Diffuse Reflectance from Rough Surfaces

Michael Oren and Shree K. Nayar

Department of Computer Science, Columbia University, New York, N.Y. 10027 *

Abstract

In this paper, we present a comprehensive model that predicts reflectance from rough diffuse surfaces. We show that diffuse reflectance from rough surfaces increases as the viewing direction approaches the source direction. This is in contrast to Lambertian surfaces, where radiance is independent on the viewing direction. The new model is a generalization of the Lambertian model, and has significant implications for machine vision, graphics, and remote sensing.

1 Introduction

A surface that obeys Lambert's Law appears equally bright from all viewing directions. This model for diffuse reflection is one of the most widely used models in computer vision. It is used used explicitly by vision techniques such as shape from shading and photometric stereo, and implicitly by techniques such as binocular stereo and motion detection.

For several real-world objects, however, the Lambertian model can prove to be a poor and inadequate approximation to the diffuse component (see Figure 1(a) and (b)). This deviation from Lambertian reflectance is significant for very rough surfaces, and increases with the angle of incidence. The topic of rough diffuse surfaces has been extensively studied in the areas of applied physics and geophysics [Minnaert-1941], [Smith-1967] and [Buhl *et al.*-1968]. However, a complete reflectance model has not been presented.

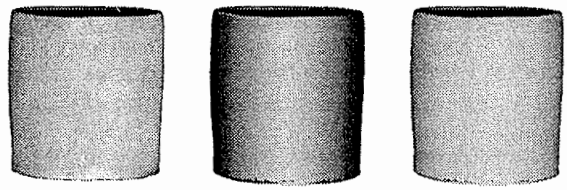

(a) Image (b) Lambertian (c) Model

Figure 1: Real image of a cylindrical clay vase compared with images rendered using the Lambertian reflectance model and the proposed model. Illumination is from the sensor direction.

We have developed a general and complete model for diffuse reflectance. This model can be applied to isotropic as well as anisotropic rough surfaces, and can handle arbitrary source and viewer directions. Further, it takes into account complex geometrical effects such as *masking*, *shadowing*, and *interreflections* between points on the rough surface. Experiments show the proposed model to be in strong agreement with reflectance measured from a variety of common surfaces, such as, clay (see Figure 1(c)), plaster, and sand.

*This research was supported in part by DARPA Contract No. DACA 76-92-C-0007 and in part by the NSF Research Initiation Award.

2 Reflectance Model

For lack of space, we will focus on the main results. Details of the model derivation can be found in [Oren and Nayar-1992]. The surface is assumed to be composed of long symmetric V-cavities. Each cavity consists of two facets. Each facet area da is small compared to the area dA of the surface patch that is imaged by a single pixel. Thus, each pixel includes a very large number of facets. Further, the facet dimensions are large compared to the wavelength of incident light and hence geometrical optics can be used to derive the reflectance model. While developing the reflectance model, each facet is assumed to be Lambertian in reflectance.

The roughness of the surface is given by the probability distribution of the facet slopes. Our approach is to compute the radiance contribution of individual facets on the surface. The total radiance of a surface patch is then determined as an aggregate of the contributions of all facets. Reflectance models have been developed for three different probability facet-slope distributions. In the most general case, the distribution is assumed to be Gaussian with mean $\mu = 0$ and standard deviation σ. The standard deviation represents the macroscopic roughness of the surface. For this case, radiance is obtained as an integral over all facets. This integral is evaluated by identifying a set of basis functions and using functional approximations. The final results are given below. Let $\hat{s} = (\theta_i, \phi_i)$ and $\hat{v} = (\theta_r, \phi_r)$ be the source and the viewer directions. We define $\alpha = Max[\theta_r, \theta_i]$ and $\beta = Min[\theta_r, \theta_i]$. The radiance of a surface with roughness σ, due to direct illumination by the source is:

$$L_r^1(\theta_i, \phi_i; \theta_r, \phi_r; \sigma) = \frac{\rho}{\pi} E_0 \cos\theta_i \Bigg[C_1(\sigma) + \tag{1}$$

$$\cos(\phi_r - \phi_i) C_2(\alpha; \beta; \phi_r - \phi_i; \sigma) \tan\beta +$$

$$\Big(1 - |\cos(\phi_r - \phi_i)|\Big) C_3(\alpha; \beta; \sigma) \tan\left(\frac{\alpha + \beta}{2}\right) \Bigg]$$

$$C_1 = 1 + \frac{\sigma^2}{\sigma^2 + 0.4}\left(\frac{2}{\pi} - 1\right)$$

$$C_2 = \begin{cases} 0.4\frac{\sigma^2}{\sigma^2+0.2}(\sin\alpha)^{\left(0.56+\frac{0.20}{\sigma^2}\right)} & \\ \qquad\qquad \text{if } \cos(\phi_r - \phi_i) \geq 0 & \\ 0.4\frac{\sigma^2}{\sigma^2+0.2}\left((\sin\alpha)^{\left(0.56+\frac{0.20}{\sigma^2}\right)} - \left(\frac{2\beta}{\pi}\right)^{\left(2+\frac{1}{2\sigma^2}\right)}\right) & \\ \qquad\qquad \text{otherwise} & \end{cases}$$

$$C_3 = 0.11\,\frac{\sigma^2}{\sigma^2 + 0.2}\left(\frac{4\alpha\beta}{\pi^2}\right)^{\left(1.55+\frac{1}{1.55\sigma^2}\right)}$$

The radiance due to interreflections (secondary illumination) between facets on the surface was found to be:

$$L_r^2(\theta_i, \phi_i; \theta_r, \phi_r; \sigma) = \rho^2 E_0 \cos\theta_i\,(0.045\,\frac{\sigma^2}{\sigma^2 + 0.2})\left[1 - \cos(\phi_r - \phi_i)\left(\frac{2\beta}{\pi}\right)^{\left(1.3+\frac{0.31}{\sigma^2}\right)}\right] \quad (2)$$

The two components are combined using the interreflection parameter χ to obtain the total surface radiance:

$$L_r = L_r^1 + \chi\, L_r^2 \quad (3)$$

Here, $0.5 \leq \chi \leq 1$ is used to account for discrepancies in the assumed surface roughness model and the real surface. It is important to note that the above radiance expression obeys Helmholtz's reciprocity principle *Also note that the above model reduces to the Lambertian model when* $\sigma = 0$.

