

Image Segmentation and Reflection Analysis Through Color

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

Computer Science Department, Carnegie-Mellon University, Pittsburgh PA 15213

Abstract

In this paper, we present an approach to color image understanding that can be used to segment and analyze surfaces with color variations due to highlights and shading. We begin with a theory that relates the reflected light from dielectric materials, such as plastic, to fundamental physical reflection processes, and describes the color of the reflected light as a linear combination of the color of the light due to surface reflection (highlights) and body reflection (object color). This theory is used in an algorithm that separates a color image into two parts: an image of just the highlights, and the original image with the highlights removed. In the past, we have applied this method to hand-segmented images. The current paper shows how to perform automatic segmentation method by applying this theory in stages to identify the object and highlight colors. The result is a combination of segmentation and reflection analysis that is better than traditional heuristic segmentation methods (such as histogram thresholding), and provides important physical information about the surface geometry and material properties at the same time. We also show the importance of modeling the camera properties for this kind of quantitative analysis of color. This line of research can lead to physics-based image segmentation methods that are both more reliable and more useful than traditional segmentation methods.¹

1. Introduction

One of the key goals of computer vision is to interpret the scene 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 such a 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^{1,2,3}. 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^{4,5}, 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 color⁶. All color pixels from one object 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 extend our reflection model with a sensor model which accounts for camera properties, such as a limited dynamic range, blooming, gamma-correction, and chromatic aberration. This allows us to apply our algorithms to real images, instead of only to synthesized images.

We have previously reported how this theory can be used to separate color images into two intrinsic images, one showing the scene without highlights, and the other showing only the highlights⁷. In the past, we have applied this method to hand-segmented images. The current paper describes an automatic segmentation method that is based on the extended dichromatic reflection model. Our approach alternates between generating hypotheses about the scene from the image data and verifying whether the hypotheses fit the image. The hypotheses relate object color, shading, highlights and camera limitations to the shapes of color clusters in local image areas. By using this control structure, driven by a physical model of light reflection, we are able to incrementally identify local and global properties of the scene, such as object and illumination colors, and to use them in interpreting pixels in the images. This allows us to adapt the image interpretation process to local scene characteristics and to react differently to color and intensity changes at different places in the image. This method stands in contrast to Gershon's approach⁸, in which he begins with a traditional segmentation and follows it with a physics-based post-processing step. His approach suffers from erroneous region boundaries created by the initial segmentation, while our

¹This 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.

method uses the physical model from the beginning and therefore does not make such unrecoverable mistakes.

The result is a combination of segmentation and reflection analysis that is better than traditional heuristic segmentation methods that base their analysis on intensity or color differences or on a fixed set of user-defined features, such as intensity, hue and saturation. Traditional algorithms also cannot distinguish reliably between different edge types, such as high-light edges, material edges and shading or shadow edges, and they cannot account for camera limitations. Moreover, our method generates important physical information about the scene. This information significantly simplifies our subsequent step to separate color images into two intrinsic reflection images⁷, since the segmentation already provides all necessary information about color variation on an object. When combined with methods to interpret the intrinsic images^{9,4}, this line of research can lead to physics-based image segmentation methods that are both more reliable and more useful than traditional segmentation methods.

2. The Dichromatic Reflection Model

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. How light is reflected from an object depends on the material of the object. It is common to distinguish between conducting materials, such as metals, and dielectric (non-conducting) materials, such as plastics, paints, papers and ceramics. In the following, we will present a model that describes the reflection processes of dielectric materials.

Most dielectric materials are inhomogeneous. They consist of a medium and some embedded pigments, as shown in Figure 2-1. The pigments selectively absorb light and scatter it by reflection and refraction. 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 high-light 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.

When light hits the surface of a dielectric material, the change in refraction indices causes the light to partially reflect back into the air and to partially refract, penetrating into the material body. We refer to the

reflection process at the material surface as *surface reflection*. It generally appears as a highlight or as gloss on an object. 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 light¹⁰. When modelling surface reflection, however, it is common to make some simplifying assumptions. Assuming that little or no light absorption occurs at the material surface, we may conclude that the light reflected from the surface has the same color as the illumination. This is a characteristic feature of highlights on most dielectric materials. Depending on the smoothness of the surface, the light may be reflected in a preferred direction (mirror reflection) or it may be scattered into many directions. Several models have been developed in the physics and computer graphics communities to describe the geometric properties of light reflection from rough surfaces^{11, 12, 13}.




Figure 2-1: Reflection from a dielectric material

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 pigments scatter the light and partially or entirely absorb it at some wavelengths^{14, 15}. Some of the light finally exits from the material body back into the air. We refer to this reflection process 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.

According to the above discussion, our reflection model describes the light, $L(\lambda, i, e, g)^2$, which is reflected from an object point as a mixture of the light $L_s(\lambda, i, e, g)$ reflected at the material surface and the light $L_b(\lambda, i, e, g)$ reflected from the material body.

