

The Measurement of Highlights in Color Images

Gudrun J. Klinker, Steven A. Shafer, and Takeo Kanade

Abstract

In this paper, we present an approach to color image understanding that accounts for color variations due to highlights and shading. We demonstrate that the reflected light from every point on a dielectric object, such as plastic, can be described as a linear combination of the object color and the highlight color. The colors of all light rays reflected from one object then form a planar cluster in the color space. The shape of this cluster is determined by the object and highlight colors and by the object shape and illumination geometry. We present a method which exploits the difference between object color and highlight color to separate the color of every pixel into a matte component and a highlight component. This generates two intrinsic images, one showing the scene without highlights, and the other one showing only the highlights. The intrinsic images may be a useful tool for a variety of algorithms in computer vision, such as stereo vision, motion analysis, shape from shading, and shape from highlights. Our method combines the analysis of matte and highlight reflection with a sensor model which accounts for camera limitations. This enables us to successfully run our algorithm on real images taken in a laboratory setting. We show and discuss the results.

1. Introduction

When we look at an image, we can interpret it as a collection of shiny and matte surfaces, smooth and rough, interacting with light, shape, and shadow. However, computer vision has not yet been successful at deriving a similar description of surface and illumination properties from an image. The key reason for this failure has been a lack of models or descriptions rich enough to relate pixels and pixel-aggregates to scene characteristics. In the past, most work with color images has considered object color to be a constant property of an object, and color variation on an object was attributed to noise [28]. However, color variation in real scenes depends to a much larger degree on the optical reflection properties of the scene. This variation causes the perception of object color, highlights, shadows and shading [14, 31], scene properties that can be determined and used by color vision algorithms.

This paper presents an approach to color image understanding that accounts for color variations due to highlights and shading. We use a reflection model which describes pixel colors as a linear combination of an object color and a highlight color [34]. All color pixels from one object then form a planar cluster in the color space. The cluster shape is determined by the object and highlight colors and by the object shape and illumination geometry. We combine our reflection model with a sensor model which accounts for camera properties, such as a limited dynamic range, blooming and gamma-correction.

By combining these explicit models, we separate color images into two intrinsic images, one showing the scene without highlights, and the other one showing only the highlights. The intrinsic images can be a useful tool for a variety of algorithms in computer vision that do not account for highlights in images, such as stereo vision and motion analysis [7, 35]. Our approach can also be used to determine the color of the illumination from highlights in the image and as a preprocessor for methods to determine object shape from shading or highlights [12, 14].

We begin by describing the major influences on the measurement of a color pixel value: the physics of light reflection in the scene, the object shape, and the sensor characteristics of the camera. We then demonstrate how our reflection model can be used for the analysis of color images. We present a method to detect and remove highlights from color images and to determine the illumination color. Finally, we discuss the implications and the future direction of our work.

2. What's in a Color Pixel?

On its path from a light source to the camera, a light ray is altered in many characteristic ways by the objects in the scene. The camera then encodes the measured light in a (color) pixel. It is the goal of image understanding methods to use properties and relationships between pixels to recover a description of the scene. It is essential to the success of such methods that they understand and model the reflection processes in the scene, as well as the sensing characteristics of the camera. We will now describe the processes that occur when light interacts with inhomogeneous, dielectric materials. We will also discuss how the measured pixel values depend on camera characteristics.

2.1. The Physics of Light Reflection from Dielectric Materials

When light interacts with dielectric materials, two kinds of reflection exist: The first takes place when light encounters a change in the refraction index at the interface between two substances; the second occurs within a substance when light is scattered and selectively absorbed [16, 19, 37]. These two kinds of reflection differ in their geometric and photometric properties, and as we will show, those differences can be exploited for computer vision.

2.1.1. Properties of Inhomogeneous Dielectric Materials

Dielectric materials are nonconducting materials. Most of them, such as paints, plastics, paper and ceramics, are inhomogeneous. They consist of a medium and some embedded pigments, as shown in Figure 1. The medium comprises the bulk of the matter and is generally approximately transparent, while the pigments selectively absorb light and scatter it by reflection and refraction.

Figure 1: Light reflection of dielectric materials

When we look at objects that are made out of dielectric material, we usually see the reflected light as composed of two distinct colors that typify the highlight areas and the matte object parts. This is a characteristic property of many dielectrics, and our reflection model capitalizes on this characteristic color change between matte and highlight areas. Our model is applicable to a large set of typical dielectric materials, as we will demonstrate in section 3. It is not applicable to non-dielectric (i.e.: conducting) materials because those reflect light in only a single color. It is also not applicable to dielectrics that imitate the reflection behavior of metals, e.g.: by containing metallic flakes in their material body. We now describe the typical characteristics of dielectrics in more detail.

2.1.2. Reflection at the Material Surface

When light hits the surface of a dielectric material, it must first pass through the interface between the surrounding medium (e.g.: air) and the material. Since the refraction index of the material is generally different from that of the surrounding medium, some percentage of the incident light is reflected at the surface of the material. We will refer to this light reflection process as *surface reflection*. This process has been referred to in the literature as interface reflection, or Fresnel reflection [13, 20, 21, 34].

If the interface between the surrounding medium and the material is completely smooth, the light is reflected from this interface in only one direction, such that the incident light beam, the exiting light beam and the surface normal are coplanar and the surface normal bisects the angle between incident and exiting light. This is the direction of perfect mirror reflection. However, the interfaces of many materials have some roughness or texture. The direction of the reflected light depends then on the orientation of the local surface normals, which varies along the object. In such cases, the light is scattered to some degree around the global angle of perfect mirror reflection. Several models have been developed in the physics and computer graphics communities to describe the geometric properties of light reflection from rough surfaces [2, 30, 33, 36]. Issues in modeling this process involve the roughness scale of the surface, compared to the wavelengths of the incident light, and self-shadowing effects on the surface which depend on the viewing direction and the direction of the incident light.

The optical properties of the material also determine the amount and color of the light that is reflected at the material surface. Fresnel's law describes how the reflected light depends on the refraction indices of the material and the surrounding medium, on the incidence angle and on the polarization of the light [25]. In general, the refraction index of a material is a function of wavelength, and the Fresnel coefficient varies over the light spectrum. Furthermore, pigment particles may be protruding from the medium at some places [17]. The reflected light may thus be also a function of the position on the object. However, it is common to assume that the pigments are completely embedded in the medium and to approximate the Fresnel coefficient of the medium as a constant over the visible spectrum. This is reasonable since the refraction coefficient of most media changes very little over the visible spectrum. Under these assumptions, the light reflected at the surface has the same color as the illuminating light.

2.1.3. Reflection in the Material Body

For dielectric materials, not all incident light is reflected at the material surface. Some percentage of the light penetrates into the material body. The refracted light beam travels through the medium, hitting pigments from time to time. The material medium is generally approximately transparent, transmitting the refracted light beam without a color or orientation change. When the light beam hits a pigment particle, some of the light is scattered, due to the change in refraction indices. The rest of the beam penetrates through the particle, where some wavelengths may be partially or entirely absorbed. The remaining light returns afterwards into the surrounding medium [16, 19]. Some part of the beam may travel then to the material surface while the rest stays within the body and hits other pigments. The part of the light beam that arrives at the material surface is there partially refracted and partially reflected, as described above for the surface reflection process. Thus, some fraction of the light exits from the material while the rest is reflected back into the body and undergoes more scattering and absorption.

We refer to the entire reflection process described above as *body reflection*. Its geometric and photometric properties depend on many factors: the transmitting properties of the medium, the scattering and absorption properties of the pigments, and the shape and distribution (including density and orientation) of the pigments. If we assume a random distribution of the pigments, the light exits in random directions from the body. In the extreme, when the exiting light is uniformly distributed, the distribution can be described by Lambert's law. The distribution of the pigments also influences the amount and the color of the reflected light. If the pigments are distributed randomly in the material body, we may expect that, on the average, the same amount and color will be absorbed everywhere in the material before the light exits. In such a case, the light that is reflected from the material body has the same color over the entire surface.

