

Using Color to Separate Reflection Components

Steven A. Shafer *
Computer Science Department
University of Rochester
Rochester, NY 14627

TR 136
April 1984

* The author's permanent address is: Computer Science Department
Carnegie-Mellon University
Pittsburgh, PA 15213

Using Color to Separate Reflection Components

Steven A. Shafer*

Computer Science Department

University of Rochester

Rochester, New York 14627

2 April 1984

Abstract

This paper presents an algorithm for analyzing a standard color image to determine intrinsic images of the amount of interface ("specular") and body ("diffuse") reflection at each pixel. The interface reflection represents the highlights from the original image, and the body reflection represents the original image with highlights removed. Such intrinsic images are of interest because the geometric properties of each type of reflection are simpler than the geometric properties of intensity in a black-and-white image.

The algorithm is based upon a physical model of reflection which states that two distinct types of reflection -- interface and body reflection -- occur, and that each type can be decomposed into a relative spectral distribution and a geometric scale factor. This model is far more general than typical models used in computer vision and computer graphics, and includes most such models as special cases. In addition, the model does not assume a point light source or uniform illumination distribution over the scene.

The properties of spectral projection into color space are used to derive a new model of pixel-value color distribution, and this model is exploited in an algorithm to derive the intrinsic images. Suggestions are provided for extending the model to deal with diffuse illumination and for analyzing the intrinsic images of reflection.

(*) The author's permanent address is: Computer Science Department, Carnegie-Mellon University, Pittsburgh, Pennsylvania 15213.

This research was sponsored by the Defense Advanced Research Projects Agency (DOD), monitored by the Office of Naval Research under Contract N00014-82-K-0193.

The views and conclusions contained in this document are those of the author and should not be interpreted as representing the official policies, either expressed or implied, of the Defense Advanced Research Projects Agency or the US Government.

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM	
1. REPORT NUMBER TR 136	2. GOVT ACCESSION NO.	3. RECIPIENT'S CATALOG NUMBER	
4. TITLE (and Subtitle) Using Color to Separate Reflection Components		5. TYPE OF REPORT & PERIOD COVERED technical report	
		6. PERFORMING ORG. REPORT NUMBER	
7. AUTHOR(s) Steven A. Shafer		8. CONTRACT OR GRANT NUMBER(s) N00014-82-K-0193	
9. PERFORMING ORGANIZATION NAME AND ADDRESS Computer Science Department University of Rochester Rochester, NY 14627		10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS	
11. CONTROLLING OFFICE NAME AND ADDRESS Defense Advanced Research Projects Agency 1400 Wilson Blvd. Arlington, VA 22209		12. REPORT DATE April 2, 1984	
		13. NUMBER OF PAGES 23	
14. MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office) Office of Naval Research Information Systems Arlington, VA 22217		15. SECURITY CLASS. (of this report) unclassified	
		15a. DECLASSIFICATION/DOWNGRADING SCHEDULE	
16. DISTRIBUTION STATEMENT (of this Report) Distribution of this document is unlimited.			
17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)			
18. SUPPLEMENTARY NOTES none			
19. KEY WORDS (Continue on reverse side if necessary and identify by block number)			
20. ABSTRACT (Continue on reverse side if necessary and identify by block number) This paper presents an algorithm for analyzing a standard color image to determine intrinsic images of the amount of interface ("specular") and body ("diffuse") reflection at each pixel. The interface reflection represents the highlights from the original image, and the body reflection represents the original image with highlights removed. Such intrinsic images are of interest because the geometric properties of each type of reflection are simpler than the geometric properties of intensity in a black-and-white image. (over)			

20. Abstract (cont.)

The algorithm is based upon a physical model of reflection which states that two distinct types of reflection--interface and body reflection--occur, and that each type can be decomposed into a relative spectral distribution and a geometric scale factor. This model is far more general than typical models used in computer vision and computer graphics, and includes most such models as special cases. In addition, the model does not assume a point light source or uniform illumination distribution over the scene.

The properties of spectral projection into color space are used to derive a new model of pixel-value color distribution, and this model is exploited in an algorithm to derive the intrinsic images. Suggestions are provided for extending the model to deal with diffuse illumination and for analyzing the intrinsic images of reflection.

Table of Contents

1. Introduction	1
1.1 Previous Work in Color Image Understanding	2
2. Reflection and the Dichromatic Model	3
2.1 Physical Properties of Reflection	3
2.1.1 An Aside on Terminology	6
2.2 The Dichromatic Reflection Model	7
2.3 Previously Used Reflection Models	9
3. Pixel Values in Color Space	11
3.1 Spectral Projection	11
3.2 The Dichromatic Model in Color Space	12
3.3 Analyzing Color Pixel Values	14
3.3.1 Assumptions of the Dichromatic Analysis Algorithm	14
3.4 Summary	15
4. Applying the Dichromatic Analysis Algorithm	16
4.1 Extending the Model for Diffuse Illumination	16
4.2 Analysis of the Intrinsic Reflection Images	17
5. Conclusions	19
5.1 Acknowledgements	20
6. Bibliography	21

List of Figures

Figure 2-1: Reflection of Light from an Inhomogeneous Material	3
Figure 2-2: Body Reflection is Not Really Isotropic (from [13])	5
Figure 2-3: A Paint Layer Producing Two Interface Reflections (from [30])	5
Figure 2-4: Photometric Angles	6
Figure 3-1: Pixel Values on a Surface Lie on a Parallelogram in Color Space	12
Figure 3-2: Position Within Parallelogram Determines Magnitudes	13
Figure 4-1: Color Space Parallelogram With Diffuse Illumination	16

1. Introduction

When we look around us, the surfaces we see are typically glossy. They may seem to be very shiny, fairly matte, or anywhere in between, but virtually all the surfaces around us exhibit highlights to varying degrees. These highlights are most pronounced when the surface normal bisects the angle between the direction of illumination and the direction of view, making the position and intensity of highlights very sensitive to viewing geometry. This causes problems with many low-level computer vision methods such as segmentation (which typically assumes uniform or smoothly varying intensity across a surface) or stereo and motion analysis (which attempt to match images taken from different viewpoints).