Qualitative Model

The following simplified model is proposed for applications where high accuracy is not critical:

$$L_r \approx \cos\theta_i(\mathbf{A} + \mathbf{B}\,Max\left[0, \cos(\phi_r - \phi_i)\right]\sin\alpha\tan\beta) \quad (4)$$

This model has only two coefficients, namely, **A** and **B**. The relationships between these coefficients and the parameters σ (roughness) and ρ (albedo) were arrived at empirically via numerous simulations:

$$A \approx \rho(\frac{1}{\pi} - 0.09\frac{\sigma^2}{\sigma^2 + 0.4}) \quad (5)$$

$$B \approx \rho(0.125\frac{\sigma^2}{\sigma^2 + 0.18}) \quad (6)$$

3 Experiments

We have conducted several experiments to verify the proposed reflectance model (equation 3). These experiments were conducted on rough surfaces with matte local reflectance properties. The results for three samples are shown in Figure 2. For each sample, radiance was measured by varying the sensor direction. For all three samples, the model predictions (solid lines) and experimental measurements (dots) match remarkably well. Notice that these radiance plots deviate substantially from Lambertian reflectance. *Surface radiance increases as the viewing direction approaches the source direction.* This is in contrast to Lambertian surfaces where radiance does not vary with sensor direction, and specular surfaces where the radiance increases in the forward (specular) direction.

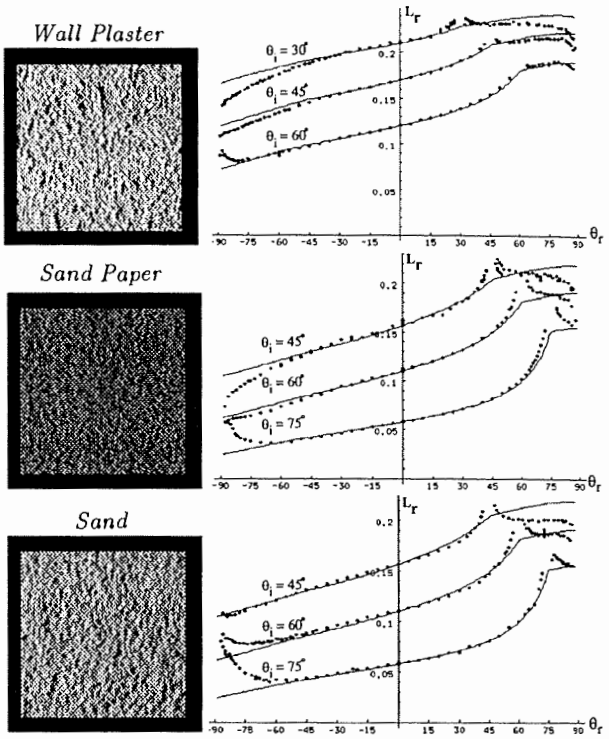

Figure 2: Reflectance plotted as function of viewing (sensor) direction. The measurements (dots) are compared with model predictions (solid lines).

References

[Buhl *et al.*, 1968] D. Buhl, W. J. Welch, and D. G. Rea. Reradiation and thermal emission from illuminated craters on the lunar surface. *Journal of Geophysical Research*, 73(16):5281–5295, August 1968.

[Minnaert, 1941] M. Minnaert. The reciprocity principle in lunar photometry. *Astrophysical Journal*, 93:403–410, 1941.

[Oren and Nayar, 1992] M. Oren and S. K. Nayar. Generalization of the Lambertian Model and Implications for Machine Vision. Technical Report CUCS-057-92, Columbia University, 1992.

[Smith, 1967] B. G. Smith. Lunar surface roughness: Shadowing and thermal emission. *Journal of Geophysical Research*, 72(16):4059–4067, August 1967.

Removal of Specularities Using Color and Polarization

Shree K. Nayar, Xi-Sheng Fang, and Terrance Boult *

Center for Research in Intelligent Systems, Department of Computer Science,
Columbia University, New York, N.Y. 10027, U.S.A.

Abstract

Specular reflections and interreflections produce strong highlights in brightness images. These highlights can cause vision algorithms, such as, segmentation, shape from shading, binocular stereo, and motion detection to produce erroneous results. We present an algorithm for separating the specular and diffuse components of reflection from images. The method uses color and polarization, simultaneously, to obtain strong constraints on the reflection components at each image point. Polarization is used to locally determine the color of the specular component, constraining the diffuse color at a pixel to a one dimensional linear subspace. This subspace is used to find neighboring pixels whose color is consistent with the pixel. Diffuse color information from consistent neighbors is used to determine the diffuse color of the pixel. In contrast to previous separation algorithms, the proposed method can handle highlights that have a varying diffuse component as well as highlights that include regions with different reflectance and material properties. We present several experimental results obtained by applying the algorithm to complex scenes with textured objects and strong interreflections.

1 Introduction

Reflection of light from surfaces can be classified into two broad categories: diffuse and specular. The *diffuse component* results from light rays penetrating the surface, undergoing multiple reflections and refractions, and re-emerging at the surface. This component is distributed in a wide range of directions around the surface normal, giving the surface a matte appearance. If the viewing direction of an image sensor is varied, diffuse reflections from scene points change slowly and in the ideal case of Lambertian surfaces, it does not change at all. The *specular component*, on the other hand, is a surface phenomenon. Light rays incident on the surface are reflected such that the angle of reflection equals the angle of incidence. Even for marginally rough surfaces, the specular reflections are concentrated in a compact lobe around the specular direction. This concentration of light energy causes strong *highlights* in brightness images of scenes. These highlights can cause vision algorithms for scene segmentation and shading analysis to produce erroneous results. If the sensor direction is varied, highlights shift, diminish rapidly, or suddenly appear in other parts of the scene. This strong directional dependence of specular reflection, poses serious problems for vision techniques such as binocular stereo and motion detection. Hence, specularities are often undesirable in images.

In this paper, we present an algorithm that separates the diffuse and specular components of brightness from images. Separation of reflection components has been a topic of active research in the past few years. Here, we discuss only those efforts that have resulted in algorithms that have been tested on real images. Most of this work is based on the dichromatic reflectance model proposed in [Shafer, 85]. The dichromatic model suggests that, in the case of dielectrics (non-conductors), the diffuse component and the specular component generally have different spectral distributions. Hence, the color of an image point can be viewed as the sum of two vectors with different directions in color space. Using this model, Klinker [Klinker, 88] and Gershon [Gershon, 87] independently observed that the color histogram of an object with uniform diffuse color takes the shape of a skewed T with two limbs. One limb corresponds to purely diffuse points on the object, which have the same color but differ in magnitude, and the second limb represents a highlight region. They proposed algorithms for automatically identifying the two limbs and used the directions of the limbs to separate the diffuse and specular components at each object point. Later, Bajcsy *et al.* [Bajcsy *et al.*, 90] showed that the color histogram of an object could have additional limbs that correspond to highlights caused by interreflections between objects. More recently, Lee [Lee, 91] proposed moving a sensor and applying spectral differencing to color histograms of consecutive images to identify specular points in the image. This method however does not compute accurate estimates of the specular component at each image point.

All of the above algorithms rely solely on color information to separate specular and diffuse reflections from images. Since the separation is not possible when an image point is treated in isolation, these methods analyze the anatomy of color histograms. Two major limitations result from the above approach. First, real scenes are complex and include objects with texture and varying reflectance. Color histograms of such scenes are generally unpredictable and a set of linear clusters such as the skewed T are unlikely. Second, all points on the highlight region are assumed to have the same diffuse component (color and magnitude). Even for an object with uniform reflectance, this assumption is valid only if the object surface is very smooth. In the case of rough surfaces, the highlights spread over a wider range of surface

*This work is supported in part by DARPA contract DACA-76-92-C-007 and NSF contract IRI-90-57951.

normals and the specular limb of the skewed T does not have a well-defined direction.