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

If we assume that there is only a single light source in the scene, with no inter-reflection between objects, and that highlights (surface reflection) have the same color as the illumination, we can then separate the spectral reflection properties of L_s from its geometric reflection properties. 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. In the infinite-dimensional vector space of spectral power distributions (each wavelength defines an independent dimension in this vector space^{16, 17}), the reflected light can thus be described as a linear combination of the two vectors $c_s(\lambda)$ and $c_b(\lambda)$.

Many reflection models, which have been developed in the physics and computer graphics communities^{18, 4, 12, 13} are special cases of the model described here⁶. 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.

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.




Figure 3-1: The shape of the spectral cluster for a cylindrical object

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, as we will now describe. 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 3-1 shows a sketch of a shiny cylinder. The left part of the figure displays the magnitudes of the body and surface

²_{i, e,} and _g describe the angles of the incident and emitted light and the phase angle; λ is the wavelength parameter.

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_{smax} , 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 3-1.

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 3-1. Accordingly, the second term of the Dichromatic Reflection Model equation (2) has a constant value, $m_{bH} c_b(\lambda)$, 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 3-1. More precisely, the highlight cluster looks like a slim, skewed wedge because of the small variation of the body reflection component over the highlight.

The combined spectral cluster of matte and highlight points thus looks like a skewed T. 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. We have investigated this relationship more precisely for the case of a spherical object which is illuminated by a point light source at some distance d and viewed by a camera from another point at the same distance from the object. If the

angle between the illumination and the viewing direction becomes too wide, the object point with (globally) maximal body reflection becomes invisible to the camera. For wider phase angles, the (local) maximum in body reflection, $m_{b,LocalMax}$, decreases as the phase angle increases. We have calculated the relationship between $m_{b,LocalMax}$ and the amount of body reflection under the highlight, m_{mH} , for varying illumination geometries. Under reasonable imaging conditions, m_{mH} has a value that is more than 48% of $m_{b,LocalMax}$. The highlight line thus generally starts in the upper 50% of the matte line. A more detailed analysis will be presented in¹⁹. We exploit this property as the *50%-heuristic* in our segmentation method.

4. A Camera Model

The previous sections have described light reflection in a theoretical, physical model. However, the observed pixels values are also influenced by the characteristics of the recording camera. 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. This section briefly describes how some camera characteristics influence the pixel values in real images. A more detailed analysis, including color pictures with color clusters from real images, can be found in⁷.

The Dichromatic Reflection Model describes light reflection using the continuous spectrum of light. When a sensing device such as a camera or the human eye records an image, the light is integrated over the spectrum using a small number of weighting functions such as color filters. This process of *spectral integration* sums the amount of incoming light, $L(\lambda, i, e, g)$, weighted by the spectral transmittance of the filter, $\tau_f(\lambda)$, 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 number 25, 58 and 47A), 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 i , e , and g are determined by x and y and by the position of the object relative to the camera.

Spectral integration is a linear transformation^{20, 17}. For this reason, 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 pixel values lie.

We will now briefly describe a few camera limitations that are important factors when images of scenes with highlights are taken.

- Real cameras have only a limited dynamic range to sense the brightness of the incoming light. This restricts our analysis of light reflection to a *color cube*, as shown in Figure 4-1. If the incoming light is too bright at some pixel position, the camera cannot sense and represent it adequately and the light is *clipped* in one or more color bands. Color clipping can be a problem for measuring the color of highlights or bright objects. In the color histograms, it causes the clusters to bend along the walls and edges of the color cube (see Figure 4-1). 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 position²¹. This causes *blooming* in the camera²² as a result of which adjacent pixels increase their values proportional to the magnitude of the overload. We call such neighboring pixels *bloomed color pixels*. We suspect color values in the upper 10% of the intensity scale in our images to be potentially influenced by color clipping or blooming.
- Cameras are generally much less sensitive to blue light than to red light. 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*. 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 1/2 and 1 1/2 f-stops higher than the aperture used for the red exposure. We also use a total IR suppressor (Corion

FR-400) in front of our camera to eliminate the CCD-camera's sensitivity to infrared light.

- The color pixels also depend on the camera response to incident light flux. Due to *gamma-correction*, the camera output is generally related by an inverse power law to the incident flux. It 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 LeClerc²³. 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. We use these functions to generate a look-up table relating the measured intensities in every color band to the incident flux.




Figure 4-1: Color cluster in the color cube, with color clipping

5. Color Image Segmentation

We will now describe the automatic segmentation method that we have developed based on the Dichromatic Reflection Model and the camera model described above. The goal of segmentation is to identify objects in an image, as delimited by material boundaries. Because most current color segmentation methods are based on a very simple interpretation of color changes in an image, they generally segment images not only along material boundaries but also along other lines exhibiting color or intensity variations, such as highlight and shadow boundaries, or object edges with significant shading changes. The Dichromatic Reflection Model provides a more sophisticated interpretation scheme relating the physics of light reflection to color changes in the image and in the color space. Our segmentation algorithm uses the Dichromatic Reflection Model and is thus able to distinguish significant color changes at material boundaries from insignificant ones that are due to shading changes or highlights. The algorithm generates hypotheses about skewed T's in the color cube while it analyzes local color variations in the image. It then uses the hypothesized orientations of the color clusters to decide whether a color change is consistent with the local skewed-T model or whether it constitutes a material change.