Highlights are not the only source of intensity variation across a surface. Even with uniform illumination of a matte, uniformly colored surface, there will be smooth shading due to the angle of incidence of the incoming illumination relative to the surface normal.

It would be very useful to be able to separate the effects of shading from highlights. This might result in *intrinsic images* telling, at each pixel, the intensity of the shading and the intensity of the highlight at that point. This separation was first suggested by Barrow and Tenenbaum [1], and would effectively produce an image of the highlights, and an image with the highlights removed. These would be useful for analysis since each of these phenomena is more simply related to the angles of illumination and viewing than is their sum (which is measured in a black-and-white image). In addition, the relative insensitivity of shading to viewpoint would make the shading intrinsic image an ideal candidate for stereo or motion analysis.

It is frequently observed that highlights have a different color from the characteristic color of a surface (which is related to shading). In this paper, we show how a simple, rather general model of reflection, called the "Dichromatic Reflection Model", can be used to determine intrinsic images of the two types of reflection from a standard color image (i.e. red-green-blue separation).

The analysis uses the properties of *spectral projection*, the process whereby color pixel values are determined from the spectral power distribution (SPD) of incoming light. Combining the Dichromatic Reflection Model with spectral projection results in a new model of pixel value distribution in R-G-B color space. The model predicts that pixel values from pixels on a single surface will lie on a parallelogram in color space, and that the position of any pixel's color within that parallelogram yields the coefficients of the two types of reflection.

A simple algorithm is presented which utilizes this model of pixel values to determine the desired intrinsic images. The model and the algorithm are then extended to deal with diffuse illumination and shadows. Additional work for the future includes implementation of the algorithm and verification of the model with real images.

For additional background information, the reader interested in radiometry in general is referred to [9], [12], and [22]; while *appearance measurement* (gloss and color) is discussed in [10], [14], [16], and [35].

1.1 Previous Work in Color Image Understanding

Color image understanding in the past has not been based on general models of reflection. Most of the work has been the application to three-dimensional color space of algorithms originally developed and used for analyzing monochrome images. This includes primarily edge detection [21] and clustering [5, 18, 20, 23] (etc.). In such work, image regions or edges are identified by distances between pixel values in color space, without appeal to any model of color generation in the scene.

Color has also been used for object labelling based on known object colors or object colors measured in the image [19, 26, 31, 32, 34, 36]. This approach, known as "spectral signature analysis" in remote sensing, uses the known reflection properties of materials of interest in a particular domain. It depends on having few different types of materials in the scene, and having prior knowledge of their spectral reflectance.

The properties of color space transformations have also been studied [17, 24], although no such transformations have sufficed to "solve" the image segmentation or labelling problems.

The only previous work in utilizing general properties of reflected color has been in the form of simple statements such as "far-away outdoor objects look bluish", "outdoor shadows are bluish", and "natural colors tend to be desaturated" [19, 31]. In the psychological field, Rubin and Richards proposed a method for using color to determine changes in material across an image [27]; however, while the method is quite interesting, their assumptions appear to be restrictive.

The approach described in this paper is significant because it is applicable to many common types of materials, without prior knowledge about their colors, and because it is based on a very general model of reflection in the real world.

Section 2.3 presents a discussion of previous work in reflectance models for computer graphics and image understanding.

2. Reflection and the Dichromatic Model

In this chapter, we present a brief account of the physics of reflection, followed by the Dichromatic Reflection Model which captures certain aspects of the reflection process. The validity and assumptions of the model are discussed, and it is compared with other reflection models used in computer graphics and image understanding.

2.1 Physical Properties of Reflection

The methods of this paper deal with materials which are optically *inhomogeneous*, meaning that light interacts both with a *medium* that comprises the bulk of the surface matter, and with particles of a *colorant* that produce scattering and coloration. Many common materials can be described this way, including most paints, varnishes, paper, ceramics (including porcelain), plastics, etc. Materials that are homogeneous are not included in this discussion; thus metals and many crystals are not amenable to the analysis presented herein. We will also limit this discussion to opaque surfaces, which transmit no light from one side to the other.




Figure 2-1: Reflection of Light from an Inhomogeneous Material

Although we sometimes think of a visual surface as a plane, this is an approximation useful only at the macroscopic level. To understand reflection from inhomogeneous materials, we will instead view the surface as having a definite thickness (Figure 2-1) [14, 16, 35].


When light strikes a surface, it must first pass through the interface between the air and the surface

medium. Because the medium's index of refraction differs from that of the air, some of the light will be reflected at the interface producing *interface reflection* as shown in figure 2-1. The direction of such reflection will be in the "perfect specular direction" relative to the local surface normal, i.e. reflected about the surface normal. Note that most materials are optically "rough", with local surface normals that differ from the macroscopic or reference surface normal. The local perfect specular direction is therefore somewhat different from the macroscopic perfect specular direction, so that the interface reflection will be somewhat scattered at the macroscopic level.