The distribution, size and scattering properties of the pigments also influence the opacity of the material body. A material is opaque if no part of the refracted light beam is able to penetrate through the entire material. Otherwise, the material is transparent or translucent.

In previous literature of computer vision, the terms "specular" and "diffuse" reflection have frequently been used to refer to surface and body reflection, respectively. However, this is not strictly correct. "Specular reflection" refers to light reflection in the direction of ideal mirror-like reflection, while "diffuse reflection" implies reflection with scattering. This is a purely geometric distinction, while the terms "surface" and "body" reflection refer to different physical processes with distinct spectra. For a very

smooth surface, the surface reflection will be specular and the body reflection will be diffuse. However, for typical objects with rough surfaces, the surface reflection will be diffused around the direction of perfect specular reflection, as modeled for example in [36]. Observing this distinction makes it possible to describe many types of material succinctly: for example, metals have only surface reflection, which may be specular or diffuse depending on the roughness.

2.2. The Dichromatic Reflection Model

We now present a mathematical model that approximates the reflection processes of dielectric, inhomogeneous, opaque materials [34]. In order for this model to be a suitable basis for work in color image understanding, it makes several restricting and simplifying assumptions about the properties of surface and body reflection. We begin by assuming that the pigments are distributed randomly in the material body and that they are completely embedded in the medium. The surface then exhibits a single spectrum of body reflection and a single spectrum of surface reflection. We also restrict the illumination conditions of the scene, allowing only one light source and no ambient light or inter-reflection between objects. These assumptions, in concept, restrict the applicability of the model. However, we will show in section 3 that the model still provides a reasonable and very useful approximation to the physics of light reflection for many dielectric materials.

We begin by describing the light, L , which is reflected from an object point as a mixture of the light L_s reflected at the material surface and the light L_b reflected from the material body.

$$L(\lambda, i, e, g) = L_s(\lambda, i, e, g) + L_b(\lambda, i, e, g) \quad (1)$$

The reflectance geometry is illustrated in Figure 2. It defines several photometric angles [14, 33]: the *incidence angle*, i , between the illumination direction, \mathbf{I} , and the surface normal, \mathbf{N} , the *exitance angle*, e , between \mathbf{N} and the viewing direction, \mathbf{V} , and the *phase angle*, g , between \mathbf{I} and \mathbf{V} . The parameter λ represents wavelengths of the light spectrum.

Figure 2: Photometric angles

Since we assume that a dielectric, inhomogeneous surface exhibits the same spectrum of surface reflection at any point on the object, we can separate the spectral reflection properties of L_s from its geometric reflection properties [26]. We thus model it as a product of a spectral power distribution, $c_s(\lambda)$, and a geometric scale factor, $m_s(i, e, g)$, which describes the intensity of the reflected light. Similarly, we separate the body reflection component L_b into a spectral power distribution, $c_b(\lambda)$, and a geometric scale

factor, $m_b(i,e,g)$. Substituting these terms into equation (1), we obtain the Dichromatic Reflection Model equation:

$$L(\lambda,i,e,g) = m_s(i,e,g) c_s(\lambda) + m_b(i,e,g) c_b(\lambda) \quad (2)$$

We thus describe the light that is reflected from an object point as a mixture of two distinct spectral power distributions, $c_s(\lambda)$ and $c_b(\lambda)$, each of which is scaled according to the geometric reflection properties of surface and body reflection. This is shown in Figure 3. In the infinite-dimensional vector space of spectral power distributions (each wavelength defines an independent dimension in this vector space [15, 32]), the reflected light can thus be described as a linear combination of the two vectors $c_s(\lambda)$ and $c_b(\lambda)$.

Figure 3: A mixture of the lights reflected from the material surface and body

Many reflection models, which have been developed in the physics and computer graphics communities [5, 14, 30, 36] are special cases of the model described here [34]. They replace the geometric variables, m_s and m_b , by specific functions that approximate the measured reflection data of a chosen set of typical materials. In our work, we concentrate on the spectral variables in equation 2, $c_s(\lambda)$ and $c_b(\lambda)$, exploiting the color difference between them. We leave the geometrical factors unspecified.

2.3. Object Shape and Color Variation

We will now discuss the relationship between the light mixtures of all points on an object. We study the spectral variation over an entire object by analyzing the histogram of the light mixtures from all object points.

The Dichromatic Reflection Model states that the light spectra of body reflection, $c_b(\lambda)$, and surface reflection, $c_s(\lambda)$, are constant over an entire object, while the geometric scale factors, $m_b(i,e,g)$ and $m_s(i,e,g)$, change with the viewing and illumination angles. Accordingly, $c_b(\lambda)$ and $c_s(\lambda)$ span a *dichromatic plane* in the infinite-dimensional vector space of spectral power distributions, and the light $L(\lambda,i,e,g)$ reflected from any object point lies in this plane. We can assume without loss of generality that $m_s(i,e,g)$ and $m_b(i,e,g)$ are variables in the range from 0 to 1. All light mixtures from an object can then be circumscribed in the dichromatic plane by a parallelogram, as shown in Figure 3.

An investigation of the geometrical properties of surface and body reflection reveals that the light mixtures form a dense spectral cluster in the dichromatic plane. The shape of this cluster is closely related to the shape of the object. We can use this relationship to determine characteristic features of the

spectral clusters. For illustration purposes, we will assume in the following discussion of spectral histograms that body reflection is approximately Lambertian and that surface reflection is describable by a function with a sharp peak around the angle of perfect mirror reflection. Note, however, that this analysis is not limited to a particular geometric reflection model.

Figure 4: The shape of the spectral cluster for a cylindrical object

Figure 4 shows a sketch of a shiny cylinder. The left part of the figure displays the magnitudes of the body and surface reflection components. The curves show the loci of constant body or surface reflection. The darker curves are the loci of constant surface reflection. Since $m_s(i,e,g)$ decreases sharply around the object point with maximal surface reflection, m_{bmax} , these curves are shown only in a small object area. We call the points in this area *highlight points*. The remaining object points are *matte points*. The right part of the figure shows the corresponding spectral histogram in the dichromatic plane. As we will describe below, the object points form two linear clusters in the histogram.

For **matte points**, the surface reflection component of the reflected light is negligible and all the reflected light comes from body reflection. The observed light at such points thus depends only on $c_b(\lambda)$, scaled by $m_b(i,e,g)$ according to the geometrical relationship between the local surface normal of the object and the viewing and illumination directions. Consequently, the matte points form a *matte line* in the dichromatic plane in the direction of the body reflection vector, $c_b(\lambda)$, as shown in the right part of Figure 4.

Highlight points exhibit both body reflection and surface reflection. However, since $m_s(i,e,g)$ is much more sensitive to a small change in the photometric angles than $m_b(i,e,g)$, the body reflection component is generally approximately constant in a highlight area, as displayed by the curve with label m_{bH} in Figure 4. Accordingly, the second term of the Dichromatic Reflection Model equation (2) has a constant value, $m_{bH}(i,e,g) c_b(\lambda)$, and all spectral variation within the highlight comes from varying amounts of $m_s(i,e,g)$. The highlight points thus form a *highlight line* in the dichromatic plane in the direction of the surface reflection vector, $c_s(\lambda)$. The line departs from the matte line at position $m_{bH}c_b(\lambda)$, as shown in Figure 4. More precisely, the highlight cluster looks like a slim, skewed wedge because of the small variation of the body reflection component over the highlight. If more than one highlight exists on a concave object, each of them describes a highlight line in the dichromatic plane (see Figure 5).

Figure 5: Cluster shapes for concave objects

The combined spectral cluster of matte and highlight points thus looks like a skewed T or a skewed comb. The skewing angle of the T depends on the spectral difference between the body and surface reflection vectors while the position and width of the highlight line depend on the illumination geometry: If the phase angle g between the illumination and viewing direction at a highlight is very small, the incidence direction of the light is close to the surface normal. The underlying amount of body reflection thus is very high. The highlight line then starts near the tip of the matte line, and the skewed T becomes a skewed L. The wider the phase angle g is, the smaller is the amount of underlying body reflection and the further away is the highlight line from the tip of the matte line. Thus, the shape of the histogram is related to the geometry of the scene.