Recently, Wolff and Boult [Wolff and Boult, 91] proposed a polarization based method for separating specular and diffuse components from gray-level (black and white) images. Details of this method will be discussed later. Their polarization-based method assumes that the diffuse component is constant over the entire highlight region. They also assume that the material type and surface normal do not vary within the highlight region. Using these assumptions, an estimate for a constant diffuse term is obtained. We will show later that the assumptions made in [Wolff and Boult, 91] are often not practical in the context of real scenes.

This paper presents a new algorithm for the separation of specular and diffuse reflection components from images. This algorithm uses color and polarization *simultaneously*, to obtain new constraints on the reflection components. As a result, it does not suffer from many of the problems associated with previous methods based either color or polarization. We assume that the scene consists of dielectric objects. This leads to two assumptions: (a) the dichromatic model is applicable, and (b) the specular component is polarized while the diffuse component is not. The restrictions imposed by these assumptions are discussed in subsequent sections. The proposed algorithm can estimate specular components that result not only from direct source illumination but also interreflections between points in the scene. We show that, under reasonable assumptions, polarization information can be used to obtain the color of the specular component independently for each point with a specular component. For each such point, the result is a line in color space on which the diffuse vector must lie. This line imposes strong constraints on the color of the diffuse component of that image point. Neighboring diffuse colors that satisfy these constraints are used to compute the diffuse component of the image point.

Since the specular color of each image point is computed independently, our approach has the following advantages over previous methods: (a) The diffuse component is not assumed to be constant under the highlight region; (b) The Fresnel ratio (which depends on the material properties and the angle of incidence) can vary over highlight regions; and (c) the diffuse component may be textured. The algorithm requires that each image point has a few (at least three) neighbors that have the same diffuse color (that is, direction in color space but not necessarily magnitude). In the experimental section, we present several results obtained by applying the algorithm to complex scenes with multiple highlights and interreflections.

2 Reflection and Interreflection

We begin by describing the mechanisms involved in the processes of reflection and interreflection. Figure 1 shows two points, A and B, in a scene. Reflection from the point A has two components, namely, diffuse and specular.[1] The diffuse component arises from the scattering of light rays that enter the surface and undergo multiple reflections and refractions. The specular component, on the other hand, is a surface phenomenon and results from single reflection of incident light rays. The surface may be assumed to be composed of several planar elements, or facets, where each facets has its own orientation. The result is a specular component that spreads around the specular direction, the width of the distribution depending on the roughness of the surface [Torrance and Sparrow, 67].

Now let us consider the phenomenon of *interreflections*. Points in the scene receive light not only from the light sources but also from other scene points. Assume that point B reflects, into the sensor, light energy from point A. The resulting image brightness value can be viewed as the linear combination of four possible interreflection components: (a) *diffuse-diffuse*; (b) *specular-diffuse*; (c) *diffuse-specular*; and (d) *specular-specular*. In each case, the first term represents the component received from point A and the second represents the component reflected by point B. In general, point B could reflect light due to both direct illumination by light sources as well as interreflections from other scene points. Without loss of generality, each brightness value in the image can be expressed as the sum of two components, diffuse and specular. The specular component can result either from direct illumination by a light source or due to the diffuse-specular or specular-specular interreflections. We assume that the specular reflection received by the sensor from any given point is either due to source illumination or due to interreflections and not both. In other words, any given scene point is positioned and oriented to specularly reflect essentially from either a light source or another scene point but not both. This assumption holds well except for very rough surfaces.

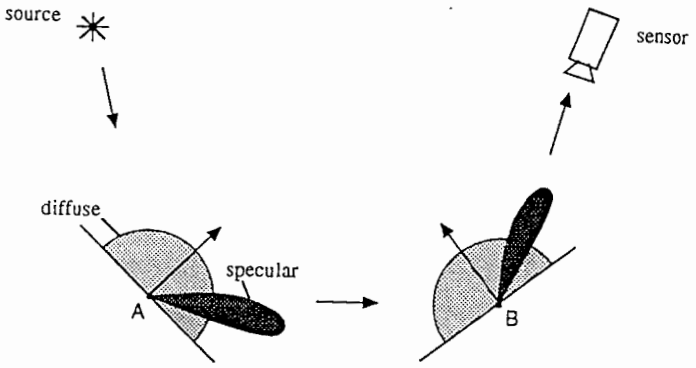

Figure 1: Components of reflection and interreflection.

3 Polarization

The method presented in this paper uses a polarization filter to determine the color of the specular component. In this section, we present a brief overview of polarization and discuss the type of surfaces for which it provides useful information. Detailed discussions on the theory of polarization can be found in [Born and Wolf, 65]. In the field of

[1] Recently, [Nayar *et al.*, 91] proposed a reflectance framework that includes three primary components of reflection: the *diffuse lobe*; the *specular lobe*; and the *specular spike*. In this paper, the two specular components are combined to yield, the *specular component*.

machine vision, polarization methods were first introduced by Koshikawa [Koshikawa, 79] who used ellipsometry for shape interpretation and recognition of glossy objects. More recently, Wolff and Boult [Wolff and Boult, 91] examined the use of linear polarization for highlight removal and material classification.

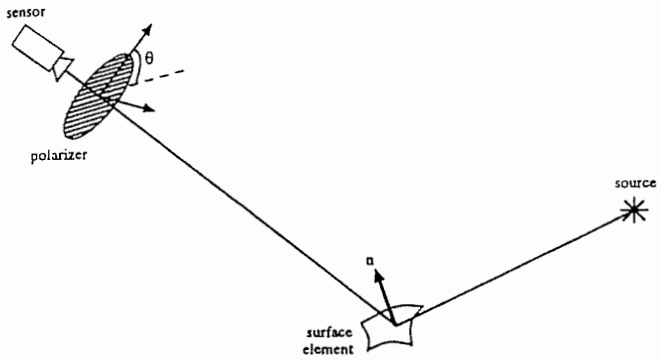

Figure 2: Surface element illuminated by a source and imaged through a polarization filter.

Figure 2 shows a surface element illuminated by a source and imaged by a sensor. A polarization filter is placed in front of the sensor. As in the previous section, let the image brightness value corresponding to the surface element be: $I = I_d + I_s$, where, I_d is the diffuse component and I_s the specular component. The linear polarization of each light wave is determined by the direction of its electric field vector. In general, the light energy (due to several light waves) reflected by a surface may be partially *polarized*. The extent of polarization depends on several factors including the material of the reflecting surface element, its orientation with respect to the image sensor, and the types of reflection mechanisms (specular or diffuse) at work.