5.1. Generating Initial Hypotheses about Color Clusters

We may use either a global (top-down) or a local (bottom-up) approach to determine skewed T's in the color histogram and the associated regions in the image. A global algorithm projects the entire image into the color space and then applies some analysis technique to the color space to identify and distinguish between the various clusters. A local algorithm, on the other hand, starts out with small image areas and merges areas with consistent interpretations into larger areas, assuming that local color variation is related to only a few scene properties. The global approach has problems when several objects with very similar colors exist in the scene; a local approach, on the other hand, has the problem of distinguishing camera noise from systematic color variation.

Figure 5-1 shows a scene with eight plastic objects under white light. The scene contains a cup and two donuts at various shades of red, a cup and a donut in different yellow colors, a green cup and a green donut, and a blue donut. The upper left quarter displays the original image. The color histogram from the entire image is shown in the upper right quarter. Note that the various clusters overlap significantly. In such an image, we feel that it is harder to decide when to split a large region with overlapping color clusters than to accommodate for camera noise. Therefore, we chose to utilize a local (region growing) scheme.

We start by dividing the image into small, non-overlapping windows of a given size (typically 10x10 pixels). We project the color pixels from one window at a time into the color space and find the principal components of the color distributions, as indicated by the eigenvectors and eigenvalues of the covariance matrix of the cluster^{24, 25}. We sort the eigenvalues by decreasing size. The eigenvalues and eigenvectors determine the orientation and extent of the ellipsoid that optimally fits the data.




Figure 5-1: Initial analysis of color variations for a scene with eight plastic objects

The shape of the ellipsoids provides information that we can use to relate the local color variation to physical interpretations. The classification is based on the number of eigenvalues that are approximately zero, within the limit of pixel value noise in the image. In order to classify the ellipsoids, we compare their eigenvalues with the estimated amount of noise, σ_0 . For our camera, we use an experimentally determined estimate of $\sigma_0 = 2.5$. We base our decision for each eigenvalue on a χ^2 -test with (n^2-1) parameters, where n is the window size, and we use a confidence level of $\alpha_0 = 0.005$. We then classify each cluster according to how many eigenvalues are determined to be significantly greater than zero.

- In zero-dimensional (*pointlike*) clusters, all three eigenvalues of the window are very small and the window probably lies on a single object which is either very flat or very dark.

- One-dimensional (*linear*) clusters are clusters for which only the first eigenvalue is significantly larger than the estimated camera noise. Pixels in such a window may stem from a matte object area or from the interior of a highlight area such that they form part of a matte or highlight cluster. The pixels may also come from two point clusters if a window overlays parts of two neighboring objects.
- Two-dimensional (*planar*) clusters have large first and second eigenvalues, and the local color data thus fits a plane in the color cube. Such clusters occur at windows that either cover some matte and some highlight pixels of one object or that overlay matte pixels from two neighboring regions.
- In three-dimensional (*volumetric*) clusters, all three eigenvalues are large. Such color clusters may arise in the middle of highlights where color clipping and blooming significantly increase the noise in the pixel measurements. Volumetric color clusters also occur along material boundaries when three or more objects in different colors share a window or when a window overlays matte pixels of one object and matte and highlight pixels of another object.

The lower left quarter in Figure 5-1 shows the classification of the color clusters from the initial 10x10 windows. Pointlike clusters are displayed in black, linear clusters in dark grey, planar clusters in light grey, and volumetric clusters in white. The image shows that the classifications relate in the expected way to scene properties. Most matte object areas are covered by linear windows, while windows at material boundaries and at highlights are planar or volumetric.

Next, we merge neighboring windows that have similar color characteristics. The algorithm proceeds in row-major order. For every area, it tests for every neighbor whether the two can be merged. In order not to merge windows across material boundaries, we only merge neighboring windows, if both of them and the resulting larger window all have the same classification. We use the χ^2 -test presented above to classify the larger window. Accordingly, we combine windows with pointlike clusters into larger pointlike windows; we merge windows with linear characteristics into larger linear windows; and we merge planar windows into larger planar windows. We do not merge neighboring volumetric color clusters since there is no constraint on the resulting cluster. We continue this process until no more areas can be merged. The results are initial hypotheses about the positions and orientations of pointlike, linear, and

planar clusters in color space and their respective approximate extents in the image. The lower right quarter in Figure 5-1 presents the results of merging neighboring windows of the same class in the image of the eight plastic objects.