2.4. Taking Color Images with Real Cameras: A Sensor Model

So far, we have described light reflection in a theoretical, physical model. However, the observed pixels values are also influenced by the characteristics of the recording camera. This section covers the influence of such camera characteristics on the image data. Since some of these influences disturb the light reflection properties stated in the Dichromatic Reflection Model, we need to provide methods that restore the physical properties of light reflection. Where this is impossible, our image analysis algorithms must be able to detect or tolerate the inaccuracies in the image data.

We will first discuss the implications of representing a continuous light spectrum as a finite dimensional vector. Next, we present problems associated with the limited dynamic range of cameras. Finally, we analyze the effects of the changes in the camera's responsivity as a function of both wavelength and intensity. Since we use a CCD-camera (SONY AVC-D1) to obtain our images, we concentrate on the issues occurring with such cameras.

2.4.1. Spectral Integration

The Dichromatic Reflection Model describes light reflection using the physically precise representation: the continuous light spectrum. It is, however, impossible to measure and process light information in this way. The human eye as well as light sensing devices use a finite set of samples to describe the spectrum. Such sample measurements are obtained by filtering the light spectrum and integrating over the filtered spectrum. We call this process *spectral integration*. It integrates the amount of incoming light, $L(\lambda, i, e, g)$, weighted by the spectral transmittance, $\tau_f(\lambda)$, of the respective filter and the spectral responsivity of the camera, $s(\lambda)$, over all wavelengths λ :

$$C_f = \int L(\lambda, i, e, g) \tau_f(\lambda) s(\lambda) d\lambda \quad (3)$$

We use a red, a green and a blue color filter (Wratten filters with numbers 25, 58 and 47), thus reducing the infinite-dimensional vector space to a three-dimensional space. The spectrum of an incoming light beam at pixel position (x,y) is represented by a triple $\mathbf{C}(x,y) = [R,G,B]$, where x and y are determined by i , e , and g and by the position of the object relative to the camera.

Spectral integration from the infinite-dimensional vector space of spectral power distributions to a three-dimensional color space is a linear transformation [10, 32]. Thus, if a light beam $L(\lambda)$ is a mixture of two lights, $L_1(\lambda)$ and $L_2(\lambda)$, in the infinite color space, then the color vector \mathbf{C} , resulting from spectral integration on $L(\lambda)$ is the same linear combination of the color vectors \mathbf{C}_1 and \mathbf{C}_2 that result from spectral integration on $L_1(\lambda)$ and $L_2(\lambda)$. For this reason, and because the photometric angles i , e , and g depend on x and y , the linear relationship between reflected light and the colors of surface and body reflection, as stated in equation (2), is maintained under spectral integration. We thus obtain the Dichromatic Reflection Model for the three-dimensional color space:

$$\mathbf{C}(x,y) = m_s(i,e,g) \mathbf{C}_s + m_b(i,e,g) \mathbf{C}_b \quad (4)$$

The color pixel value $\mathbf{C}(x,y)$ is thus a linear combination of the two color vectors, $\mathbf{C}_s = [R_s, G_s, B_s]$ and $\mathbf{C}_b = [R_b, G_b, B_b]$, which describe the colors of surface and body reflection on an object in the scene. Within the three-dimensional color space, \mathbf{C}_s and \mathbf{C}_b span a dichromatic plane which contains a parallelogram in which the color clusters lie. This is shown in Figure 6.

Figure 6: Color cluster in a three-dimensional color space

Note that, in principle, spectral integration introduces a nonuniqueness, called *metamerism* [18, 19], in the resulting finite representation: the infinite-dimensional color space is partitioned into a set of equivalence classes, each of which contains an infinite number of spectral curves that are all represented by the same color. However, this is of no concern to us since the finite color space maintains the linear relationships between the reflected light mixtures from one object.

2.4.2. Limited Dynamic Range

Real cameras have only a limited dynamic range to sense the brightness of the incoming light. This restricts our analysis of light reflection, as measured in real color images, to a *color cube*, as shown in Figure 7. Its walls denote the upper and lower limits in the dynamic range of the respective color bands.

Figure 7: Color cluster in the color cube, with color clipping

If the incoming light is too bright at some pixel position, thus exceeding the dynamic range of the

camera, the camera cannot sense and represent it adequately. This effect may occur in one, two, or all three color bands, depending on the color of the light. We refer to this effect as *color clipping*. It can be a problem for measuring the color of highlights or bright objects. In the color histograms, clipping causes the clusters to bend along the walls and edges of the color cube (see Figure 7). In a very bright area, the dynamic range of the camera may be exceeded in all three color bands, resulting in white pixels even though the color of the light incident at the camera may not be white. Such *clipped color pixels* do not follow the characteristics of the dichromatic reflection model and must thus be distinguished from matte pixels and highlight pixels.

If a CCD-camera is used to obtain the images, too much incident light at a pixel may completely saturate the sensor element at this position [4]. This may cause *blooming* in the camera [3] which means that more charges are generated at a pixel than the relevant sensor element can hold. The excess carriers then spread out to the adjacent pixels and change their values proportional to the magnitude of the overload. In case of very bright spots in the scene, blooming may travel quite far from the sensor element that was originally overloaded, thus causing extended white areas in an image. We call pixels which have increased color values because of blooming *bloomed color pixels*. Their color values may be increased arbitrarily in one, two or all three color bands.

Both color clipping and color blooming are a problem for image sensing with most currently available cameras; especially, for the sensing of scenes with highlights. They interact badly with another effect that is common in CCD-cameras: spatial averaging between neighboring pixels that is commonly built into CCD-cameras to reduce the variance in the sensing characteristics in the array. Because of such spatial averaging, color clipping and blooming is "smoothed" over an image area. As a consequence, we cannot detect these effects by merely thresholding the color image at maximal intensity: because of the averaging, the effect is spread to neighboring pixels, increasing their values and decreasing the values of the clipped pixels. For this reason, we suspect color values in the upper 10% of the intensity scale in our images to be potentially influenced by color clipping or blooming.

2.4.3. Color Balancing and Spectral Linearization

We now describe how the measured color pixel values are influenced by the responsivity of CCD-cameras to varying wavelengths and incident flux. This analysis applies to black-and-white CCD-cameras.

Figure 8: Spectral responsivity of CCD-cameras and the human eye

A typical curve of the spectral responsivity of CCD-cameras is shown in Figure 8 [4]. Because of the ramp-like spectral responsivity of CCD-cameras within the visible spectrum, the blue band registers much less contrast than the red band and thus lowers the signal to noise ratio. In order to provide an equal scaling on the three color axes in the color space, we need to rescale the pixel data separately in the color bands. We refer to this procedure as *color balancing*. A simple multiplication of the color data by constant factors, however, does not improve the dynamic range in the color bands and thus does not improve the quality of our color data; it only introduces gaps into the color distribution in the color histogram. Instead, we balance the color bands by controlling the camera aperture during the picture taking process, using apertures for green and blue exposures under tungsten light that are $\frac{1}{2}$ and $1\frac{1}{2}$ f-stops (respectively) higher than the aperture used for the red exposure. The upper left quarter of Figure 9 shows the image of a Macbeth Color Checker color chart that we obtained using color balancing by aperture control. The lower left quarter shows an image of the same color chart taken without color balancing. The unbalanced image looks more orange, since the blue and green color bands do not contribute appropriately to the pixel colors. For this reason, its grey scale in the lowest row of the chart is brown. The grey scale of the color balanced image, on the other hand, is approximately grey. The pictures shown hereafter have been color balanced by aperture control.