The diffuse component of reflection tends to be *unpolarized*.[2] In contrast, the specular component tends to be partially polarized; rotation of the polarization filter varies the specular component as a cosine function, as shown in Figure 3. The specular component can be expressed as the sum of a *specular constant* I_{sc} and a *specular varying* term that is a cosine function with amplitude I_{sv}:

$$I = I_d + I_{sc} + I_{sv} \cos 2(\theta - \alpha) \qquad (1)$$

where, θ represents the angle of the polarization filter and α is the phase angle determined by the projection of the normal of the surface element onto the plane of the polarization filter. The exact values of I_{sc} and I_{sv} depend on the material properties and the angle of incidence. This dependence is determined by the *Fresnel reflection coefficients* $F_\perp(\eta, \psi)$ and $F_\parallel(\eta, \psi)$ which represent the polarization of the reflected light waves in the directions perpendicular and parallel to the plane of incidence, respectively. The relationship between

[2]Note that this assumption does not hold near the occluding contour of an object, see [Boult and Wolff, 91]. That paper addresses the classification of scene edges based on their polarization characteristics.

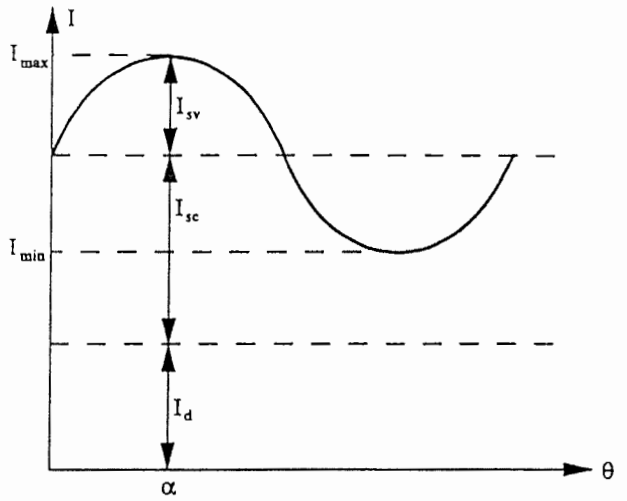

Figure 3: Image brightness plotted as a function of polarization filter position.

I_{sc} and I_{sv} and the Fresnel coefficients is:

$$\frac{I_{sc} + I_{sv}}{I_{sc}} = \frac{F_\perp(\eta, \psi)}{F_\parallel(\eta, \psi)} \qquad (2)$$

The parameter η is the *complex index of refraction* of the surface medium, and depends on the physical properties of the reflecting material. The parameter ψ is the angle of incidence.[3]

Note that the terms I_d and I_{sc} in equation 1 are constant and can be represented by a single component $I_c = I_d + I_{sc}$ to obtain: $I = I_c + I_{sv} \cos 2(\theta - \alpha)$. For any given position θ_i of the polarization filter we have:

$$I_i = I_c + I_{sv} \cos 2(\theta_i - \alpha). \qquad (3)$$

This can also be represented as the dot product of two vectors:

$$\begin{aligned} \mathbf{f}_i &= (1, \cos 2\theta_i, \sin 2\theta_i) \\ \mathbf{v} &= (I_c, I_{sv} \cos 2\alpha, I_{sv} \sin 2\alpha) \\ I_i &= \mathbf{f}_i \cdot \mathbf{v}. \end{aligned} \qquad (4)$$

Let M be the total number of discrete filter positions used to obtain the image brightness values $\{I_i \mid i = 1, 2, ..., M\}$. If $M = 3$, equation 4 yields a linear system of equations that can be solved to obtain the parameters I_c, I_{sv}, and α. If $M > 3$, we have an over-determined linear system that can be solved to obtain more robust estimates of I_c, I_{sv}, and α, in the presence of image noise.[4]

From I_c and I_{sv}, we can obtain the maximum and minimum values of image brightness as:

$$I_{min} = I_c - I_{sv}, \qquad I_{max} = I_c + I_{sv} \qquad (5)$$

[3]For metals, the two Fresnel coefficients are nearly equal except close to the grazing angle (when ψ lies between 70 and 90 degrees). Thus, linear polarization based methods are generally not effective for metals. For dielectrics (non-conductors), however, the two Fresnel coefficients differ substantially except for near-normal angles of incidence (when ψ is less than 10-15 degrees).

[4]This formulation of the problem using vectors saves substantial computations compared to the non-linear formulation of the type $a + b \sin^2(\theta - \alpha)$ used in [Boult and Wolff, 91] and [Wolff, 90] that requires the use of iterative non-linear estimation techniques.

The *degree of polarization* at a scene point can be determined as [Born and Wolf, 65]:

$$\rho = \frac{I_{\max} - I_{\min}}{I_{\max} + I_{\min}} \quad (6)$$

The degree of polarization lies between 0 and 1 and can be used during highlight removal to classify points into those that are only diffuse ($\rho \ll 1$) and those that include a specular component. However, this measure must be used with care as both $I_{\min}$ and $I_{\max}$ include the constant specular component I_{sc} as well as the diffuse component I_d; varying either I_{sc} or I_d has the same effect on ρ.

We conclude this section with a note on previous work on highlight removal using polarization. A method for computing I_d and I_s by rotating the polarization filter, is presented in [Wolff and Boult, 91]. From the above discussion, we know that I_d and I_{sc} are both constant. They can be computed from I_c only if we know the ratio of the Fresnel coefficients, $q = F_\perp(\eta, \psi)/F_\parallel(\eta, \psi)$, for the corresponding scene point. The Fresnel coefficients are determined by the material properties of the scene point as well as the angle of incidence. Neither of these factors are known. To constrain the problem, [Wolff and Boult, 91] use all points (pixels) on a segmented highlight. They assume that both the diffuse component I_d as well as the Fresnel ratio q are constant under the highlight region, and estimate them using all points within the highlight region. This assumption, however, is unrealistic for two reasons. First, in real scenes, the diffuse component I_d within the highlight region may vary due to the curvature of the surface or due to texture on the surface. Secondly, the Fresnel ratio q cannot be assumed to be constant since the angle of incidence can vary substantially over large highlight regions resulting from extended sources in the scene. The latter of these problems was discussed in [Wolff, 90] but no solutions were proposed.

4 Color

We now discuss the role of color in the removal of specularities. We present an overview of the representation of diffuse and specular components in color space and discuss previous work on the removal of highlights using color. In contrast to gray-level images, color images represent wavelength (λ) dependence of the light reflected by a scene. Let $x(\lambda)$ be the spectral distribution of the light reflected by a scene point, and $s(\lambda)$ represent the response of the sensor to wavelength.