5.2. Exploiting Linear Hypotheses

So far, the algorithm has used a bottom-up approach to extract information about the scene from the image. In a top-down step, the algorithm now uses the generated hypotheses to segment the image. We start by using the hypothesis with the largest linear cluster. The chosen linear hypothesis provides us with a model of what color variation to expect in the associated image area. The mean value and the first eigenvector describe the position and orientation of a linear color cluster, while the second and third eigenvalues determine the extent of the color cluster perpendicular to the major direction of variation. According to the Dichromatic Reflection Model, we attribute color variation along the major axis to a physical property of the scene, such as a changing amount of body or surface reflection or a material boundary.




Figure 5-2: Resolving a color conflict between two matte color clusters

Our model assumes that color variation perpendicular to the first eigenvector is caused by random noise. We thus approximate a linear color cluster by a cylinder. The radius depends on the estimated camera noise; we generally choose a constant multiple of σ_0 (typically $4\sigma_0$). We exclude dark color pixels from our color analysis because the matte clusters merge near the dark corner of the color cube. The cylinder is thus delimited at its dark end by a sphere which is centered at the black corner. We typically use a radius for the sphere of 23, which is approximately 10% of the maximum pixel value.

We then use the linear hypothesis to locally resegment the image. We select a start pixel from the area associated with the color cluster: this pixel must have a color that is contained within the color cylinder of the current hypothesis. We then examine its four neighbors, testing whether their colors lie on the cylinder. If

so, we recursively examine the next fringe of neighbors and so on. We thus grow a region of four-connected image pixels that are consistent with the current hypothesis.

Since the matte clusters converge towards the dark corner of the color cube, there exists a potential conflict between neighboring matte areas. Our heuristic of excluding very dark pixels from the analysis eliminates the most difficult cases; still, the cylinders of neighboring clusters may intersect beyond the selected dark threshold. This depends on the cylinder radius and on the angle between the two cylinders. Thus, neighboring objects with similar colors will have conflicts even at fairly bright colors. We assign pixels with a color conflict to the cluster with the closest axis, as shown in Figure 5-2.

Since there is generally more than one object in the scene, our algorithm iterates the above steps for each image area with a large linear cluster, selecting the areas by decreasing size. It stops when the next selected area is too small (typically less than 500 pixels).




Figure 5-3: Segmentation steps of the scene with eight plastic objects

In principle, the resulting "linear" regions may be related in any of several different ways to the physical processes in the scene. As discussed in section 5.1., a linear color cluster may be a matte cluster or a highlight cluster or even a combination of two clusters across a material boundary. However, we expect that

linear color clusters from highlights and across material boundaries are generally much smaller than matte clusters. Since we use only hypotheses connected with large linear clusters, we assume in the following that all of the linear regions correspond to matte linear clusters.

The upper right quarter of Figure 5-3 shows the results of selecting and applying linear hypotheses to the image with the eight plastic objects. The region boundaries outline the matte object parts in the scene, with the material boundaries being well observed. The highlight areas on the objects have not yet been analyzed. This sometimes divides the matte pixels on an object into several matte areas, as shown on the middle cup and on the lower right donut.

5.3. Generating Planar Hypotheses

We will now extend the above generated linear hypotheses into planar hypotheses that describe dichromatic planes and skewed T's. For every linear region, we examine all its neighbors to determine a prospective highlight candidate. In order to decide whether a neighboring region comes from another object (across a material boundary) or whether it is a highlight on the same object, we test whether both linear clusters point in positive directions and whether the two clusters form a skewed T and meet in the upper 50% of the matte cluster. A similar idea was used by Gershon⁸ to identify highlight segments in a postprocessing step after an initial, traditional segmentation was performed. He tested whether the two clusters are nearly parallel or whether they intersect, suspecting that parallel clusters are neighboring matte clusters. In our method, the physical model is used during the segmentation process rather than as a post-processing step.

Each candidate cluster that we test is a group of pixels on the border of a linear region. If the cluster is in fact a highlight, the first eigenvector of the cluster's color distribution will generally be in the direction of the highlight color. We start by testing whether the first eigenvector of a highlight candidate describes a positive color change, i.e: $dR \geq 0$ and $dG \geq 0$ and $dB \geq 0$. If this is not the case, the window of the highlight candidate probably overlays a material boundary and thus partially covers pixels from the current matte region and partially covers pixels from a neighboring matte region.

Next, we exploit the T-shape of the dichromatic cluster and the 50%-heuristic. For this purpose, we determine the brightest matte point in the color cluster. In order not to select a highlight point that fell into the matte cylinder, as shown in Figure 5-4, we exploit the observation that highlight clusters always grow inwards into the color cube, due to the additive nature of body and surface reflection. We thus choose the brightest matte point only from pixels with color values that are closer to the walls of the cube

than the matte line is. Next, we determine the intersection points of the two clusters. We use their color means and first eigenvectors to describe a matte and a prospective highlight line and determines the two points on the two lines that are closest to one another. If the distance between them is larger than a multiple of the estimated camera noise (typically $4\sigma_0$), we decide that the clusters do not meet in a skewed T and the neighboring area is discarded. Similarly, if the clusters intersect in the lower 50% of the matte cluster, we suspect that we are looking at two matte clusters from neighboring objects, and we discard this neighbor.