The amount of light reflected at the interface is governed by Fresnel's laws, which relate interface reflectance to the angle of incidence (relative to the local surface normal), the index of refraction of the material, and the polarization of the incoming illumination [15]. Interface reflectance is thus a function of wavelength of light, since index of refraction generally depends on wavelength, but the amount is typically constant to within a few percent across the visible spectrum; acrylic plastic, for example, has an index of refraction of 1.485 at 400 nm (the blue end of the visible spectrum) and 1.505 at 700 nm (the red end) [3, 16, 35], producing Fresnel reflection coefficients of 3.8% and 4.1% at the ends of the spectrum. Interface reflection is frequently assumed to be constant with respect to wavelength because of the small magnitude difference, and is thus said to have the same color (relative SPD) as the illuminant [7, 14, 16]. The effect of polarization is more severe, and interface reflection tends to be highly polarized, especially for large angles of incidence ("grazing angles") [15, 35].

The light that penetrates through the interface passes through the medium, where it undergoes scattering from the colorant, and eventually is either transmitted through the material (if it is not opaque), absorbed by the colorant, or re-emitted through the same interface by which it entered, producing *body reflection* as shown in figure 2-1. The geometric distribution of this body reflection is sometimes assumed to be isotropic, i.e. independent of viewing direction, although some work is under way to produce improved models of scattering [7, 10]. The color of the body reflection is generally different from that of the illumination, since interactions with colorant particles result in absorption with a probability dependent on wavelength. Body reflection is usually considered to be unpolarized.


Real reflection tends to be more complex than described above. For example, real body reflection is not isotropic (figure 2-2 shows curves of reflection as a function of viewing direction) [13], and some materials exhibit several interfaces producing interface reflection (Figure 2-3) [30]. However, the explanation provided above is a very useful approximation.



Goniophotometric curves of directional reflectance as a function of angle

Figure 2-2: Body Reflection is Not Really Isotropic (from [13])

This figure shows some curves of reflection as a function of viewing angle for various materials. In each case, the direction of illumination is held constant as the distribution of reflection is measured.



Electron micrograph of section of an air-drying polyurethane paint film, illustrating the two boundary reflectances (r_1 and r_2) which contribute to gloss.

Figure 2-3: A Paint Layer Producing Two Interface Reflections (from [30])

2.1.1 An Aside on Terminology

This paper adopts the terms "interface" and "body" reflection rather than the more common terms "specular" and "diffuse" reflection. While the latter have gained some popularity in the literature of appearance measurement and computer imaging (graphics and vision), they have some severe problems which render them controversial. "Specular reflection" may mean any of three things:

- *interface reflection* -- the most common usage
- *interface reflection in perfect specular direction at macroscopic level* -- the usage most common in gloss measurement
- *interface reflection where surface is locally optically smooth* -- the usage in scattering theory and laser optics [2].

Similarly, the term "diffuse reflection" means "reflected light scattered over a large solid angle", and thus may apply equally well to body and to interface reflection. The author, having experienced numerous communication problems due to the use of the terms specular and diffuse, has ultimately adopted the terminology used in this paper, which is sometimes found in the more technical optics literature.




Figure 2-4: Photometric Angles

The terminology used throughout this paper for describing reflectance geometry is illustrated in figure 2-4, which defines the following angles [11]:

- *the angle of incidence, i* -- the angle between the illumination direction I and the surface normal N
- *the angle of emittance, e* -- the angle between N and the viewing direction V
- *the phase angle, g* -- the angle between I and V

- the off-specular angle, s -- the angle between V and the direction of (macroscopic) perfect specular reflection J .

We also use the standard symbol λ to refer to the wavelength of light.

2.2 The Dichromatic Reflection Model

We now propose a simple mathematical model of reflectance, based on the above discussion, called the *Dichromatic Reflection Model*.

The Dichromatic Reflection Model states:

$$L(\lambda, i, e, g) = L_i(\lambda, i, e, g) + L_b(\lambda, i, e, g) \quad (2-1)$$

$$= m_i(i, e, g) c_i(\lambda) + m_b(i, e, g) c_b(\lambda) \quad (2-2)$$

This model represents two statements about reflected light, as expressed by the two parts of the equation:

1. Equation (2-1) says that the total radiance L of the reflected light is the sum of two independent parts: the radiance L_i of the light reflected at the interface and the radiance L_b of the light reflected from the surface body.
2. Equation (2-2) says that each of these components of the light can be decomposed into two parts:
 - a. *composition* -- a relative spectral power distribution c_i or c_b which depends only on wavelength but is independent of geometry, and
 - b. *magnitude* -- a geometric scale factor m_i or m_b which depends only on geometry and is independent of wavelength.

Intuitively, the Dichromatic Reflectance Model says that there are two independent reflection processes, and that each has a characteristic color whose magnitude, but not spectral distribution, varies with the directions of illumination and view.

In the remainder of this chapter we will address the scope of this model and its validity; in the next chapter we will see how this model may be exploited to determine intrinsic images of m_i and m_b , the amount of interface and body reflection at each pixel.

The Dichromatic Reflection Model assumes the following:

- The surface is an opaque, inhomogeneous medium with one significant interface.
- The surface is not optically active, i.e. it has no fluorescence or thin-film properties, and it is uniformly colored, i.e. it has a uniform distribution of the colorant.
- Reflection from the surface is isotropic with respect to rotation about the surface normal.

- There is no inter-reflection among surfaces.
- There is a single light source, i.e. no diffuse ("ambient") illumination, and the relative spectral power distribution $S(\lambda)$ of the illumination is constant across the scene.

The assumptions about the surface are typical for reflectance models and not too unrealistic. The assumption of no inter-reflection is also typical, but unfortunately not realistic at all. Finally, the assumption of no ambient light is not at all realistic, but will be relaxed (in fact, eliminated) in section 4.1 below.