Typically, color images are obtained by using three filters with responses $r(\lambda)$, $g(\lambda)$, and $b(\lambda)$ that have peaks close to the wavelengths that humans perceive as "red," "green," and "blue." The resulting three brightness values measured by a sensor element constitute the color vector $\mathbf{I} = [I^r, I^g, I^b]$ for the corresponding point in scene. The three brightness values in the color vector are related to the spectral distribution of the reflected light as:

$$\begin{aligned} I^r &= \int x(\lambda)\, r(\lambda)\, s(\lambda)\, d\lambda \\ I^g &= \int x(\lambda)\, g(\lambda)\, s(\lambda)\, d\lambda \\ I^b &= \int x(\lambda)\, b(\lambda)\, s(\lambda)\, d\lambda \end{aligned} \quad (7)$$

Each brightness value includes a diffuse component and a specular component. Hence, in three-dimensional color space we have the following decomposition: $\mathbf{I} = \mathbf{I}_d + \mathbf{I}_s$.

The *dichromatic reflectance model* [Shafer, 85] suggests that, for dielectrics, the spectral distribution of the diffuse component is determined by the colorant in the surface whereas the specular component preserves the spectral distribution of the incident light. As a result, the two vectors $\mathbf{I}_d$ and $\mathbf{I}_s$ generally have different directions in color space. The two vectors will however have the same direction if, for instance, a gray object is illuminated by white light.

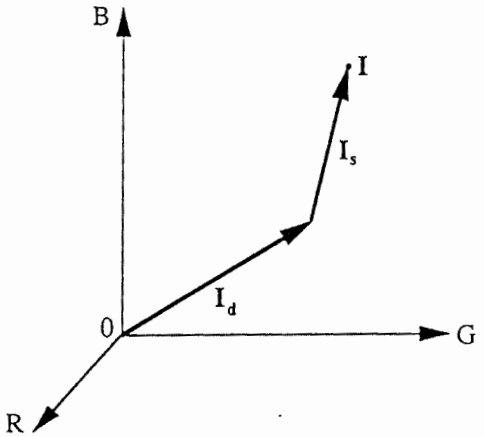

Figure 4: Diffuse and specular components in color space.

As discussed in the introduction, Klinker [Klinker, 88] and Gershon [Gershon, 87] independently used the dichromatic reflectance model to remove highlights from images using the two limbs of the skewed T in color space. These methods are based on two main assumptions: (a) the object is segmented away from the scene and has a single uniform diffuse color and (b) the diffuse component is near constant within the highlight region. These two assumptions must hold for a skewed T to be formed in color space. The first assumption is often violated in real scenes where objects may be textured or have patches with different reflectance properties. The second assumption can be valid only if the surface is very smooth, thus producing a very compact highlight. Even for marginally rough surfaces, the highlight is expected to include object points with a range of surface orientations. In such cases, the specular limb of the skewed T spreads into a wide cluster in color space that is difficult to separate from the diffuse limb. This last observation has also been made by Novak and Shafer [Novak and Shafer, 92] and Lee [Lee, 91].

5 Removal of Specularities using Color and Polarization.

In this section, we develop a method for removing specular reflections from images. The following assumptions are made.

(A) The scene consists of dielectric objects. Hence, the dichromatic reflectance model applies at each point, and the specular reflections and interreflections are polarized while the diffuse reflections are not. Further, the color

of the incident light at each scene point is different from the color of the material.

(**B**) The specular interreflections result from either the diffuse-specular mechanism *or* the specular-specular mechanism and not both. In the first case, the incident light is unpolarized and hence the Fresnel ratio is simply $q = F_\perp(\eta,\psi)/F_\parallel(\eta,\psi)$. In the second case, the incident light is partially polarized and the effective Fresnel ratio for the surface point is $q = aF_\perp(\eta,\psi)/bF_\parallel(\eta,\psi)$. The parameters a and b account for the partial polarization of the incident light.

(**C**) Fresnel coefficients $F_\perp(\eta,\psi)$ and $F_\parallel(\eta,\psi)$ are independent of the wavelength of incident light. This assumption is reasonable (see [Driscoll, 78]) since we are assuming dielectrics and operating in the visible-light spectrum. Assumptions (B) and (C) result in the Fresnel ratio q being equal for all three color bands.

The color of the specular component is computed *locally* at each pixel. This places strong constraints on the diffuse component $\mathbf{I}_d$. Neighboring diffuse points that satisfy these constraints are then used to compute the diffuse component. This approach has the following advantages over all of the previous methods for highlight removal:

- In contrast to the previous methods based either on color or polarization, we do not assume that the diffuse component $\mathbf{I}_d$ is constant within each highlight region. In fact, the surfaces could be textured with patches of different materials underlying the highlights.
- The specular color of each point is computed independently. This *local* approach does not require prior segmentation of either highlights or objects in the scene. Further, the highlights need not be compact; the method can handle substantial surface roughness conditions.

We describe the technique for specularity removal by focusing on a single image point $\mathbf{x}$. The same procedure is applied independently to all image points. The color vector for the image point is: $\mathbf{I} = \mathbf{I}_d + \mathbf{I}_s$. Given the above assumptions, the Fresnel ratio q is the same for all three color bands. Hence the cosine term in equation 4 will be in phase for the 3 color bands and for the polarization filter position θ_i we have the color vector: $\mathbf{I}_i = \mathbf{I}_c + \mathbf{I}_{sv} \cos 2(\theta_i - \alpha)$. In our experiments, we have used 6 or more polarizer positions. This gives us an over-determined linear system of equations that are solved to obtain robust estimates of $\mathbf{I}_c$, $\mathbf{I}_{sv}$, and α. The color vectors corresponding to maximum and minimum polarization (see equation 5) are: $\mathbf{I}_{max} = \mathbf{I}_c + \mathbf{I}_{sv}$ and $\mathbf{I}_{min} = \mathbf{I}_c - \mathbf{I}_{sv}$.

Two tests are used to determine if the image point $\mathbf{x}$ is to be processed any further. First, a degree of polarization ρ (expression 6) is computed for each of the three color bands. If the largest of three ρ estimates is less than a threshold value T_1, the point is not sufficiently polarized and is assumed to be purely diffuse. In this case, the next image point is examined.

If the degree of polarization of the point $\mathbf{x}$ is greater than T_1, the angle β subtended by the vector $\mathbf{k} = \mathbf{I}_{max} - \mathbf{I}_{min}$ from the origin $\mathbf{0}$ is computed (see Figure 5). If β is less than a threshold T_2, the color of the specular component is very similar to that of the diffuse component $\mathbf{I}_d$, and the dichromatic model cannot be used with confidence.

On the other hand, if the point $\mathbf{x}$ is polarized and its β value is not small, we proceed to compute its diffuse component $\mathbf{I}_d$. If we can determine $\mathbf{I}_{sc}$, then the diffuse component can be computed as $\mathbf{I}_d = \mathbf{I}_c - \mathbf{I}_{sc}$. Recall that the specular components $\mathbf{I}_{sc}$ and $\mathbf{I}_{sv}$ satisfy $\mathbf{I}_{sc} + \mathbf{I}_{sv} = \mathbf{I}_{sc}\, q$. Unfortunately, the Fresnel coefficient q is not known as it depends on the material properties and the angle of incidence. Though we have estimates of $\mathbf{I}_c$ and $\mathbf{I}_{sv}$, we do not have a simple way of determining $\mathbf{I}_d$. The color measurements $\mathbf{I}_i$ obtained by rotating the polarization filter lie on a straight line L (see Figure 5) in color space. The diffuse component $\mathbf{I}_d$ is unaffected by rotations of the polarizer; only the specular component $\mathbf{I}_s$ varies. The specular component varies along a straight line since the cosine functions in the three color bands are in phase (assumptions **B** and **C**).