Figure 9: The effect of color balancing and gamma correction on color images

Note that CCD-cameras are very sensitive to infrared light. Since most color filters in the visible spectrum are nearly transparent in the infrared, camera responsivity to the infrared range can add a significant amount of measured intensity to the color values. This may "wash out" the colors in the image significantly. To overcome this, we always use a total IR suppressor (Corion FR-400) in front of our camera.

The color pixels also depend on the camera response to incident light flux. Most modern sensors have a linear response to flux [4]. However, since the luminous output of display devices is related by a power law to the driving voltage, $L=kV^\gamma$, it is common practice in the television industry to produce a signal proportional to $Y^{1/\gamma}$ in the camera, with 2.2 as a typical value for γ . This is known as *gamma-correction*. As a result, the camera output is related by this inverse power law to the incident flux. We have measured this relationship for our camera in the Calibrated Imaging Laboratory at Carnegie-Mellon University by comparing the known reflectance values of the grey scale of the Macbeth color chart with the measured, averaged intensity values in the corresponding image areas in the three color bands. The

resulting diagrams are shown in the right quarters of Figure 9. The upper diagram displays the measurements from a color image that was taken with color balancing by aperture control. The horizontal axis represents the known reflectance values of the grey blocks. The vertical axis represents the measured intensities in the three color bands. The points in the diagram approximately lie on the typical curve for gamma-correction. The lower diagram displays the same measurements from the unbalanced color image. It demonstrates the decreasing spectral responsivity of our camera to short wavelengths.

Figure 10: Curvature in the color space, caused by gamma-corrected cameras

Figure 10 demonstrates how the gamma-correction of our camera disturbs the linear properties of the Dichromatic Reflection Model. The upper left quarter shows the image of an orange cup. The cup does not reflect a significant amount of surface reflection into the camera. Thus, all pixels in the outlined image area are matte pixels and should form a straight matte line in the color space. The upper right quarter of Figure 10 demonstrates that this is not the case: gamma-correction introduces curvature into the color space, distorting the linear properties of the Dichromatic Reflection Model.

To linearize the color data, we fit interpolating cubic splines to the measured responsivity data, separately in each color band, as suggested by LeClerc [22]. To model very bright light reflection, as can occur in highlights, we rescale the linearization function and extrapolate it from the brightest measurement outwards by a square root curve. The rescaling is necessary because linearizing very bright spots in the original image produces values beyond the representable range of pixel values, due to the slope of the fitted curve. We then use these functions to generate a look-up table relating the measured intensities in every color band to the incident flux. The upper right box of Figure 11 shows the cubic splines that were fitted to the responsivity data of Figure 9. The lower left box of Figure 11 displays the linearized image of the color chart. When compared with an unlinearized image (upper left box), the colors of the linearized image look more saturated. The image is also darker, due to the rescaling of the intensities. The lower half of Figure 10 shows that when our linearization method is applied to the image of the orange cup, the resulting color histogram is approximately linear, pointing in the direction of the body reflection vector. All pictures shown hereafter are linearized.

Figure 11: Spectral linearization with cubic spline functions

3. Analysis of Real Color Images

After discussing in section 2 what influences the pixel values in a color image, we will now show how the described properties can be used for color image understanding. We first demonstrate that the Dichromatic Reflection Model accounts well for color variation found in real color images. We will then present an algorithm that uses the described properties to generate intrinsic reflection images and discuss the results. Such images show only the body reflection component or only the surface reflection component of every pixel. Since the intrinsic reflection images have simpler relationships to the illumination geometry than the original image, we expect that this algorithm will be a suitable preprocessor for other image understanding methods, such as stereo vision, shape from shading and shape from highlights.

3.1. Color Clusters from Real Color Images

There exist some inherent assumptions in the Dichromatic Reflection Model. We will now demonstrate that real color data from dielectric, inhomogeneous materials under varying illumination colors follows the theory of this model. We have taken a series of color images in the Calibrated Imaging Laboratory of plastic cups, glossy paper folders and a ceramic cup under white, yellow and pink light. Black curtains on the walls and a black ceiling were used to eliminate indirect illumination in the visible spectrum of the scene.

Figure 12: Three plastic cups under yellow light

Figure 13: Color histogram of three plastic cups under yellow light

Figure 12 shows the image of a green, a yellow and an orange plastic cup, illuminated by yellow light. Figure 13 shows the color histogram that results from projecting the color pixels from the outlined areas of Figure 12 into the color cube. The color histogram consists of several linear clusters. There are three linear clusters which start from the origin of the color cube at the black corner. These clusters contain the matte pixels of the three cups. They lie along the orange, yellow and green matte lines of the three cups. The matte clusters are connected to linear highlight clusters which lie along the respective highlight lines of the cups. For the yellow cup, the matte and highlight clusters are nearly collinear because the material and illumination color are nearly the same. Note that the highlight clusters from all three cups are approximately parallel to each other. Furthermore, they are parallel to the color vector describing the yellow illumination because the refraction indices of the cups are fairly constant over the visible spectrum. We thus can use the direction of the highlight clusters to determine the illumination color. This is useful

for color constancy algorithms [6, 23, 24] which try to remove the influence of the illumination color from the body reflection component, thus "normalizing" the image to a standard white illumination.

Further linear clusters in the color histogram contain clipped color pixels. These are displayed in Figure 13 in their complementary color (blue). We identify such color clipping through a high threshold on the measured intensities in the three color bands. The clipped color pixels come from the centers of the highlights where the reflected light exceeds the dynamic range of the camera. Accordingly the centers of the highlights contain white pixels, whereas the pixels closer to the highlight boundaries exhibit yellow surface reflection.

Figure 14: Orange cup under yellow light

Figure 15: Color histogram of the orange cup under yellow light

Figures 14 and 15 demonstrate how blooming influences the color distribution in the color histogram. The figures show an image and color histogram of an orange cup under yellow light. The highlight center on the cup is so bright that charges spread out to adjacent pixels, increasing their values. Because the cup is orange and the illumination is yellow, blooming occurs mainly in the red color band. Bloomed color pixels thus have an increased red color value, while the green and blue values are correct. As a result, the highlight cluster is spread out along the red axis of the color cube. The same effect is present - to a smaller degree - in the green highlight cluster in Figure 13. Since blooming cannot always be completely avoided in pictures with highlights, we need to account for it in our color histograms. We assume that blooming occurs only at a few pixels in the image and that it increases the color pixel values arbitrarily. It is then unlikely that such color pixel values occur often in the image. We thus suspect that colors with a low frequency count in the color histogram are caused by blooming or noise. Figure 15 displays such color pixels in white.

Figure 16: Images and color clusters of glossy paper and a ceramic cup under white light

Figure 16 shows that the Dichromatic Reflection Model applies to a variety of dielectric, inhomogeneous materials. The figure displays the images and the color histograms of a ceramic cup and of two folders made out of glossy paper. The folders are slightly bent such that they have highlights in only a small object area. The corresponding color histograms consist of matte and highlight clusters, as described above, thus supporting the theory of the Dichromatic Reflection Model.

3.2. Detecting and Removing Highlights

We will now show how the Dichromatic Reflection Model can be used to detect and remove highlights in color images. The Dichromatic Reflection Model describes the color of every pixel in a color image as a mixture of two colors, the color of the body reflection component of the viewed object, \mathbf{C}_b , and the color of the surface reflection component, \mathbf{C}_s . We have developed an algorithm that exploits the T-shape of a color cluster to reverse this color mixing process. It fits a matte and a highlight line to the cluster and decomposes the color pixels into their constituent matte and highlight components.

Currently, our algorithm does not have the capability to distinguish between several color clusters in the color cube, since some of its operations (such as plane fitting) are applied globally to all entries in the color cube. If there are several objects in the scene, our algorithm requires that the image be previously segmented into image areas which each correspond to a single color cluster. It then applies its analysis of color variations separately to each segmented area. Most current color image segmentation methods [27, 28] cannot provide the required segmentation since they separate the highlights from the surrounding matte object areas and thus cut the color clusters into pieces. Gershon has reported encouraging first results in merging matte and highlight areas in a postprocessing step after an initial segmentation [9]. We are currently developing a segmentation algorithm that exploits the Dichromatic Reflection Model for its analysis of local color variations in an image, generating local hypotheses about dominant linear color clusters and dichromatic planes. This research is still under development. For the work described here we have segmented the images by hand.