Equally interesting is a list of assumptions *not* made by the model, which express the scope or generality of the model:

- *Imaging geometry* -- the model makes no assumption that orthography or perspective is being used. Either projection satisfies the model.
- *Planar surface* -- the model applies equally to curved and planar surfaces. It also applies to textured surfaces, i.e. surfaces with macroscopic roughness (but see the note below about analyzing intrinsic images).
- *Specific reflectance model* -- the model does not assume specific functions m_i , c_i , m_b , or c_b ; in particular, there is no specific geometric model of highlights, no assumption that the highlights have the same color as the illumination, and no assumption that the body reflection is isotropic.
- *Point light source* -- the model applies equally well to a point light source, an extended light source, or a light source infinitely far away.
- *Uniform distribution of illumination* -- the model does not assume that the amount of illumination is the same everywhere in the scene; only that the SPD is the same (see above). This is important, since real (especially extended) light sources produce nonuniform amounts of illumination in different areas of the scene.

It must be pointed out that any complexity of the forms mentioned above will cause great difficulty in *analyzing* the resulting intrinsic images of m_i and m_b . However, the Dichromatic Reflection Model will make it possible to *compute* these intrinsic images regardless of such complexity.

In spite of the apparent simplicity of the Dichromatic Reflectance Model, it suffers from some flaws. In practice, they ought to have only a minimal impact on the usefulness of the model.

- There is no obvious way to decide how to scale the magnitude functions m_i and m_b against the composition functions c_i and c_b . Therefore, the resulting intrinsic images may be only of relative, rather than absolute, reflection magnitudes. At least, it is reasonable to require that $0 \leq m_i, m_b \leq 1$ for all i, e , and g .
- Interface reflection exhibits an interdependence between wavelength and geometry, as expressed by Fresnel's equations. As stated above, this is a small effect; the author's

estimate is an effect no greater than 2% in pixel value errors if the Dichromatic Model is assumed. An error so large would only occur when $i \geq 60^\circ$ or so and the viewing direction is nearly the perfect specular direction; in most situations, therefore, the effect should be negligible.

- Body reflection also exhibits an interdependence between wavelength and geometry. If $c_i(\lambda)$ is not constant, then the color of the light passing through the interface will differ somewhat from the color of the illumination $S(\lambda)$. Since the total amount of light reflected at the interface ($\int_e \int_g L_i[\lambda, i, e, g] dg de$) varies with the angle of incidence i , the color of the light passing through the interface into the material body will also vary with the angle of incidence. Thus, the color of the body reflection should vary with geometry. However, since c_i is generally nearly constant, this should also be a negligible effect in real images.

In summary, the Dichromatic Reflection Model is expressed in a simple equation which should be approximately true over a very wide range of circumstances.

2.3 Previously Used Reflection Models

Previous models of reflection used in computer graphics and computer vision have proposed specific reflectance functions to predict the exact amount of reflection.

The most frequently cited model is that of Phong [25, 8], also applied in computer vision by Horn [11]:

$$L(i, e, g) = t \frac{n+1}{2} \cos^n s + (1-t) \cos i$$

Here, t and n are parameters of the material, with t representing the total amount of light reflected at the interface and n representing the clustering of the interface reflection about the perfect specular direction. The equation shown above is only a monochrome model, but is frequently applied for color images with the color of the interface reflection being white and the color (R-G-B values) of the body reflection being a parameter of the material. Phong's model is easily seen to be an instantiation of the Dichromatic Reflection Model in this paper as follows:

$$m_i = t \frac{n+1}{2} \cos^n s$$

$$m_b = (1-t) \cos i$$

$$c_i = 1$$

$$c_b = \text{surface parameter}$$

A later model was Blinn's [4], in which the interface reflection was characterized by a facet model of surface structure (from [33]) and Fresnel's equations were used for the color of the interface reflection. Cook and Torrance, in more recent work [6], used Beckmann's model of reflection from a rough surface [2] for the distribution interface reflection, and similarly applied Fresnel's equations for the color of the interface reflection. In both cases, the body reflection was assumed to be isotropic,

i.e. $m_b = \cos i$. These models of reflection are also examples of the Dichromatic Reflection Model, with the exception of the use of Fresnel's equations to govern the color of the interface reflection (usually a small effect, as noted above).

In computer vision, Shibata and Frei [29] used the following model to predict reflection in aerial images:

$$L(i, e, g) = t_1 \cos^n s + t_2 \cos^m g + (1 - t_1 - t_2) \cos i$$

Here, the first and last terms are similar to Phong's, and the middle term represents "backscattering" of light in the direction of illumination. The constants t_1 , t_2 , n , and m are parameters of the material being viewed. Extended to color images, this model would also be an instantiation of the Dichromatic Reflection Model, with the backscattering term added to m_1 .

3. Pixel Values in Color Space

This chapter begins with a discussion of spectral projection, the relationship between an SPD of light and its color coordinates. When this process is applied to the Dichromatic Reflection Model, a new model for pixel values in color space is the result. A simple algorithm is presented for exploiting this model to determine the intrinsic images of m_i and m_b .

3.1 Spectral Projection

Spectral projection is the process whereby pixel values are computed from the spectral power distribution (SPD) of the measured light. The process in a monochrome camera is quite simple, with the pixel value p being just a summation of the amount of light at each wavelength $X(\lambda)$, weighted by the responsivity of the camera to the various wavelengths, $s(\lambda)$:

$$p = \int X(\lambda) s(\lambda) d\lambda$$

The interval of summation is determined by $s(\lambda)$, which is non-zero over a bounded interval of wavelengths λ .