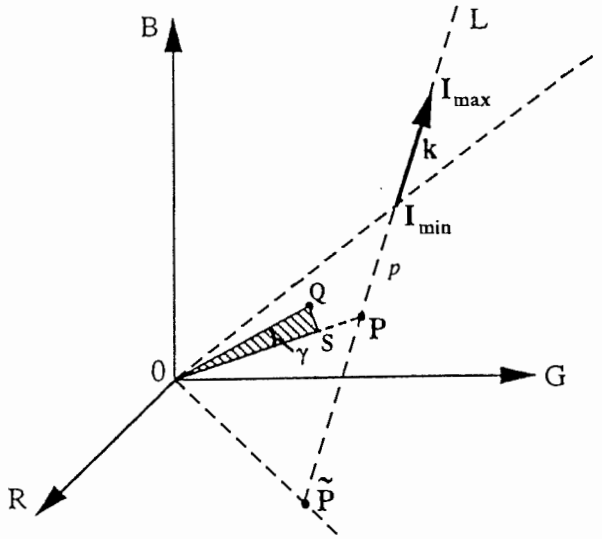

Figure 5: Using neighboring points and specular line constraint to compute the diffuse component $\mathbf{I}_d$.

Though we are unable to compute the diffuse component locally, the specular line gives useful constraints on the diffuse component. Assume that the diffuse component of $\mathbf{x}$ corresponds to the point $\mathbf{P}$. The position of $\mathbf{P}$ on the specular line L can be parametrized as follows: $\mathbf{P} = \mathbf{I}_{min} - p\mathbf{k}$ where p ($0 \le p \le \tilde{p}$) is the distance of $\mathbf{P}$ from $\mathbf{I}_{min}$ as shown in Figure 5 and $\tilde{p}$ is defined as:

$$\tilde{p} = Min\left(\frac{I_{min}{}^{r}}{k^{r}}, \frac{I_{min}{}^{g}}{k^{g}}, \frac{I_{min}{}^{b}}{k^{b}}, \right). \qquad (8)$$

$\tilde{p}$ determines the point $\tilde{\mathbf{P}}$ that is the intersection of the specular line L with one of the three planes of the color space. In the example shown in Figure 5, the specular line intersects with the $R-G$ plane. In general, however, L could intersect any one of the three planes that constitute the color space. Define $\mathcal{P}_1$ to be the plane that passes through the points ($\mathbf{0}$, $\mathbf{I}_{min}$, $\mathbf{I}_{max}$). The expression for $\mathcal{P}_1$ is: $A\,I^r + B\,I^g + C\,I^b = 0$,

where:

$$\begin{aligned} A &= k^{g} {I_{min}}^{b} - k^{b} {I_{min}}^{g} \\ B &= k^{b} {I_{min}}^{r} - k^{r} {I_{min}}^{b} \\ C &= k^{r} {I_{mjn}}^{g} - k^{g} {I_{mjn}}^{r}. \end{aligned} \quad (9)$$

Since we do not have sufficient constraints to compute the diffuse component $\mathbf{I}_d$ of the point $\mathbf{x}$ from the color measurements $\mathbf{I}_i$, we use neighboring image points that satisfy the following conditions:

(1) A neighboring image point $\mathbf{y}$ can be used if we know its the diffuse component $\mathbf{Q}$. This occurs if $\mathbf{y}$ has a low degree of polarization ρ and hence can be assumed to be purely diffuse, or if its diffuse component has already been computed.

(2) In color space, the vector $\mathbf{Q}$ must lie on the plane $\mathcal{P}_1$. Further, it must lie between the vectors $\mathbf{I}_{min}$ and $\tilde{\mathbf{P}}$, since the diffuse vector $\mathbf{I}_d$ of the point $\mathbf{x}$ lies on the line L between the points $\mathbf{I}_{min}$ and $\tilde{\mathbf{P}}$. Note, however, that $\mathbf{Q}$ need not lie on line L; it can lie inside or outside the triangle $(\mathbf{0}, \mathbf{I}_{min}, \tilde{\mathbf{P}})$.

If these conditions are satisfied, the neighboring point $\mathbf{y}$ is *assumed* to have the same diffuse color as the point $\mathbf{x}$. Then, the line passing through $\mathbf{Q}$ and the specular line L intersect to give $\mathbf{P}$, an estimate of the diffuse component of $\mathbf{x}$.

Due to noise in the color and polarization measurements, the diffuse component $\mathbf{Q}$ of a neighboring point is not expected to exactly satisfy the above conditions. To accommodate for such discrepancies, we compute the angle γ subtended by $\mathbf{Q}$ with respect to $\mathcal{P}_1$ (see Figure 5):

$$\sin\gamma = \frac{A\,Q^{r} + B\,Q^{g} + C\,Q^{b}}{\sqrt{A^2 + B^2 + C^2}} \quad (10)$$

If γ is larger than a threshold value T_3, $\mathbf{Q}$ is not used any further.[5] If γ is small ($\gamma \leq T_3$), the point $\mathbf{P}$ is computed by extending the vector $\mathbf{S}$ to intersect the line L as shown in Figure 5. This can be done without computing the projection $\mathbf{S}$ of $\mathbf{Q}$ on the plane $\mathcal{P}_1$. Consider the plane $\mathcal{P}_2$ that passes through the points $(\mathbf{0}, \mathbf{Q}, \mathbf{P})$. It can be expressed as $D\,I^{r} + E\,I^{g} + F\,I^{b} = 0$, where:

$$\begin{aligned} D &= Q^{g} P^{b} - Q^{b} P^{g} \\ E &= Q^{b} P^{r} - Q^{r} P^{b} \\ F &= Q^{r} P^{g} - Q^{g} P^{r} \end{aligned} \quad (11)$$

Since the planes $\mathcal{P}_1$ and $\mathcal{P}_2$ are perpendicular, we have: $A\,D + B\,E + C\,F = 0$, which can be expanded using equations 9 and 11. By substituting the expression for $\mathbf{P}$ given by equation 8 in the expansion, a solution for the line parameter p is directly obtained:

$$p = \frac{\begin{pmatrix} A\left(Q^{g}{I_{min}}^{b} - Q^{b}{I_{min}}^{g}\right) \\ +B\left(Q^{b}{I_{min}}^{r} - Q^{r}{I_{min}}^{b}\right) \\ +C\left(Q^{r}{I_{min}}^{g} - Q^{g}{I_{min}}^{r}\right) \end{pmatrix}}{\begin{pmatrix} A\left(Q^{g}k^{b} - Q^{b}k^{g}\right) \\ +B\left(Q^{b}k^{r} - Q^{r}k^{b}\right) \\ +C\left(Q^{r}k^{g} - Q^{g}k^{r}\right) \end{pmatrix}} \quad (12)$$

[5] Note that we have used the angle γ rather than the distance of $\mathbf{Q}$ from $\mathcal{P}_1$. This is because a neighboring point may have a very small diffuse component that lies close to the origin $\mathbf{0}$ and as a result is also close to $\mathcal{P}_1$. Such a point could have relatively large errors due to image noise and must not be used in the computation of $\mathbf{I}_d$.