3.2.1. Determining the Body and Surface Reflection Vectors of an Image Area

We will now describe our method to determine the body and surface reflection vectors of the segmented image areas. The algorithm projects the color pixels from one chosen image area into the color cube, ignoring clipped or bloomed pixels as described in section 3.1. According to the Dichromatic Reflection Model, we may expect that the color pixels lie in a dichromatic plane in the color cube. Our algorithm thus fits a plane to the cluster, using principal component analysis to determine the eigenvalues and eigenvectors of the covariance matrix of the cluster [1, 11]. We sort the eigenvalues by decreasing size. The third eigenvector then determines the direction of minimal color variation in the cluster. The physically relevant color variation occurs in the plane which is perpendicular to the third eigenvector. We thus project all color pixels perpendicularly into this dichromatic plane.

Our analysis of the relationship between object shape and cluster shape in section 2.3. predicts that

the color cluster forms a skewed T in the dichromatic plane and that the orientations of the linear sections of the T determine the body and surface reflection vectors of the cluster. Determining the corresponding linear matte and highlight subclusters in the cube is not trivial because neither a classification of the pixels into matte and highlight pixels nor the orientation and position of the subclusters are known. Our algorithm determines the two subclusters in a two-stage process.

The first stage is intended to render an initial, rough estimate of the matte and highlight line. We fit a convex hull of line segments to the contour of the cluster, using a recursive line spitting method [29]. To do so, we choose the line L between the brightest and darkest points in the cluster as the starting line. We divide the plane into buckets which have the form of slim stripes running perpendicularly to L (see figure 17). We compute the color distance of all pixels from L and store the maximal distance within every stripe s as a function $d(s)$. This function approximates the contour of the color cluster. Its shape, however, is very susceptible to noise. In order to have it reflect only the large scale color variations, we smooth $d(s)$, averaging its values over a window. We then approximate the smoothed contour by a convex polygon. For this, we determine the stripe s_i with maximal distance $d(s_i)$ from L . If $d(s_i)$ exceeds a given threshold, L is split into two segments, and the distances of the color pixels from the new line segments are recomputed. Both segments are then split recursively into smaller segments until all distance measurements are below the threshold. The choice of the threshold determines how significant an orientation change of the contour needs to be before it influences the shape of the polygon. For all results shown in this paper, we have required $d(s_i)$ to exceed 20.

Figure 17: Fitting a convex polygon to the color cluster

Our program next classifies the line segments. We know from the discussion of the cluster shape in section 2.3., that matte pixels depend only on the body reflection component. We thus expect that the matte pixels form a matte line that approaches the origin of the color cube. Accordingly, we assume that the line segment closest to the black corner of the cube approximates the matte line. Section 2.3. further states that the highlight line starts from the matte line. Accordingly, our algorithm classifies the next line segment connected to the matte line as the highlight line. If there are further line segments connected to the highlight line, we assume that they are related to clipped color pixels that were not detected by the heuristics of section 3.1. All further line segments are thus classified as clipped color lines.

The line fitting method described so far provides us with an initial estimate for the matte and highlight lines. But for several reasons, it cannot be expected to generate very accurate results. Most importantly,

the fitted convex polygon sweeps across the concavity at the corner between the matte and highlight line. As a result, the estimated highlight line will generally describe an orientation with too little slope. Secondly, our initial plane fitting step was applied blindly to all pixels in the color cluster. If some clipped color pixels escaped the discarding heuristics, they may have biased the orientation of the dichromatic plane towards the white corner of the color cube. Since the fitted lines lie in the plane, they are then also biased towards white. For these reasons, we consider the matte and highlight lines to be only an initial approximation to the linear parts of the color cluster. Their orientations and positions will be improved in the second stage of the line fitting process.

In the second stage, we use the preliminary result of the first stage to coarsely classify the color pixels into matte pixels, highlight pixels and clipped pixels, depending on the line they are closest to (see figure 18). We then refit the matte and highlight lines separately to the matte pixels and to the highlight pixels, using principal component analysis. The first eigenvectors describe the directions of the largest color variation in the two clusters. They determine our new estimates for the matte and the highlight line. The corresponding body and surface reflection vectors are parallel to the matte and highlight lines, but start at the origin of the color space. In the above described steps, we assume that the classification assigns correct labels to most of the pixels in the cluster. We know, however, that some highlight pixels may have been misclassified as matte pixels because the initial estimate of the highlight line may have been wrong, as discussed above. Some highlight pixels near the starting point of the highlight line (i.e.: in the concavity) may be closer to the initial matte line than to the initial highlight line (see figure 18). We expect that this will generally not be a problem because the number of such false matte pixels is generally very small, when compared to the number of correct matte pixels. The orientation of the matte line is thus not influenced significantly by false matte pixels. This argument relies on the fact that highlights are generally much smaller than the surrounding matte object area. For flat objects, this may not be true. We may then have to alternate several times between reclassifying the pixels and refitting the lines.

Figure 18: Classifying the color pixels

Figure 19: Color histogram and fitted vectors of the three plastic cups under yellow light

Table 1: Matte and highlight lines of the three plastic cups under yellow light

Table 2: Body and surface reflection vectors of the three plastic cups under yellow light

Figure 19 shows the matte and highlight lines that were fitted to the color histogram of the three cups

under yellow light. It displays a matte line and a highlight line (M0, H0) for the green cup, a matte and a highlight line (M1, H1) for the yellow cup, and a matte line, a highlight line and a clipped color line (M2, H2, C2) for the orange cup. A list of the start and end points of all lines is given in Table 1. Table 2 contains the corresponding body and surface reflection vectors, normalized to unit length. As expected, the body reflection vector of the green cup is dominated by its green component while the body reflection vectors of the yellow and orange cup are increasingly dominated by their red component. Currently, we do not have the facilities to evaluate these results on a quantitative basis. The surface reflection vectors of all three cups are approximately the same, thus supporting the observation that the three highlight clusters are parallel. To obtain an estimate of the illumination color, we have taken a picture of a grey card (Munsell N 8.5/ with 68.4% reflectivity) under yellow light. The card reflects the illuminating light without changing its spectral curve significantly. We can thus use the averaged color pixel values of the image as an indication of the color of the light source. The resulting, normalized estimate for our yellow illumination is also shown in Table 2. It is very close to the fitted surface reflection vectors of the three cups. This suggests that our method for determining the body and surface reflection vectors of an image area is also a useful estimator of the illumination color.

3.2.2. Generating Intrinsic Reflection Images

Once we have determined the body and surface reflection vectors, \mathbf{C}_b and \mathbf{C}_s , of an image area, we can decompose the color pixels into their constituent body and surface reflection components. Equation 4 in section 2.4.1. describes the color of every pixel as linear combination of these two color vectors:

$$\mathbf{C}(x,y) = m_s \mathbf{C}_s + m_b \mathbf{C}_b$$

\mathbf{C}_b , \mathbf{C}_s and their cross product, $\mathbf{C}_b \times \mathbf{C}_s$, define a new (not necessarily orthogonal) coordinate system in the color cube. This coordinate system describes every color in the cube in terms of their amounts of body reflection, surface reflection and noise ϵ , as given by the color distance from the dichromatic plane (see figure 20). There exists an affine transformation, and thus a linear transformation matrix T , which transforms any color vector $\mathbf{c} = [R,G,B]^T$ from the initial coordinate system into a vector $\mathbf{d} = [m_b, m_s, \epsilon]^T$ in the new coordinate system

$$\mathbf{d} = T\mathbf{c} \quad (5)$$

After computing T from \mathbf{C}_b and \mathbf{C}_s , we can thus transform every color pixel in the image area into its constituent body and surface reflection components, m_b and m_s . By selecting the m_b -component of every transformed pixel, we generate the *body reflection image* of the image area. By selecting the m_s -component, we generate the corresponding *surface reflection image*.