Figure 5-4: Finding the brightest matte pixel

There may be several highlight candidates because the highlight on the object may consist of a series of small windows that could not be merged, due to color clipping and blooming in the middle of the highlight. In order to select a good representative of the entire highlight, we average the intersection points of all highlight candidates, weighted by the number of pixels in the various regions. We then select the highlight region whose intersection point is closest to the average.

Since we assume in the Dichromatic Reflection Model that highlights have the same color as the illumination, all highlight clusters are parallel to one another and to the illumination color vector. As a consequence, all dichromatic planes intersect along one line, which is the illumination vector. We use this constraint to further reduce the error estimating the orientations of the highlight clusters, combining all highlight hypotheses into a single hypothesis about the color of the illumination. Observing that any imprecision of the highlight vector orientation within the dichromatic plane is irrelevant we base our computation on the normals \mathbf{n}_i of the dichromatic planes of all objects. We then determine the intersection vector $\mathbf{i}_{i,j}$ between each pair of planes:

$$\begin{aligned} \mathbf{n}_i &= \mathbf{C}_{bi} \times \mathbf{C}_{si} \\ \mathbf{i}_{i,j} &= \mathbf{n}_i \times \mathbf{n}_j \end{aligned} \quad (5)$$

Each intersection vector $\mathbf{i}_{i,j}$ provides us with a vector describing a hypothesized illumination color. The strength of the hypothesis depends on the angle between the two dichromatic planes under consideration. If the planes intersect nearly perpendicularly, the intersection line will not vary much with small amounts of error in estimating the orientations of the planes. If, on the other hand, the planes are nearly parallel, small amounts of noise may arbitrarily influence whether and where the planes intersect. We use a weighted average of all intersection vectors to generate a hypothesis on the illumination color, in which every intersection vector contributes according to the dihedral angle enclosed by the respective dichromatic planes.

5.4. Exploiting Planar Hypotheses

The Dichromatic Reflection Model states that color variations on a single object lie in a plane in the color space. We therefore use the planar hypotheses to resegment the image, expecting that the resulting image areas cover entire objects and will not be interrupted by highlight boundaries. We select one hypothesis at a time and apply it to the image, proceeding iteratively until no more unprocessed hypotheses with sufficient support from the image exist.

The chosen matte cluster and the illumination hypothesis determine a planar hypothesis predicting all significant color variation on the associated object. We use the cross product of the illumination vector and the first eigenvector of the matte cluster to determine the normal to this dichromatic plane. We use the color mean of the matte cluster to position the plane in the color cube, and we extend the plane into a slice of fixed thickness to account for camera noise. A typical choice for the width of the slice is $4\sigma_0$ in the positive and negative direction of the normal.

The algorithm uses the chosen planar hypotheses to locally resegment the image. In principle, we start from the selected matte region and expand it, recursively including pixels at the region boundaries which are consistent with the planar hypothesis. However, the algorithm must be augmented with special provisions to handle coplanar color clusters from neighboring objects. If the illumination vector lies in the plane spanned by the matte vectors of two neighboring objects, the color clusters of the objects lie in the same dichromatic plane and thus cannot be distinguished by a mere planar region growing approach. The resulting segmentation would generally be quite counterintuitive since the matte object colors may be very different and even complementary. To avoid such segmentation errors, we exploit the previously gathered knowledge about existing matte color clusters. When the planar region process encounters pixels from a previously grown matte region other than the starting region, it only continues growing if

the pixel lies within the matte color cylinder of the starting region. We thus apply the unrestricted planar growing criterion only to pixels that have not been previously recognized as matte pixels, while we fall back to the linear region growing method when matte pixels are concerned. This reflects the observation that if several matte areas exist in an object area, separated by highlight regions, all such matte areas form a single matte cluster. We also apply the proximity heuristic described in the previous section to resolve ambiguities for color pixels at the intersection of dichromatic planes.

The lower left quarter of Figure 5-3 displays the results of segmenting the scene using the generated planar hypotheses. In comparison to the linear segmentation in the upper right quarter, the segmented image areas have grown into the highlight areas. As a result, the two matte image areas on the middle cup that were previously separated by a highlight are now united. The same occurred on the lower right donut. Due to camera limitations, not all pixels in the centers of the highlights are yet integrated into the object areas. This will be discussed and remedied in the next section.

5.5. Accounting for Camera Limitations

Unfortunately, real images generally do not comply with the Dichromatic Reflection Model. Color clipping and blooming may distort the color of pixels in and around highlights significantly. As a result, the color pixels in the centers of highlights will generally not fall into the planar slice defined for the dichromatic plane of an object area and the planar segmentation will exclude such pixels, as can be observed in the upper right quarter of Figure 5-3.

Since color information is so unreliable for these pixels, we do not use it. Instead, we use a geometric heuristic to include them into the region. As mentioned above, pixels with distorted colors generally occur in the middle of the highlight areas. Thus, starting from highlight pixels we expand the planar regions into areas that are next to highlight pixels and contain very bright pixels (i.e.: brighter than the intersection point between the matte and highlight cluster).