The above process is repeated for all neighboring diffuse components $\mathbf{Q}_j$ that satisfy conditions (1) and (2). The result is a set of estimates $\{p_j \mid j = 1, 2, ..., N\}$. If $N < T_4$ (we use $T_4 = 3$ in our implementation) there are not enough neighboring diffuse components to compute a robust estimate of $\mathbf{I}_d$ for the point $\mathbf{x}$. If $N \geq T_4$, the mean and standard deviation of p_j are computed as:

$$\begin{aligned} \overline{p} &= \frac{\sum_{j=1}^{N} w_j\, p_j}{\sum_{j=1}^{N} w_j} \\ \sigma_p &= \frac{\sum_{j=1}^{N} \left(\overline{p} - p_j\right)^2}{N-1} \end{aligned} \quad (13)$$

where the weight w_j given to each p_j equals the magnitude of the corresponding diffuse component, $\| \mathbf{Q}_j \|$. The mean value $\overline{p}$ is accepted if the standard deviation σ_p is less than a threshold T_5, i.e. the estimates p_j form a compact cluster on the line L. This constraint is used to ensure that different diffuse colors in the neighborhood of $\mathbf{x}$ that happen to lie close to the plane $\mathcal{P}_1$, are not used together to obtain an erroneous estimate of $\mathbf{I}_d$. Once $\overline{p}$ has been determined, the diffuse component $\mathbf{I}_d = \mathbf{P}$ of the image point $\mathbf{x}$ is obtained using equation 8.

The algorithm proposed in this section is applied to all points in the image. Not all image points end up with an $\mathbf{I}_d$ estimate. An image point may lie in the middle of a very large highlight, in which case, it may not have a sufficient number of neighbors with diffuse colors that satisfy conditions (1) and (2) or produce a compact cluster of intersection points on the line L. Hence, we apply the algorithm repeatedly to the image points. This iterative approach is effective in the case of complex scenes; each iteration provides a new set of computed diffuse colors thus increasing the likelihood of finding neighboring diffuse colors in the next iteration. The iterations are discontinued when no new diffuse estimates are obtained.

6 Experimentation

This section presents experimental results obtained using the proposed algorithm. Details of our experimental setup, implementational problems, and calibration procedures can be found in [Nayar *et al.*, 93]. Here, we discuss some details of the algorithm implementation. This is followed by results obtained by applying the algorithm to scenes with textured objects, primary (source) specularities, and secondary (inter-reflection) specularities.

6.1 Implementation Details

For each scene, a set of color images are obtained by rotating the polarization filter. The polarization parameters I_{min}, I_{max}, and α are computed for each color channel. These parameters are computed using the linear least squares (LS) fitting method. The results of the polarization fitting are 6 images for each color channel, I_{min}, I_{max}, I_{avg} ($= \frac{I_{max}+I_{min}}{2}$), ρ (percent polarization), phase (the angle α), and RMSE (root mean square error in fitting). Of these, only the I_{max} and I_{min} images are directly used by the specularity removal algorithm. The others are used only to debug the algorithm

and analyze the results. The I_{avg} image is what would be obtained without a polarizer but with a 50% neutral density filter instead.

The algorithm requires labeling of points as purely diffuse or partially specular, which is done by using the degree of polarization ρ. Since ρ depends on I_{min}, the noise level in ρ varies with I_{min}. Rather than using a fixed threshold T_1 on ρ to identify partially specular points, our implementation uses a threshold that varies with I_{min}. T_1 varies from around 5% for points with $I_{min} > 200$ to around 10% for points when $I_{min} \approx 20$.

The algorithm has three other thresholds that affect its performance; the threshold T_2 for the angle β, the threshold T_3 for the angle γ, and T_4 for the standard deviation σ_p. The algorithm is not too sensitive to T_2 and T_4, which have been set at 0.08 and $arccos(0.99)$, respectively. The angle threshold T_3, which determines if points lie close to the plane $\mathcal{P}_1$, has a strong affect on the quality of computed results as well as the computation time. The current implementation starts with a relatively small threshold value (T_3=0.02), and doubles it after every 10 iterations.

6.2 Experimental Results

In the each of the following examples, the images obtained after fitting the polarization parameters are used to computed the diffuse color image $\mathbf{I}_d$. This image is then subtracted from the average color image $\mathbf{I}_{avg}$ to obtain the specular color image $\mathbf{I}_s$. $\mathbf{I}_{avg}$ is the image we would obtain if the polarization filter were not used. All the results discussed here are presented in color in [Nayar *et al.*, 93]. The first example is shown in Figure 6(a). It is the $\mathbf{I}_{avg}$ image of a cup with a flowered pattern on it. The petals of the flower are of different colors and within each petal there is a moderate amount of diffuse color variation. Along the middle of the cup is a large highlight. Figure 6(b) shows the diffuse image $\mathbf{I}_d$ computed using the proposed algorithm. The algorithm was clearly successful in removing the highlight despite the texture underlying the highlight.

The second example includes strong interreflection effects. Figure 6(c) shows the original image ($\mathbf{I}_{avg}$) of a scene including a blue plastic plate and a part of the McBeth color chart. Color patches on the chart are reflected by the plate. There are pieces of plastic tape (some dark reddish and others black) on the plate. Also visible is a film canister, which interreflects portions of the color chart as well as the surrounding environment. Figure 6(d) and Figure 6(e) show the diffuse and specular components computed by the algorithm. We see that, despite the strong interreflections, these images are quite accurate. We do, however, see that a primary highlight on the left side of plate has not been completely removed. This may have been caused by very high brightness values in the highlight region for which the sensor calibration is not reliable.

The final example is shown in Figure 6(f)–(h). Several objects with different colors are placed close to one another in a large red box. Figure 6(f) shows the $\mathbf{I}_{avg}$ image of this complex scene. Figure 6 (g) and (h) show the diffuse and specular components computed by the algorithm. The primary highlights on the blue torus (in front) and red torus (top left corner) are accurately removed. Errors in the separation are however seen on the occluding boundary of the torus. This results from the strong polarization of the diffuse component on the occluding boundary; the assumption that the diffuse component is unpolarized is violated. Also note that the results are not as good for the interreflections of the blue (front) torus and the marker pen. This is due to interreflections between the two walls of the red box that cause the dichromatic reflectance assumption to be violated (see, [Nayar *et al.*, 93] for details).