Figure 20: Decomposing a color pixel into its constituent body and surface reflection components

Figure 21: Intrinsic body reflection image of the three plastic cups under yellow light; clipped pixels and bloomed pixels are shown in black

Figure 22: Intrinsic surface reflection image of the three plastic cups under yellow light; clipped and bloomed pixels are shown in black

Figures 21 and 22 show the intrinsic body and surface reflection images that are thus generated from Figure 12. The colors of clipped and bloomed pixels have not been separated into their reflection components. They are displayed in black, as are the pixels outside the marked image areas in Figure 12. We expect that flagging clipped and bloomed pixels in a separate intrinsic *error image* will provide useful information to other image understanding algorithms like shape from shading. It helps them to base their analysis only on undistorted image data. If necessary, they may then extrapolate the derived shape information into areas with clipped or bloomed pixels. Our method of detecting and discarding distorted color pixels may also enable us to automatically control the camera aperture so that we can obtain color images with minimal clipping and blooming.

The intrinsic body reflection image in Figure 21 demonstrates the success of our algorithm as a method for removing highlights in undistorted image areas. It also shows that the algorithm detects very well the areas in which the pixel values are distorted. The salt-and-pepper noise on the yellow cup is due to the approximate collinearity of its body and surface reflection vectors. Small amounts of noise in its color pixels have a large effect on the computed amounts of underlying body reflection. The intrinsic surface reflection image in Figure 22 shows that our algorithm is able to reliably detect highlights in color images. It also determines the illumination color, as reflected from the surfaces of the cups. Note that this approach to the detection of highlights in color images provides quantitative information on the amount of surface reflection that is present at every pixel, as opposed to a boolean decision. Our highlight images may thus be useful input for methods that determine object shapes from highlights [12].

3.2.3. Restoring the Colors of Clipped and Bloomed Pixels

In addition to flagging clipped and bloomed color pixels, we try to restore their correct color values. For this operation, we distinguish between bloomed matte pixels and bloomed highlight pixels, depending on the reflection line they are closest to. To restore the color values of clipped or bloomed pixels, we exploit the observation that in many cases, clipping and blooming occurs only in one or two color bands. The pixels may thus have correct data in the other color bands. We assume that the smallest of the three

values of a color pixel stems from a color band without clipping and blooming. We then replace the clipped or bloomed pixels by pixels on the matte or highlight line that have the same value in the undistorted band.

Figure 23: Intrinsic body reflection image of the three plastic cups under yellow light; with restored clipped and bloomed pixels

Figure 24: Intrinsic surface reflection image of the three plastic cups under yellow light; with restored clipped and bloomed pixels

Figures 23 and 24 show the intrinsic images of Figure 12 after clipped and bloomed pixels were restored. In the intrinsic body reflection image, the pixels were restored reasonably well. However, there are visible intensity changes between some distorted and undistorted pixels. There is a limit to how much can be done. Since all clipped pixels and all bloomed highlight pixels are restored to lie exactly on the highlight line, they are all projected into the same body reflection component. We thus lose all shading information for these pixels, which is normally encoded in the distance of a highlight pixel from the highlight line. This may cause a perceptible intensity change between distorted and undistorted pixels in the body reflection image. We expect this data to be used only in combination with the corresponding intrinsic error image. It may serve as additional, approximate information to infer the object shape in areas with distorted color data.

The intrinsic surface reflection image, on the other hand, has gained significantly from our restoring procedure. The clipped and bloomed pixels fill the centers of the highlights vividly with the proper surface reflection color.

3.3. Results and Discussion

We have applied our algorithm to several images of plastic cups, glossy paper folders and a ceramic cup under white, yellow and pink illumination. Some results have already been presented and discussed in previous sections. We will now show further results.

Figure 25: Intrinsic reflection images of the glossy paper folders under white light

Figure 26: Intrinsic reflection images of the ceramic cup under white light

Figure 27: Intrinsic reflection images of the three plastic cups under pink light

Figure 28: Intrinsic reflection images of the three plastic cups under white light

Table 3: Body and surface reflection vectors of the paper folders under white light

Table 4: Body and surface reflection vectors of the stripes on the ceramic cup under white light

Table 5: Body and surface reflection vectors of the three plastic cups under pink light

Table 6: Body and surface reflection vectors of the three plastic cups under white light

Figures 25 through 28 demonstrate that our algorithm is able to detect and remove highlights well on a variety of dielectric materials under varying illumination colors. In all cases, the body reflection images have been restored very well. When judging by eyeball inspection, the body reflection images of the paper folders and of the ceramic cup do not seem to contain any indication of a highlight. We will evaluate this statement quantitatively in the future by testing whether shape from shading methods determine the correct object shape from these images. In the images of the plastic cups, the highlights could not be perfectly removed at clipped and bloomed pixels for reasons that were discussed in section 3.2.3.

The surface reflection images locate very well the highlights in the images. Note the gradual change of the amount of surface reflection on the paper folders. Since the surface reflection image captures this gradual change, it may provide the means for a quantitative approach to determining object shape from highlights.

Tables 3 through 6 show the normalized reflection vectors of the objects in all images. Our algorithm performed reasonably well in detecting the white illumination color from the highlights of the paper folders and the ceramic cup and the pink illumination from the highlights of the cups. We are currently investigating what caused the larger variations in the surface reflection vectors of the cups under white light (Table 6). There may be several contributing factors. Most importantly, color clipping can alter the orientation of a surface reflection vector significantly. This is most evident for the surface reflection vector of the yellow cup. In this case, the starting point of the highlight cluster is located at very high color values, especially in the red component of the color pixels. This does not leave a lot of room for the highlight cluster before it is bent at the red wall of the cube. Furthermore, our clipping threshold does not capture all clipped pixels on this object; and it provides a cutoff plane which is parallel to the wall of the cube, but at a skewed angle to the highlight cluster. As a result, disproportionately many color pixels are close to the clipping threshold and on the inner side of the skewed T, and the angle between the fitted

highlight line and the normal to the wall is too large.

Another reason for wrong orientations of the surface reflection vectors may be that the classifications of pixels into matte and highlight pixels is imprecise (see section 3.2.1, Figure 18). This can result in selecting a bad starting point for the highlight line. The determination of the starting point also depends on the selection of the threshold for recursive line splitting and on the width of the averaging window for smoothing the distance function $d(s)$.

There seems to be a general trend in Table 6 of overestimating the blue component and underestimating the red component of the illumination vector. There may be measurement errors involved, concerning the color balancing and spectral linearization parameters. They might have been caused, for example, by a shading gradient on the color chart.

All of these hypotheses are subject to further investigations, as well as ways of avoiding such problems or improving the results. One way to improve the orientations of the surface reflection vectors is given by the observation that the vectors should be the same for all highlights. This constraint can be useful, if there exist several approximately correct surface reflection vectors and only a few bad ones. We may then determine a single, good surface reflection vector that best fits a weighted combination of all highlight clusters.

Our examples show that we can detect and remove highlights reasonably well in the given color images. We will now discuss what the limitations of our approach are. The main feature of our algorithm is that it exploits the color difference between surface reflection and body reflection. There needs to be a certain angle between the orientations of the body and surface reflection vectors of an object, in order for our program to detect them correctly. How big this angle needs to be depends on the lengths of the color clusters and on the selected thresholds for recursive line splitting and smoothing. These, in turn, depend on the expected amount of noise in the color clusters. Our example of the yellow cup under yellow light in section 3.2.1. has demonstrated that the algorithm was able to distinguish the surface reflection vector correctly from the body reflection vector when the angle between them was only 15 degrees. In this example, however, both color clusters were fairly long. In general, we may have problems with very dark or very bright objects (when either of the two color clusters is very short) and with objects with unsaturated colors, such as pastels.