The lower right quarter of Figure 5-3 displays the results of segmenting the scene using the generated planar hypotheses. Nearly all pixels in the highlight centers have now been integrated into the segmented regions, and the image segments correspond very well to the objects in the scene. Very few pixels on the highlights have been excluded, due to our heuristic of only integrating very bright pixels. Any image processing method for filling holes in image segments should be able to include these pixels. The lowest of the three donuts in the upper left exhibits two interesting properties. First, a small area at its upper left edge has wrongly been assigned to the donut above it. We can see in the original image that this area is

covered by a shadow, resulting in very dark pixel values. Since the color clusters of the two donuts (yellow and red) are already merged at these color values, the assignment was based on the distance from the two cylinder axes, which happened to be smaller for the upper donut. Second, there is a small area in the lower right part of the donut that was not included into the large image segment covering the donut. The color of these pixels has been significantly altered by inter-reflection from the middle cup which reflected a part of its body color (orange) onto the (yellow) donut, thus influencing the reflected colors of the donut. Inter-reflection may also be a factor in the shadowy part in the upper left of the donut. This demonstrates the sensitivity of our algorithm to the reflection processes in the scene. Since inter-reflection is not yet a part of our reflection model, our current algorithm cannot identify it in the image and process it correctly.

6. Removing Highlights

As one application of our above method to analyze and segment color images, we can now use the gathered information about the color clusters to split every color pixel into its two reflection components. We thus generate two intrinsic images of the scene, one showing the objects as if they were completely matte, and the other showing only the highlights.

We have previously reported a method to detect and remove highlights from hand-segmented images⁷. That method projected the pixels from a selected image area into the color space and fitted a skewed T to the entire color cluster, thus determining the body and surface reflection vectors, C_b and C_s , of the area. Since our new segmentation algorithm already provides this information as a result of its analysis of local color variations, we can now skip this step. However, due to possible estimation errors in the segmentation process and due to camera problems such as blooming and chromatic aberration, the vectors may not yet fit perfectly. In order to obtain a more precise fit to the data, we retest every pixel in the segmented area and label it as a matte or highlight pixel, depending on whether it is closer to the matte line or to the highlight line, which is the illumination vector starting from the intersection point of the skewed T. We then refit the matte and highlight line to the matte and highlight pixels by determining the first eigenvectors and the color means of the clusters.

The algorithm uses the reflection vectors C_b and C_s that are defined for every segmented region to split every pixel in each region into its two reflection components. C_b , C_s , and their cross product, $C_b \times C_s$, define a new (not necessarily orthogonal) coordinate system in the color cube. This coordinate system describes every color in the cube in terms of the amounts of body reflection, surface reflection and noise ϵ , as given by the color distance from the

dichromatic plane (see figure 6-1). 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. 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 6-1: Decomposing a color pixel into its constituent reflection components

We cannot apply this method to bloomed and clipped pixels because such pixels generally do not satisfy the conditions of the Dichromatic Reflection Model: since their color is altered by the sensing process in the camera, they generally do not lie on the skewed T or even in the dichromatic plane. In order to restore the physically correct color of such 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 pixel by a pixel on the matte or highlight line that has the same value in the undistorted band.

The upper quarters in Figure 6-2 display the resulting intrinsic images of the scene with the eight plastic objects. The images demonstrate that we are able to determine the body and surface reflection com-

ponents of the various objects reasonably well. The orientations of the corresponding reflection vectors are shown in Table 9-1. When evaluated under eyeball inspection, the body reflection image seems to generally provide a smooth shading component across the highlight area. We expect that this image may therefore be a useful tool to determine object shapes from shading information. We will investigate this application in the future. There exist thin dark rings in the body reflection image around the highlight on the lower right donut. This error is due to chromatic aberration in the camera which currently limits the performance of our algorithm.

The surface reflection image exhibits very well the highlight components on the objects. All highlights have been detected, and beyond that, the image provides precise quantitative information on the varying amount of surface reflection. Notice the gradient in the surface reflection component on the rightmost cup and the small amount of gloss that was detected on the handle of the middle cup. A careful inspection of the surface reflection image reveals that the surface reflection component also increases at the material boundaries. This effect is related to aliasing occurring at pixels that integrate light from two neighboring objects. The colors of such pixels are a linear combination of two matte object colors. Depending on the orientations of the dichromatic planes, they may be included in either of the two object areas in the image, or they may be left unassigned, as at the edge between the middle and the rightmost cup. If they are assigned to one of the object areas, they generally do not lie close to the matte line, thus resulting in a higher surface reflection component.