In a color camera, color filters are interposed between the incoming illumination and the camera. Each filter has a transmittance function $\tau(\lambda)$, specifying the fraction of light transmitted at each wavelength; thus, spectral projection with a filter is specified by the above integral, with $s(\lambda)$ replaced by $\tau(\lambda) s(\lambda)$. Typically, three filters (red, green, and blue) are used with transmittances τ_r , τ_g , and τ_b , resulting in a vector of three color values, $\mathbf{C} = [r, g, b]$. If we let $\bar{r}(\lambda)$ be the responsivity of the camera combined with the red filter, $\bar{r}(\lambda) = \tau_r(\lambda) s(\lambda)$, (etc. for \bar{g} and \bar{b}), then the color value of any SPD $X(\lambda)$ is given by:

$$\mathbf{C}_X = \begin{bmatrix} r_X \\ g_X \\ b_X \end{bmatrix} = \begin{bmatrix} \int X(\lambda) \bar{r}(\lambda) d\lambda \\ \int X(\lambda) \bar{g}(\lambda) d\lambda \\ \int X(\lambda) \bar{b}(\lambda) d\lambda \end{bmatrix}$$

Spectral projection is a linear transformation, as shown in [28]; in other words, $\mathbf{C}_{aX+bY} = a\mathbf{C}_X + b\mathbf{C}_Y$ where a and b are scalars and $X(\lambda)$ and $Y(\lambda)$ are SPD's. To see this, consider first the red component r_{aX+bY} ; it is easily seen that:

$$\begin{aligned} r_{aX+bY} &= \int [aX(\lambda) + bY(\lambda)] \bar{r}(\lambda) d\lambda = a \int X(\lambda) \bar{r}(\lambda) d\lambda + b \int Y(\lambda) \bar{r}(\lambda) d\lambda \\ &= a r_X + b r_Y \end{aligned}$$

With similar equations for green and blue, we have the complete result. The linearity property of

spectral projection is important because it says that a mixture of two SPD's of light results in a sum of the corresponding pixel values, taken in the same proportion.

3.2 The Dichromatic Model in Color Space

When the linear property of spectral projection is combined with the Dichromatic Reflection Model, a powerful new model of pixel color values results.

First, consider a specific point on a surface. At that point, the geometry (angles i , e , and g) is determined, and the magnitudes m_i and m_b may be considered as scalars. So, the Dichromatic Reflectance Model may be rewritten at a specific point as:

$$L(\lambda) = m_i c_i(\lambda) + m_b c_b(\lambda)$$

This defines the SPD of the light reflected from the surface. Now, applying the linearity of spectral projection, we have:

$$\mathbf{C}_L = m_i \mathbf{C}_i + m_b \mathbf{C}_b$$

where \mathbf{C}_L is the color (pixel value) measured, m_i and m_b are the magnitudes of reflection at the point in question, and \mathbf{C}_i and \mathbf{C}_b are the colors of the interface and body reflection of the material.

Consider the colors \mathbf{C}_L of a set of points on the same uniformly colored surface. Because the geometry is different at each point, the scale factors m_i and m_b vary from point to point. However, the colors \mathbf{C}_i and \mathbf{C}_b of the interface and body reflection are the same at all points on the same surface, because they are simply the spectral projections of $c_i(\lambda)$ and $c_b(\lambda)$ which do not vary with geometry. In other words, the pixel values are a linear combination of \mathbf{C}_i and \mathbf{C}_b , with the coefficients determined by m_i and m_b at each point.

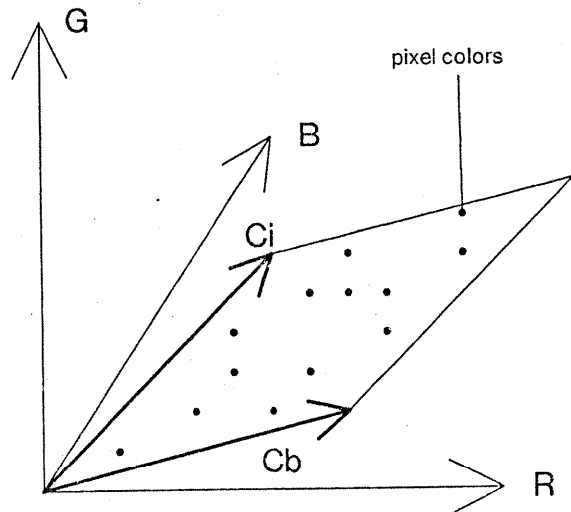


Figure 3-1: Pixel Values on a Surface Lie on a Parallelogram in Color Space




Figure 3-2: Position Within Parallelogram Determines Magnitudes

Recalling that we can assume $0 \leq m_i, m_b \leq 1$ without loss of generality, we see that the pixel values C_L for a set of points on a single surface must lie within a parallelogram in color space, bounded by the colors C_i and C_b of the interface and body reflection of the surface (Figure 3-1). One corner of this parallelogram will be located at the origin, $[r,g,b] = [0,0,0]$. Further, within this parallelogram, the position of any color is determined by the values of m_i and m_b at the corresponding point (Figure 3-2).

The Dichromatic Model in Color Space makes the following assumptions:

- Validity of the Dichromatic Reflection Model, i.e. all the assumptions made therein.
- Prior segmentation of the image into groups of pixels known (or believed) to lie on a single surface.
- Pixel values returned by the camera are linearly related to the irradiance on the sensor plane, i.e. the camera is photometrically calibrated. If the camera is calibrated for monochrome response, it is not necessary to re-calibrate it with each color filter separately since the relative SPD of a filter's transmittance is constant with respect to total intensity of illumination.

The model of pixel values does not assume:

- *Specific C_i* -- it is not assumed that C_i is achromatic (i.e. $r_i = g_i = b_i$), nor that it is the same for all surfaces in the image.
- *Specific color responses* -- no assumption is made about the color filters $\tau_r(\lambda)$, $\tau_g(\lambda)$, and $\tau_b(\lambda)$, except that they are linearly independent, nor is any assumption made about the camera's spectral responsivity $s(\lambda)$, except that the relative spectral responsivity is constant with respect to the total amount of irradiance and constant across the sensor plane.

