpled red velvet cloth in the background. The crumpled cloth produces self shadows and the bust casts shadow onto the cloth. [(a),(b)] Two input images under different lighting conditions and their intensity histograms. [(c)] The G-invariant image computed (eq. 13) from the input images. The image is brightened for display. The histogram of the invariant image in (c) clearly has two peaks - one corresponding to the red velvet cloth and the other to the bust. [(d)] A simple thresholding of the image in (c) produces correct material segmentation (Black: Cloth, White: Statue of David).


all our algorithms need are changes in illumination and accurate (high dynamic range) irradiance values. In future work, we wish to do a detailed statistical analysis of the discriminability and sensitivity of the invariant class for low quality images (say, a 8-bit noisy surveillance camera).

References


- [1] P. Belhumeur and D. Kriegman. What is the set of images of an object under all possible lighting conditions? *In Proc. CVPR*, 96.
- [2] J. Burns, R. Weiss, and E. Riseman. The non-existence of general-case view-invariants, geometric invariance in computer vision, edited by j. mundy, a. zisserman. *MIT Press Cambridge*, 1992.
- [3] H. F. Chen, P. N. Belhumeur, and D. W. Jacobs. In search of illumination invariants. *In Proc. IEEE CVPR*, 2000.




(a) Input image 1 (contrast stretched) and its intensity histogram



(b) Input image 2 (contrast stretched) and its intensity histogram



(c) G-Invariant image using (a) and (b) and its histogram.



(d) Material Segmentation

Figure 5: Second example of finding materials in a scene with similar colors. The scene consists of two whitish cubes made of sand-blasted aluminum and marble. Notice the good segmentation despite cubes being of similar colors.

- [4] H. Farid and E. H. Adelson. Separating reflections and lighting using independent components analysis. *In Proc. of CVPR*, 1999.
- [5] S. D. Hordley G. D. Finlayson and M. S. Drew. Removing shadows from images. *In Proc. ECCV*, 2002.
- [6] J.M. Geusebroek and A.W.M. Smeulders. Measurement of color invariants. *In CVPR00*, pages I: 50–57, 2000.
- [7] T. Gevers and A.W.M. Smeulders. Color constant ratio gradients for image segmentation and similarity of textured objects. *In CVPR01*, pages I:18–25, 2001.
- [8] P. Hallinan. A low-dimensional representation of human faces for arbitrary lighting conditions. *In Proc. CVPR*, 1994.
- [9] R.M. Haralick. Propogating covariance in computer vision. *PRAI*, 10:561–572, 1996.
- [10] G.J. Klinker, S.A. Shafer, and T. Kanade. A physical approach to color image understanding. *IJCV*, 4(1):7–38, January 1990.
- [11] J. J. Koenderink and A. J. van Doorn. Photometric invariants related to solid shape. *Optica Acta*, 27(7):981–996, 1980.
- [12] Stephen Lin and Sang W. Lee. Estimation of diffuse and specular appearance. *In Proc. IEEE CVPR*, 1999.
- [13] Stephen Lin and Sang W. Lee. An appearance representation for multiple reflection components. *In Proc. CVPR*, 2000.
- [14] Y. Matsushita, K. Ikeuchi, M. Sakauchi, and K. Nishino. Shadow elimination for robust video surveillance. *IEEE Workshop on Motion and Video Computing*, 2002.
- [15] S.J. Maybank. Probabilistic analysis of the application of the cross ratio to model-based vision. *IJCV*, 16(1):5–33, September 1995.
- [16] S.K. Nayar and R.M. Bolle. Reflectance based object recognition. *IJCV*, 17(3):219–240, March 1996.



(a) Input image 1



(b) Input image 2 (c) G-Invariant : from 1 and 2



(d) Input image 3 (e) G-Invariant : from 1 and 3



(f) Input image 4 (g) G-Invariant : from 1 and 4 (h) Variance

Figure 6: Separating Material from Geometry using the G-invariant in the presence of texture and soft shadows. [(a),(b),(d),(f)] Scene 2: A textured doll made of cotton is imaged under different lighting conditions. Notice the subtle changes in positions of soft shadows on the doll. [(c),(e),(g)] Computed G-invariant images. Notice the flat texture appearance of the invariant images without any shading or shadows due to 3D shape. [(h)] Variance in the invariant for the head region computed using a standard covariance propagation method [9]. The hair and eye region on the doll are dark and the variance is high. So the G-invariants computed in this region are not reliable.


- [17] S.K. Nayar, K. Ikeuchi, and T. Kanade. Surface reflection: Physical and geometrical perspectives. *PAMI*, 13(7):611–634, July 1991.
- [18] M. Oren and S.K. Nayar. Generalization of the lambertian model and implications for machine vision. *IJCV*, 14(3):227–251, April 1995.
- [19] Y. Sato and K. Ikeuchi. Temporal-color space analysis of reflection. *JOSA A*, 11, 1994.
- [20] Y. Sato and K. Ikeuchi. Reflectance analysis under solar illumination. In *PBMCV95*, page SESSION 6, 1995.



(a) Input image 1



(b) Input image 2




(c) G-Invariant from images 1 and 2.



(d) Input image 3



(e) G-Invariant from images 1 and 3.



(f) Material or K-Invariant using images 1, 2, 3.

Figure 7: Separating Material from Geometry using the G-invariant in the presence of strong specularities. [(a), (b), (d)] Scene 3: A green plastic pear is imaged under different lighting conditions. Notice the strong changes in positions of specularities on the doll. [(c), (e)] G-invariant images, each computed (eq. 13) using a different pair of input images. Notice the flat or constant appearance of the invariant images without any specularities. [(f)] K-invariant computed using eq. 14 from the 3 images in (a), (b), and (d). Notice the positions of specularities seen in the K-invariant image.

- [21] Y. Sato, M.D. Wheeler, and K. Ikeuchi. Object shape and reflectance modeling from observation. In *MjR01*, pages Chapter II–4, 2001.
- [22] C. Schmid and R. Mohr. Local grayvalue invariants for image retrieval. *IEEE Transactions on PAMI*, 19(5):530–535, 1997.
- [23] S. Shafer. Using color to separate reflection components. *Color Research and Applications*, pages 210–218, 1985.
- [24] D.A. Slater and G. Healey. The illumination-invariant recognition of 3d objects using local color invariants. *PAMI*, 18(2), Feb 1996.
- [25] K.E. Torrance and E.M. Sparrow. Theory for off-specular reflection from roughened surfaces. *JOSA*, 57:1105–1114, 1967.
- [26] Y. Weiss. Deriving intrinsic images from image sequences. In *Proc. ICCV*, 2001.
- [27] L. Wolff and J. Fan. Segmentation of surface curvature with a photometric invariant. *JOSA A*, 11(11):3090–3100, 1994.
- [28] M. Zerroug and R. Nevatia. Three-dimensional descriptions based on the analysis of the invariant and quasi-invariant properties of some curved-axis generalized cylinders. *IEEE Trans. PAMI*, 18(3), 1996.