Figure 6-2: Intrinsic reflection images of the scene with eight plastic objects

7. Results and Discussion

Figure 7-1 shows the results of applying our segmentation algorithm and the reflection analysis to a scene of an orange, a green and a yellow cup under white light. These same images were analyzed in our earlier work, using hand-segmentation⁷. The seg-

mented image in the upper right quarter demonstrates that our method is capable of correctly identifying the object areas. Due to its modeling of matte and highlight color variations, our program ignores color changes along the highlight boundaries. The lower quarters in the figure display the intrinsic body and surface reflection images. The orientations of the matte and highlight vectors are given in Table 9-2. Table 9-3 shows the vector orientations of the cups under yellow light. The intrinsic images and the reflection vectors are comparable or better than the ones that we obtained earlier by fitting skewed Ts to the entire histogram of hand-segmented image areas ⁷.




Figure 7-1: Analysis results for three plastic cups under white light

Our examples show that our algorithm analyzes and segments the given images quite well. We will now discuss the conceptual and implementational limitations of our approach. First, our current implementation depends on the window size used for determining the initial local color variation in the image. If the windows are too large, many of them overlap several objects, and the clusters are volumetric. If the windows are too small, color variation on a relatively flat or dark object may be overlooked, as the result of pointlike clusters. Second, our cylindrical model for the linear hypotheses does not account for secondary effects on pixel values, such as color clipping, blooming, chromatic aberration, aliasing and inter-reflection between objects. Because such effects influence the shape of the color cluster, a

non-circular model of the cross-section of linear clusters may be needed. We deal with this problem by choosing a fairly large cylinder radius. It may be desirable to adaptively fit a tighter model to the linear cluster. Along the same lines, the hypothesized orientation of the linear cluster may be inexact and could be improved by an adaptive method that refits the cylinder until it optimally fits a given cluster. We are currently investigating these issues.

The conceptual limitations of our approach are related to the basic principles of our model. Since we attribute any linear color variation to the changing illumination geometry of a single material, we are unable to find material boundaries between objects with collinear matte clusters. We will need a geometrical analysis, linking intensity gradients to object shape, to distinguish between such objects. The same will be needed to analyze dark image areas which are currently excluded because their color information is too unreliable.

The model also makes simplifying assumptions about the illumination conditions and the materials in the scene. A color cluster from a single object in an unconstrained scene will generally not be a skewed T composed of linear subclusters because the illumination color may vary on different parts on the object surface, and the reflection properties of the object may also change, due to pigment variations in the material body. The necessary extensions to the model will be the subject of future work.

Furthermore, our method to split color pixels into their reflection components (but not the segmentation) relies on a characteristic color change between the matte object color and the highlight color. There needs to be a certain angle between the orientations of the body and surface reflection vectors of an object. How big the angle needs to be depends on the estimated camera noise (as related to the width of the matte cylinder). If the matte and highlight clusters are approximately collinear, we cannot separate the reflection components. Similarly, we have problems when one of the two linear clusters does not exist or is very small for an object. The matte cluster is missing, if the viewed object is very dark, or if the scene is illuminated with a narrow-band illuminant that does not overlap with the wavelengths at which the material reflects light. Matte clusters also do not exist for metallic objects. On the other hand, the highlight cluster may be missing, if an object does not reflect a highlight into the camera, due to its position in the scene and the illumination geometry. As a third case, we need to consider objects with very rough surfaces such that every pixel in the image area has both a significant body and surface reflection component. The color cluster may then fill out the entire dichromatic plane. A common special case of this are so-called "matte" or "lambertian" materials - as opposed to glossy materials - which reflect a constant

amount of surface reflection in every direction and thus never exhibit a highlight in the common sense of the word. The corresponding color clusters are linear clusters, translated from the origin of the color space according to the constant surface reflection component. Our current method is not capable of distinguishing between all these cases that result in a single, linear cluster. In combination with exploiting previously determined scene properties, such as the illumination color, we will need to analyze the intensity gradients along the linear axes and relate them to the properties of m_s and m_b , as described in a geometrical model of light reflection.

8. Conclusions

In this paper, we have demonstrated that it is possible to analyze and segment real color images by using a physics-based 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 automatically segment an image while generating hypotheses about the scene. We then use the knowledge we have gained 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^{26, 9, 4, 27}. 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^{28, 29, 30, 31}.

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^{28, 8, 30, 5}, previous work on image segmentation and highlight detection 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 distinguish highlight boundaries from material boundaries and 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 hypothesis-based approach towards image segmentation may provide a new paradigm for low-level image understanding. Our method gains its strength from using an intrinsic model of physical processes that occur in the scene. The result are intrinsic images and hypotheses which are closely related in their interpretation to the intrinsic model, being instantiations of concepts formulated in the model. Our system alternates between a bottom-up step which generates hypotheses and a top-down step which applies the hypotheses to the images. Our analysis thus consists of many small, complete interpretation cycles that combine bottom-up processing with feed-back in top-down processing. This approach stands in contrast to traditional image segmentation methods which do not relate their analysis to intrinsic models and that also generally have a strictly bottom-up control structure. We feel that many low-level image understanding methods such as shape-from-x methods, stereo and motion analysis may be viewed and approached under this paradigm. We hope to extend our approach into a more complete low-level image analysis system which combines color analysis with a geometrical analysis of the scene, exploiting the body and surface reflection images. Along these lines, we may generate hypotheses about object shapes and about the object materials²⁹. The highlight image may also provide strong evidence for the position of the light source.