References

[Bajcsy *et al.*, 90] R. Bajcsy, S. W. Lee, and A. Leonardis "Color image segmentation with detection of highlights and local illumination induced by interreflections," *Proc. of International Conf. on Pattern Recognition,* Atlantic City, NJ, June 1990.

[Born and Wolf, 65] M. Born and E. Wolf, *Principles of Optics,* London:Pergamon, 1965.

[Boult and Wolff, 91] T. Boult and L.B. Wolff, "Physically based edge labeling," *Proc. CVPR,* pp. 656-663 1991.

[Gershon, 87] R. Gershon, *The Use of Color in Computational Vision,* PhD thesis, University of Toronto, 1987.

[Marvin, 88] *Handbook of Laser Science and Technology,* edited by M.J. Marvin, CRC Press Inc., Vol. 4, 1988.

[Driscoll, 78] *Handbook of Optics,* edited by W.G. Driscoll, McGraw Hill Inc., 1978.

[Klinker, 88] G. J. Klinker, *A Physical Approach to Color Image Understanding,* PhD thesis, Carnegie Mellon University, 1988.

[Koshikawa, 79] K. Koshikawa, "A polarimetric approach to shape understanding," *Proc. IJCAI,* pp. 493-495, 1979.

[Nayar *et al.*, 91] S. K. Nayar, K. Ikeuchi, T. Kanade, "Surface Reflection: Physical and Geometrical Perspectives," *IEEE Trans. on Pattern Analysis and Machine Intelligence*, Vol. 13, No. 7, pp. 611-634, July 1991.

[Nayar *et al.*, 93] S.K. Nayar, X.S. Fang, and T.E. Boult. Separation of reflection components using color and polarization. Technical Report CUCS-058-92, Center for Research in Intelligent Systems, Columbia University, Nov 1992.

[Novak and Shafer, 92] C. Novak and S. Shafer, "Anatomy of a Color Histogram," *Proc. CVPR,* pp. 599-605, 1992.

[Shafer, 85] S. Shafer, "Using color to separate reflection components," *Color Research and Applications,* Vol. 10, pp. 210-218, 1985.

[Lee, 91] S.W. Lee, *Understanding of Surface Reflection in Computer Vision by Color and Multiple Views,* PhD thesis, University of Pennsylvania, 1991.

[Torrance and Sparrow, 67] K. Torrance and E. Sparrow, "Theory for Off-Specular Reflection from Roughened Surfaces," *Journal of the Optical Society of America*, No. 57, pp. 1105-1114, 1967.

[Wolff and Boult, 91] L.B. Wolff and T. Boult, "Constraining Object Features using a Polarization Reflectance Model," *IEEE Trans. on Pattern Analysis and Machine Intelligence*, Vol. 13, No. 7, pp. 635-657, July 1991.

[Wolff, 90] L.B. Wolff, *Polarization methods in computer vision,* PhD thesis, Columbia University, 1990.

(a) I_{avg} for cup

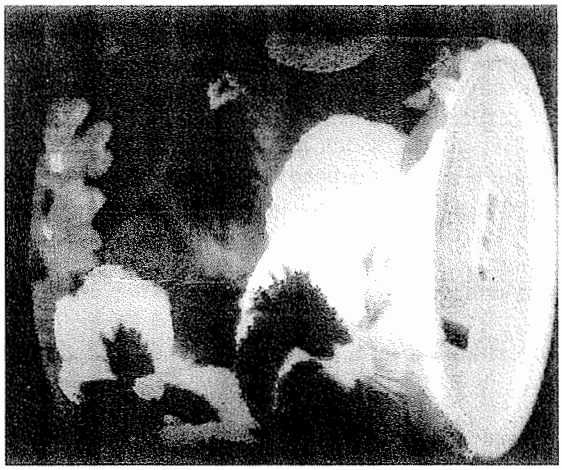

(b) I_d (diffuse image)

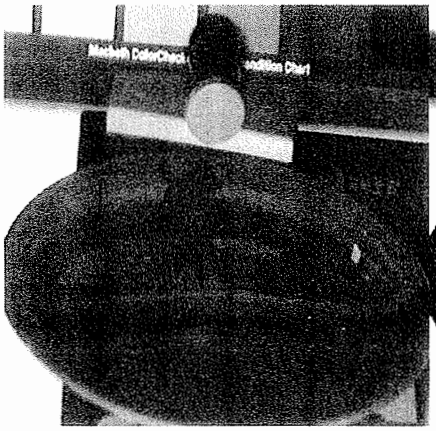

(c) I_{avg} for plate scene

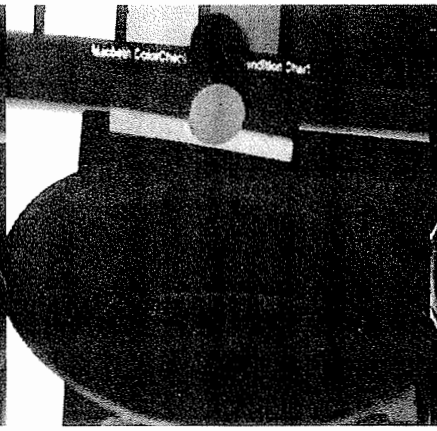

(d) I_d (diffuse image)

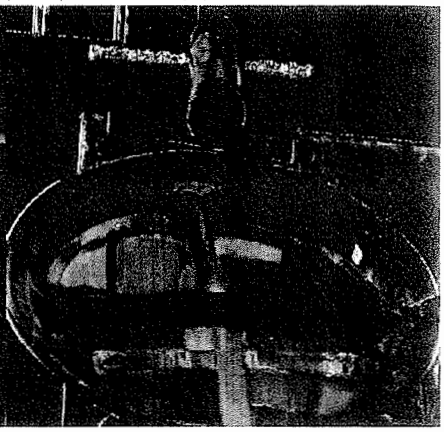

(e) I_s (specular image)

(f) I_{avg} for complex scene

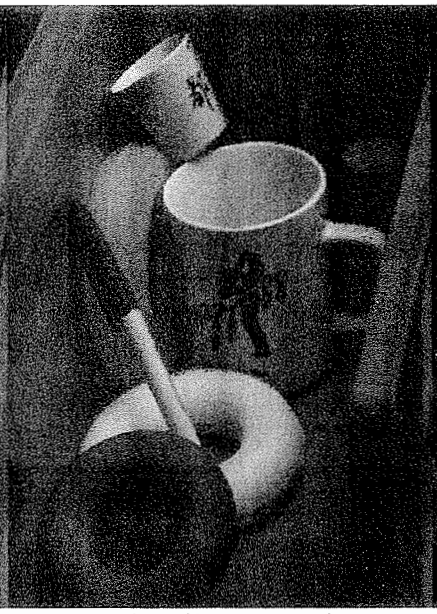

(g) I_d (diffuse image)

(h) I_s (specular image)

Figure 6: Experimental results: (a)-(b) textured cup; (c)-(e) scene with plastic plate and color chart; (f)-(h) complex scene with multiple highlights and strong interreflections.