3.3 Analyzing Color Pixel Values

We can now present a simple algorithm for computing the intrinsic images of m_i and m_b over a set of pixels corresponding to a single surface in the image:

1. Histogram the pixel values in color space.
2. Fit a plane in color space to these points, with the restriction that the plane must pass through the origin.
3. Fit a parallelogram on this plane to the points, with the restriction that one vertex must lie at the origin. The sides of this parallelogram are C_i and C_b . If it is necessary to distinguish one from the other, either geometric criteria may be used or the fact that C_i probably lies close to the achromatic axis of the color space.
4. At each pixel, express its color as a linear combination of C_i and C_b . This can be performed with a simple linear transformation, i.e. a matrix multiplication of the $[r, g, b]$ color values by a 2×3 matrix (not 3×3 since the direction orthogonal to the plane in color space is of no interest except as an error term).
5. The resulting values are the magnitudes m_i and m_b .

3.3.1 Assumptions of the Dichromatic Analysis Algorithm

In addition to the assumptions of the Dichromatic Model in Color Space, this algorithm depends upon a number of additional assumptions for its effectiveness:

- C_i and C_b must be linearly independent, i.e. the interface and body reflection must have different colors.
- The noise in measuring pixel values must be small enough that plane-fitting and parallelogram-fitting can proceed.
- The distribution of pixel values within the parallelogram must not be pathological. In general, the author expects that many points will lie close to the C_b axis, i.e. have m_i close to zero, and the remaining points will lie close to the C_i axis. These latter points may have pixel values "clipped" by the extent of the color cube itself, since highlights may be very bright and thus cause truncation of the measured pixel value at the maximum response of the camera.

It should also be noted that the Dichromatic Analysis Algorithm uses pixel colors to determine m_i and m_b , but it does not interpret m_i and m_b in terms of the photometric angles i , e , and g . Thus, this algorithm does not tell how to analyze the values of m_i and m_b .

3.4 Summary

In this chapter, we have seen a mathematical description of spectral projection -- the process by which an SPD of light is converted into color values. Spectral projection is a linear transformation, a fact that can be applied to the Dichromatic Reflection Model to produce a model of pixel values in color space. The distribution of pixel values is expected to form a parallelogram, bounded by the colors C_i and C_b of the interface and body reflection from the surface. The position of any color within this parallelogram can be used to determine the coefficients m_i and m_b at the corresponding pixel in the image.

4. Applying the Dichromatic Analysis Algorithm

In this chapter we discuss the extension of the Dichromatic Reflection Model for dealing with diffuse illumination, and some possible methods of analyzing the intrinsic images of m_i and m_b .

4.1 Extending the Model for Diffuse Illumination

The Dichromatic Reflection Model as presented above depends on the assumption that the illumination at any point comes from a single (point or extended) light source. It is more realistic to model the illumination as consisting of a light source, plus "ambient" or "diffuse" light, of lower intensity, coming from all directions in equal amounts, and possibly with a different SPD than the small source. This model is a better approximation of daylight, which contains light from a point source (the yellowish sun) and light from a hemisphere (the bluish sky), and of light in a room, which comes from light fixtures and from inter-reflections off walls and other objects. All the previously mentioned models of reflection in computer graphics presume some sort of diffuse illumination, although the assumption made here (diffuse illumination from all directions equally) is admittedly highly idealized.




Figure 4-1: Color Space Parallelogram With Diffuse Illumination

The light reflected by diffuse illumination contains a part due to interface reflection and a part due to body reflection. By assuming that this light is incident from and reflected into all directions equally, it can be modelled by adding a single term, $L_a(\lambda)$, to the Dichromatic Reflection Model:

$$L(\lambda, i, e, g) = m_i(i, e, g) c_i(\lambda) + m_b(i, e, g) c_b(\lambda) + L_a(\lambda)$$

In color space, the model becomes:

$$\mathbf{C}_L = m_i \mathbf{C}_i + m_b \mathbf{C}_b + \mathbf{C}_a$$

where \mathbf{C}_a is the color of the light reflected from diffuse illumination $L_a(\lambda)$. Since $L_a(\lambda)$ does not vary with geometry, the effect of this change is to translate the parallelogram of pixel colors for a single surface away from the origin by the vector \mathbf{C}_a , as seen in figure 4-1.

The Dichromatic Analysis Algorithm can still be applied with this change, but note that the plane-fitting and parallelogram-fitting operations specified in the algorithm will be less constrained and therefore more error-prone. If the relative SPD of the diffuse illumination is the same as that of the small source, then \mathbf{C}_a is a linear combination of \mathbf{C}_i and \mathbf{C}_b [6]; in such a case the *plane* containing the points should pass through the color space origin although the *parallelogram* will not.

It is interesting to note that a point whose color lies at exactly $\mathbf{C}_L = \mathbf{C}_a$ has $m_i = m_b = 0$, and might therefore be suspected of lying within a shadow. This is a far more precise description of colors within shadows than statements of the form "shadows tend to be bluish", which have been seen in previous work in image understanding. In fact, since for different surfaces \mathbf{C}_a is very likely to differ, it may be possible to associate shaded parts of surfaces with illuminated portions, by constructing \mathbf{C}_a for each adjacent illuminated area and finding which one matches best with the color of the pixels in the shaded area. Unfortunately, since pixel values in shaded areas of an image tend to be poorly measured by current digitizing cameras, this kind of analysis may prove to be unreliable.

The description presented here is intended to model diffuse illumination; however, it will be a poor substitute for a detailed model of inter-reflection when surfaces are close to each other.