We cannot apply our algorithm to metals because they do not have a significant body reflection component. This means that the color pixels from a metal form only a highlight cluster in the color cluster; no matte cluster exists. We thus do not need to separate its surface reflection component from its body reflection component. Problems may arise when a scene contains both metallic and dielectric objects and we need to determine whether a single linear cluster in the color cube should be interpreted as a highlight cluster, generated by a metal, or as a matte cluster, generated by a dielectric object without a highlight. To answer this question, we have to evaluate geometric information. One viable approach might be to use shape from shading to generate the shape of a hypothetical dielectric object and to compare it with the shape of a hypothetical metal, as generated by a shape from highlights method.

Another important requirement in our algorithm is that the subclusters form lines. There may be several reasons why this may not be the case for the matte clusters of all objects in real scenes. Firstly, the matte color clusters may be curved because the illumination color changes in the scene. The most important factor here is inter-reflection between objects. If inter-reflection exists between two objects, A and B , the illumination at some points on object B is a mixture of the direct illumination and the reflected light from parts of A . Consequently, the light reflected from those points on B depends on the color of A . The color cluster then contains entries that spread towards the color of object A .

The matte cluster may also be nonlinear because the reflection properties of the object vary, either gradually or abruptly. This happens if there are color texture patterns on the object. The ceramic cup in Figure 26 is an example of an object that has abrupt color changes. The resulting color clusters are widely separated in the color cube. Thus, if we can somehow determine object areas belonging to a single color cluster, our current algorithm is able to detect and remove highlights in each of these areas. If the color changes are more gradual, as in a rainbow, or if the object areas with constant reflection properties are too small, the color pixels do not form distinctive, separate color clusters in the color cube and our method breaks down.

4. Conclusions and Further Work

In this paper, we have demonstrated that it is possible to analyze real color images by using a color reflection model. Our model accounts for highlight reflection and matte shading, as well as for some characteristics of cameras. By developing a physical description of color variation in color images, we have developed a method to separate highlight reflection from matte object reflection. The resulting

intrinsic reflection images have a simpler relationship to the illumination geometry than the original image and may thus improve the results of many other computer vision algorithms, such as motion analysis, stereo vision, and shape from shading or highlights [7, 12, 14, 35]. Since the surface reflection component of dielectric materials generally has the same color as the illumination, we can also determine the illumination color from the intrinsic surface reflection image, information which is needed by color constancy algorithms [6, 23, 24].

The key points leading to the success of this work are our modeling of highlights as a linear combination of both body and surface reflection and our modeling of the camera properties. With few exceptions [6, 8, 23, 31], previous work on highlight detection in images has assumed that the color of highlight pixels is completely unrelated to the object color. This assumption would result in two unconnected clusters in the color space: one line or ellipsoid representing the object color and one point or sphere representing the highlight color. Our model and our color histograms demonstrate that, in real scenes, a transition area exists on the objects from purely matte areas to the spot that is generally considered to be the highlight. This transition area determines the characteristic shapes of the color clusters which is the information that we use to detect and remove highlights. This view of highlights should open the way for quantitative shape-from-gloss analysis, as opposed to binary methods based on thresholding intensity.

By modeling the camera properties, we are able to obtain high quality color images (through color balancing and spectral linearization) in which most pixels maintain the linear properties of light reflection, as described in the Dichromatic Reflection Model. We can also detect most distorted color pixels in an image and thus generate an intrinsic error image which then guides our algorithm to separate only undistorted color pixels into their reflection components. We expect that the intrinsic error image will be similarly useful in guiding other computer vision algorithms, such as shape from shading. It may also enable us to automatically control the camera aperture so that we can obtain color images with minimal clipping and blooming.

Our approach to color image analysis may influence research in other areas of color computer vision. The color histograms demonstrate that all color pixels from one object form a single cluster in the color space. A color cluster thus groups pixels into areas of constant material properties, independently of influences from the illumination geometry, such as shading or highlights. This finding can be used for color image segmentation to distinguish material changes from shading or highlight boundaries.

Furthermore, the shape of the color clusters reveals some information about the illumination geometry and the object shapes. We are currently investigating these implications of the model to improve color image understanding methods, and we are also working on extending our model to account for other material types, such as metals, and more complex illumination conditions including ambient light, light sources in several colors, shadow casting and inter-reflection between objects.

Although the current method has only been applied in a laboratory setting, its initial success shows the value of modeling the physical nature of the visual environment. Our work and the work of others in this area may lead to methods that will free computer vision from its current dependence on statistical signal-based methods for image segmentation.

Acknowledgements

We would like to thank Ruth Johnston-Feller for her comments and suggestions. Richard Szeliski and Chuck Thorpe provided valuable comments on previous drafts of the article. We are also grateful to the reviewers. They helped to improve the paper significantly.

Footnotes

Affiliation of authors:

Department of Computer Science, Carnegie-Mellon, Pittsburgh, PA 15213, USA

Acknowledgement of financial support:

This material is based upon work supported by the National Science Foundation under Grant DCR-8419990 and by the Defense Advanced Research Projects Agency (DOD), ARPA Order No. 4976, monitored by the Air Force Avionics Laboratory under contract F33615-84-K-1520. Any opinions, findings, and conclusions or recommendations expressed in this publication are those of the authors and do not necessarily reflect the views of the National Science Foundation, the Defense Advanced Research Projects Agency, or the US Government.

References

- [1] D.H. Ballard, and C.M. Brown.
Computer Vision.
Prentice Hall, Inc., Englewood Cliffs, NJ 07632, 1982.
- [2] P. Beckmann and A. Spizzichino.
The Scattering of Electromagnetic Waves from Rough Surfaces.
MacMillan, New York, 1963.
pp 1-33, 70-96.
- [3] J.D.E. Beynon and D.R. Lamb (eds.).
Charge-coupled devices and their application.
McGraw-Hill, London, 1980.
- [4] W. Budde.
Optical Radiation Measurements. Volume 4: Physical Detectors of Optical Radiation.
Academic Press, New York, 1983.
- [5] R.L. Cook and K.E. Torrance.
A Reflectance Model for Computer Graphics.
ACM Transactions on Graphics 1(1):7-24, January, 1982.
also published in *Computer Graphics* 15(3), SIGGRAPH'81.
- [6] M. D'Zmura and P. Lennie.
Mechanisms of color constancy.
Journal of the Optical Society of America A (JOSA-A) 3(10):1662-1672, October, 1986.
- [7] L. Dreschler and H.-H. Nagel.
Volumetric Model and 3D Trajectory of a Moving Car Derived from Monocular TV Frame
Sequences of a Street Scene.
Computer Graphics and Image Processing 20:199-228, 1982.
- [8] R. Gershon.
The Use of Color in Computational Vision.
PhD thesis, Department of Computer Science, University of Toronto, 1987.

- [9] R. Gershon, A.D. Jepson, and J.K. Tsotsos.
Highlight Identification Using Chromatic Information.
In *Proc. of the first International Conference on Computer Vision (ICCV), London*, pages 161-171.
Computer Society Press (IEEE), Juni 8-11, 1987.
- [10] H. Grassmann.
On the Theory of Compound Colors.
Phil. Mag. , April, 1854.
- [11] H.H. Harman.
Modern Factor Analysis, second edition.
University of Chicago Press, Chicago and London, 1967.
- [12] G. Healey and T.O. Binford.
Local Shape from Specularity.
In J.M. Brady and A. Rosenfeld (editor), *Proceedings of the First International Conference on Computer Vision (ICCV)*, pages 151-161. Computer Society Press (IEEE), London, June 8-11, 1987.
Also appeared in DARPA-Image Understanding Workshop, pages 874-887, Los Angeles, CA, February 1987.
- [13] G. Healey and T.O. Binford.
The Role and Use of Color in a General Vision System.
In L.S. Bauman (editor), *DARPA-Image Understanding (IUS) workshop*, pages 599-613. Los Angeles, CA, February, 1987.
- [14] B.K.P. Horn.
Understanding Image Intensities.
Artificial Intelligence 8(11):201-231, 1977.
- [15] B.K.P. Horn.
Exact Reproduction of Colored Images.
Computer Vision, Graphics, and Image Processing (CVGIP) 26:135-167, 1984.
- [16] R.S. Hunter.
The Measurement of Appearance.
John Wiley and Sons, New York, 1975.
- [17] R.M. Johnston.
Geometric Metamerism.
Color Engineering 5(3):42, May-June, 1967.
- [18] R.M. Johnston.
Color Theory.
Pigment Handbook.
John Wiley & Sons, Inc., New York, 1974, pages 229-287, Chapter III-D-b.
- [19] D.B. Judd and G. Wyszecki.
Color in Business, Science and Industry.
John Wiley and Sons, New York, 1975.
- [20] G.J. Klinker, S.A. Shafer, and T. Kanade.
Measurement of gloss from color images.
In *Inter-Society Color Council (ISCC) 87 conference on Appearance*, pages 9-12. Williamsburg, VA, February 8-11, 1987.