The author continues to call the above equation a "Dichromatic Reflection Model" (rather than a trichromatic model) because the essence of the model is that reflection occurs from two places in a surface: the interface and the surface body.

4.2 Analysis of the Intrinsic Reflection Images

The potential utility of the intrinsic images of m_i and m_b stems from two facts:

- Both m_i and m_b have simpler geometric properties than does intensity in a monochrome image, which represents a weighted sum of the two.
- m_b , representing diffusely distributed light, is relatively insensitive to changes in viewpoint from one image to the next.

Here are some possible methods for exploiting these properties of the intrinsic images of m_i and m_b :

1. In stereo or motion analysis, image matching using m_b might be more reliable because of

the elimination of interface reflection, whose position in the image is highly viewpoint-dependent. *

2. If a specific reflectance model such as Phong's is assumed, then the two values m_i and m_b provide two constraints on the surface normal at any point. For example, using Phong's model, the angle s can be calculated from m_i and the angle i can be calculated from m_b . With such constraints, unique surface normals can be established from a color image. (In practice, this kind of analysis would be limited by the applicability of such reflectance models and by the distribution of points on the parallelogram in color space.)
3. If, as above, a specific reflectance model is assumed but the parameters (t, n in Phong's model) are not known, it may be possible to determine or estimate these from the image. At a point with a known orientation, the coefficients m_i and m_b can be determined by color analysis; using these, the surface parameters may be calculated. This might form the basis for a technique of looking at parts of a surface (say, the spots with pronounced highlights) to estimate the nature of the reflectance properties of the surface, then using this estimated reflectance model to analyze the rest of the surface in detail.
4. Using differentials, since the photometric angles i, e , and g are functions of x and y in the image, we can relate the differentials of the intrinsic image of m_i to the geometric properties of m_i and the imaging geometry by:

$$[\partial m_i / \partial x \quad \partial m_i / \partial y] = [\partial m_i / \partial i \quad \partial m_i / \partial e \quad \partial m_i / \partial g] \begin{bmatrix} \partial i / \partial x & \partial i / \partial y \\ \partial e / \partial x & \partial e / \partial y \\ \partial g / \partial x & \partial g / \partial y \end{bmatrix}$$

This equation tells how the intrinsic image and a reflectance model yield constraints on the imaging geometry using first derivatives.

5. Conclusions

In this paper, we have seen the development of a fairly simple algorithm for determining intrinsic images of reflection from a color image. The presentation included a fairly general reflection model based on a physical description of the reflection process, a model of pixel value distribution based on the properties of spectral projection, and an algorithm for analyzing color values in an image. This paper also presented a discussion of some related issues such as diffuse illumination, image segmentation, and analysis of the resulting intrinsic images. These intrinsic images promise to be useful for a variety of image understanding tasks, including stereo matching, geometric surface analysis, and analysis of surface reflectance parameters.

This work represents a quantum advance over previous efforts in color image understanding, because it presents a quantitative theory based on the physics of the real world. In this respect, it follows an increasingly important paradigm of three-dimensional image understanding research:

1. Describe a physical (geometric or optical) phenomenon in the scene.
2. Develop a mathematical statement of this scene property.
3. Apply an imaging model to this statement, to produce a new mathematical statement relating image properties to scene properties.
4. Develop an algorithm to analyze the scene properties using the above relationship.

As is frequently the case, the complex intermediate mathematics are used in this paper only to derive the algorithm; they are not necessarily directly incorporated in any actual program.

This paper also displays one of the dilemmas of current work in 3D computer vision: we frequently see simple theories and algorithms developed to analyze isolated phenomena such as shadows and highlights, but we rarely see them applied to real images. While some readers may feel that this is because the theories are too complex, the opposite is a more tenable point of view: that such theories fail in practice because they are too simple, i.e. their assumptions are too restrictive. In order to apply them to real images, it is first necessary to relax as many assumptions as possible, as was done here in discussing diffuse illumination and image segmentation. This produces the striking paradox of complex theories whose validity should be tested by applying them to real images; but they cannot be applied to real images until they are made yet more complex!

We have yet to see how researchers in three-dimensional image understanding will resolve this problem, but the utility of such research depends upon its solution. The current work may help by demonstrating a theory which depends upon relatively few special assumptions, and by providing

"simpler" intrinsic images useful for estimating surface and illumination parameters from examination of the images themselves.

Further efforts in developing this work will include implementation (when a calibrated camera is obtained), pursuing the ideas for intrinsic image analysis suggested above, and developing the theory that will ultimately allow analysis of nonuniform illumination, extended light sources, and perhaps inter-reflection among surfaces. In addition, a segmentation scheme based on the Dichromatic Reflection Model might be investigated.

5.1 Acknowledgements

The author would like to thank Fred Billmeyer, of Rensselaer Polytechnic Institute, and David MacAdam, of the Optics Institute at the University of Rochester, for their comments and suggestions. Takeo Kanade at Carnegie-Mellon University provided much guidance for this work, and Marty Herman and Chris Brown provided comments on the paper. Frank O'Donnell at Rensselaer Polytechnic Institute provided a wealth of papers on gloss measurement from which figures 2-2 and 2-3 were drawn. Those figures are reproduced by permission from their respective publishers.

6. Bibliography

- [1] Barrow, H. G. and Tenenbaum, J. M.
Recovering Intrinsic Scene Characteristics from Images.
In Hanson, A. R. and Riseman, E. M. (editor), *Computer Vision Systems*, pages 3-26.
Academic Press, New York, 1978.
- [2] Beckmann, P. and Spizzichino, A.
The Scattering of Electromagnetic Waves from Rough Surfaces.
MacMillan, 1963.
Pages 1-33, 70-98.
- [3] Billmeyer, Fred W. Jr.
Private communication.
- [4] Blinn, J. F.
Models of Light Reflection for Computer Synthesized Pictures.
Computer Graphics 11(2):192-198, 1977.
SIGGRAPH '77.
- [5] Coleman, G. B. and Andrews, H. C.
Image Segmentation by Clustering.
Proc. IEEE 67(5):773-785, May, 1979.
- [6] Cook, R. L. and Torrance, K. E.
A Reflectance Model for Computer Graphics.
Computer Graphics 15(3):307-316, August, 1981.
SIGGRAPH '81.
- [7] Egan, W. G. and Hilgeman, T. W.
Optical Properties of Inhomogeneous Material.
Academic Press, New York, 1979.
- [8] Foley, J. D. and Van Dam, A.
Fundamentals of Interactive Computer Graphics.
Addison-Wesley, Reading, Mass., 1983.
- [9] Grum, F. and Becherer, R. J.
Optical Radiation Measurements, Volume 1: Radiometry.
Academic Press, New York, 1979.
- [10] Grum, F. and Bartleson, C. J., editors.
Optical Radiation Measurements, Volume 2: Colorimetry.
Academic Press, New York, 1980.
Chapter 7, "Colorant Formulation and Shading", by E. Allen.
- [11] Horn, B. K. P.
Understanding Image Intensities.
Artificial Intelligence 8:201-231, 1977.
- [12] Horn, B. K. P. and Sjoberg, R. W.
Calculating the Reflectance Map.
Applied Optics 18:1770-1779, 1979.

- [13] Hunter, R. S.
High Gloss Measurements.
Official Digest of the Federation of Societies for Paint Technology 36:348-356, April, 1964.
- [14] Hunter, R. S.
The Measurement of Appearance.
J. Wiley and Sons, New York, 1975.
- [15] Jenkins, F. A. and White, H. E.
Fundamentals of Optics.
McGraw-Hill, New York, 1976.
- [16] Judd, D. B. and Wyszecki, G.
Color in Business, Science and Industry.
J. Wiley and Sons, New York, 1975.
- [17] Kender, J. R.
Instabilities in Color Transformations.
In *PRIP-77*, pages 266-274. IEEE Computer Society, Troy NY, June, 1977.
- [18] Levine, M. D. and Shaheen, S. I.
A Modular Computer Vision System for Picture Segmentation and Interpretation, Part I.
In *PRIP-79*, pages 523-533. IEEE Computer Society, Chicago, Ill., August, 1979.
- [19] Nagao, M., Matsuyama, T., and Ikeda, Y.
Region Extraction and Shape Analysis in Aerial Photographs.
Computer Graphics and Image Processing 10:195-223, 1979.
- [20] Nagin, P. A., Hanson, A. R., and Riseman, E. M.
Region Extraction and Description Through Planning.
COINS TR 77-8, U. Mass., May, 1977.
- [21] Nevatia, R.
A Color Edge Detector and Its Use in Scene Segmentation.
IEEE Trans. Systems, Man, and Cybernetics TSMC-7(11):820-826, November, 1977.
- [22] Nicodemus, F. E., Richmond, J. C., Hsia, J. J., Ginsberg, I. W., and Limperis, T.
Geometrical Considerations and Nomenclature for Reflectance.
NBS Monograph 160, National Bureau of Standards, October, 1977.
- [23] Ohlander, R.
Analysis of Natural Scenes.
PhD thesis, Carnegie-Mellon Univ. Computer Science Dept., 1975.
- [24] Ohta, Y., Kanade, T., and Sakai, T.
Color Information for Region Segmentation.
Computer Graphics and Image Processing 13:222-241, 1980.
- [25] Phong, Bui Tuong.
Illumination for Computer Generated Pictures.
Communications of the ACM 18:311-317, 1975.
- [26] Rubin, S.
The ARGOS Image Understanding System.
PhD thesis, Carnegie-Mellon Univ. Computer Science Dept., 1978.

- [27] Rubin, J. M. and Richards, W. A.
Color Vision and Image Intensities: When are Changes Material?
AI Memo 631, MIT, May, 1981.
- [28] Shafer, S. A.
Describing Light Mixtures Through Linear Algebra.
J. Optical Soc. Am. 72(2):299-300, February, 1982.
- [29] Shibata, T., Frei, W., and Sutton, M.
Digital Correction of Solar Illumination and Viewing Angle Artifacts in Remotely Sensed Images.
In *Proc. 7th Symposium on Machine Processing of Remotely Sensed Data*, pages 169-177.
1981.
- [30] Simpson, L. A.
Measuring Gloss and Factors that Affect It.
Technical Report, BTP Tioxide Limited, Stockton-on-Tees, England, 1981.
- [31] Sloan, K.
World Model Driven Recognition of Natural Scenes.
PhD thesis, U. Penna. Moore School of Electrical Engineering, June, 1977.
- [32] Tenenbaum, J. M. and Weyl, S.
A Region-Analysis Subsystem for Interactive Scene Analysis.
In *Proc. 4th IJCAI*, pages 682-687, September, 1975.
- [33] Torrance, K. E. and Sparrow, E. M.
Theory for Off-Specular Reflection from Roughened Surfaces.
J. Optical Soc. Am. 57:1105-1114, September, 1967.
- [34] Weymouth, T. E., Griffith, J. S., Hanson, A. R., and Riseman, E. M.
Rule Based Strategies for Image Interpretation.
In *Proc. NCAI-83*, pages 429-532. AAAI, August, 1983.
- [35] Williamson, S. J. and Cummins, H. Z.
Light and Color in Nature and Art.
J. Wiley and Sons, New York, 1983.
- [36] Young, I. T.
The Classification of White Blood Cells.
IEEE Trans. Biomedical Engineering BME-19(4):291-298, July, 1972.