- [21] G.J.Klinker, S.A. Shafer, and T. Kanade.
Using a Color Reflection Model to Separate Highlights from Object Color.
In J.M. Brady and A. Rosenfeld (editor), *Proceedings of the First International Conference on Computer Vision (ICCV)*, pages 145-150. Computer Society Press (IEEE), London, June 8-11, 1987.
Also appeared in DARPA-Image Understanding Workshop, pages 614-619, Los Angeles, CA, February 1987.
- [22] Y. LeClerc.
A Method for Spectral Linearization.
1986.
Private communication.
- [23] H.-C. Lee.
Method for computing the scene-illuminant chromaticity from specular highlights.
Journal of the Optical Society of America A (JOSA-A) 3(10):1694-1699, October, 1986.
- [24] L.T. Maloney and B.A. Wandell.
Color constancy: a method for recovering surface spectral reflectance.
Journal of the Optical Society of America A (JOSA-A) 3(1):29-33, January, 1986.
- [25] P. Moon.
A Table of Fresnel Reflection.
J. Math. Phys. 19(1), 1940.
- [26] F.E. Nicodemus, J.C. Richmond, J.J. Hsia, I.W. Ginsberg, and T. Limperis.
Geometrical Considerations and Nomenclature for Reflectance.
Technical Report NBS Monograph 160, National Bureau of Standards, October, 1977.
- [27] R. Ohlander, K. Price, and D.R. Reddy.
Picture Segmentation Using a Recursive Region Splitting Method.
Computer Graphics and Image Processing 8:313-333, 1978.
- [28] Y. Ohta, T. Kanade, and T. Sakai.
Color Information for Region Segmentation.
Computer Graphics and Image Processing 13:222-231, 1980.
- [29] T. Pavlidis.
Structural Pattern Recognition.
Springer Verlag, Berlin, Heidelberg, New York, 1977.
- [30] B.T. Phong.
Illumination for Computer Generated Pictures.
Communications of the ACM 18:311-317, 1975.
- [31] J.M. Rubin and W.A. Richards.
Color Vision and Image Intensities: When are Changes Material ?
Biological Cybernetics 45:215-226, 1982.
- [32] S.A. Shafer.
Describing Light Mixtures Through Linear Algebra.
Journal of the Optical Society of America (JOSA) 72(2):299-300, February, 1982.
- [33] S.A. Shafer.
Optical Phenomena in Computer Vision.
In *CSCSI-84*. Canadian Society for Computational Studies of Intelligence, Ontario, May, 1984.
Also available as technical report TR-135, Computer Science Department, University of Rochester, March 1984.

- [34] S.A. Shafer.
Using Color to Separate Reflection Components.
COLOR research and application 10(4):210-218, Winter, 1985.
Also available as technical report TR-136, Computer Science Department, University of Rochester, NY, April 1984.
- [35] C.E. Thorpe.
FIDO: Vision and Navigation for a Robot Rover.
PhD thesis, Computer Science Department, Carnegie-Mellon University, December, 1984.
Available as technical report CMU-CS-84-168.
- [36] K.E. Torrance, and E.M. Sparrow.
Theory for Off-Specular Reflection from Roughened Surfaces.
Journal of the Optical Society of America 57:1105-1114, September, 1967.
- [37] S.J. Williamson and H.Z. Cummins.
Light and Color in Nature and Art.
John Wiley and Sons, New York, 1983.

Table of Contents

1. Introduction	1
2. What's in a Color Pixel?	2
2.1. The Physics of Light Reflection from Dielectric Materials	2
2.1.1. Properties of Inhomogeneous Dielectric Materials	2
2.1.2. Reflection at the Material Surface	3
2.1.3. Reflection in the Material Body	4
2.2. The Dichromatic Reflection Model	5
2.3. Object Shape and Color Variation	6
2.4. Taking Color Images with Real Cameras: A Sensor Model	8
2.4.1. Spectral Integration	8
2.4.2. Limited Dynamic Range	9
2.4.3. Color Balancing and Spectral Linearization	10
3. Analysis of Real Color Images	13
3.1. Color Clusters from Real Color Images	13
3.2. Detecting and Removing Highlights	15
3.2.1. Determining the Body and Surface Reflection Vectors of an Image Area	15
3.2.2. Generating Intrinsic Reflection Images	18
3.2.3. Restoring the Colors of Clipped and Bloomed Pixels	19
3.3. Results and Discussion	20
4. Conclusions and Further Work	23

List of Figures

Figure 1: Light reflection of dielectric materials	2
Figure 2: Photometric angles	5
Figure 3: A mixture of the lights reflected from the material surface and body	6
Figure 4: The shape of the spectral cluster for a cylindrical object	7
Figure 5: Cluster shapes for concave objects	7
Figure 6: Color cluster in a three-dimensional color space	9
Figure 7: Color cluster in the color cube, with color clipping	9
Figure 8: Spectral responsivity of CCD-cameras and the human eye	10
Figure 9: The effect of color balancing and gamma correction on color images	11
Figure 10: Curvature in the color space, caused by gamma-corrected cameras	12
Figure 11: Spectral linearization with cubic spline functions	12
Figure 12: Three plastic cups under yellow light	13
Figure 13: Color histogram of three plastic cups under yellow light	13
Figure 14: Orange cup under yellow light	14
Figure 15: Color histogram of the orange cup under yellow light	14
Figure 16: Images and color clusters of glossy paper and a ceramic cup under white light	14
Figure 17: Fitting a convex polygon to the color cluster	16
Figure 18: Classifying the color pixels	17
Figure 19: Color histogram and fitted vectors of the three plastic cups under yellow light	17
Figure 20: Decomposing a color pixel into its constituent body and surface reflection components	19
Figure 21: Intrinsic body reflection image of the three plastic cups under yellow light; clipped pixels and bloomed pixels are shown in black	19
Figure 22: Intrinsic surface reflection image of the three plastic cups under yellow light; clipped and bloomed pixels are shown in black	19
Figure 23: Intrinsic body reflection image of the three plastic cups under yellow light; with restored clipped and bloomed pixels	20
Figure 24: Intrinsic surface reflection image of the three plastic cups under yellow light; with restored clipped and bloomed pixels	20
Figure 25: Intrinsic reflection images of the glossy paper folders under white light	20
Figure 26: Intrinsic reflection images of the ceramic cup under white light	20
Figure 27: Intrinsic reflection images of the three plastic cups under pink light	20
Figure 28: Intrinsic reflection images of the three plastic cups under white light	20

List of Tables

Table 1: Matte and highlight lines of the three plastic cups under yellow light	17
Table 2: Body and surface reflection vectors of the three plastic cups under yellow light	17
Table 3: Body and surface reflection vectors of the paper folders under white light	21
Table 4: Body and surface reflection vectors of the stripes on the ceramic cup under white light	21
Table 5: Body and surface reflection vectors of the three plastic cups under pink light	21
Table 6: Body and surface reflection vectors of the three plastic cups under white light	21