Although the current method has only been applied in a laboratory setting, its 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.

9. References

1. J.R. Kender, "Instabilities in Color Transformations", *Pattern Recognition and Image Processing*, Vol. PRIP-77, June 1977, pp. 266-274, IEEE Computer Society, Troy NY

2. R. Ohlander, K. Price, and D.R. Reddy, "Picture Segmentation Using a Recursive Region Splitting Method", *Computer Graphics and Image Processing*, Vol. 8, 1978, pp. 313-333.
3. Y. Ohta, T. Kanade, and T. Sakai, "Color Information for Region Segmentation", *Computer Graphics and Image Processing*, Vol. 13, 1980, pp. 222-231.
4. B.K.P. Horn, "Understanding Image Intensities", *Artificial Intelligence*, Vol. 8, No. 11, 1977, pp. 201-231.
5. J.M. Rubin and W.A. Richards, "Color Vision and Image Intensities: When are Changes Material?", *Biological Cybernetics*, Vol. 45, 1982, pp. 215-226.
6. S.A. Shafer, "Using Color to Separate Reflection Components", *COLOR research and application*, Vol. 10, No. 4, Winter 1985, pp. 210-218, Also available as technical report TR-136, Computer Science Department, University of Rochester, NY, April 1984
7. G.J. Klinker, S.A. Shafer, and T. Kanade, "The Measurement of Highlights in Color Images", to appear in the *International Journal on Computer Vision (IJCV)*, 1988.
8. R. Gershon, *The Use of Color in Computational Vision*, PhD dissertation, Department of Computer Science, University of Toronto, 1987.
9. G. Healey and T.O. Binford, "Local Shape from Specularity", *Proceedings of the First International Conference on Computer Vision (ICCV)*, IEEE, London, June 1987, pp. 151-161.
10. P. Moon, "A Table of Fresnel Reflection", *J. Math. Phys.*, Vol. 19, No. 1, 1940.
11. P. Beckmann and A. Spizzichino, *The Scattering of Electromagnetic Waves from Rough Surfaces*, MacMillan, New York, Vol. 15, 1963, Pages 1-33, 70-96
12. B.T. Phong, "Illumination for Computer Generated Pictures", *Communications of the ACM*, Vol. 18, 1975, pp. 311-317.
13. K.E. Torrance, and E.M. Sparrow, "Theory for Off-Specular Reflection from Roughened Surfaces", *Journal of the Optical Society of America*, Vol. 57, September 1967, pp. 1105-1114.
14. R.S. Hunter, *The Measurement of Appearance*, John Wiley and Sons, New York, 1975.
15. D.B. Judd and G. Wyszecki, *Color in Business, Science and Industry*, John Wiley and Sons, New York, 1975.
16. B.K.P. Horn, "Exact Reproduction of Colored Images", *Computer Vision, Graphics, and Image Processing (CVGIP)*, Vol. 26, 1984, pp. 135-167.
17. S.A. Shafer, "Describing Light Mixtures Through Linear Algebra", *Journal of the Optical Society of America (JOSA)*, Vol. 72, No. 2, February 1982, pp. 299-300.
18. R.L. Cook and K.E. Torrance, "A Reflectance Model for Computer Graphics", *ACM Transactions on Graphics*, Vol. 1, No. 1, January 1982, pp. 7-24, Also published in *Computer Graphics 15(3)*, SIGGRAPH'81
19. G.J. Klinker, *A Physical Approach to Color Image Understanding*, PhD dissertation, Computer Science Department, Carnegie-Mellon University, Summer 1988, In preparation
20. H. Grassmann, "On the Theory of Compound Colors", *Phil. Mag.*, April 1854.
21. W. Budde, *Physical Detectors of Optical Radiation*, Academic Press, New York, Optical Radiation Measurements, Vol. 4, 1983.
22. J.D.E. Beynon and D.R. Lamb (eds.), *Charge-coupled devices and their application*, McGraw-Hill, London, 1980.
23. Y. LeClerc, "A Method for Spectral Linearization", Private communication
24. D.H. Ballard, and C.M. Brown, *Computer Vision*, Prentice Hall, Inc., Englewood Cliffs, NJ 07632, 1982.
25. H.H. Harman, *Modern Factor Analysis, second edition*, University of Chicago Press, Chicago and London, 1967.
26. 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*, Vol. 20, 1982, pp. 199-228.
27. C.E. Thorpe, *FIDO: Vision and Navigation for a Robot Rover*, PhD dissertation, Computer Science Department, Carnegie-Mellon University, December 1984, Available as technical report CMU-CS-84-168
28. M. D'Zmura and P. Lennie, "Mechanisms of color constancy", *Journal of the Optical Society of America A (JOSA-A)*, Vol. 3, No. 10, October 1986, pp. 1662-1672.

29. G. Healey and T.O. Binford, "The Role and Use of Color in a General Vision System", *DARPA-Image Understanding (IUS) workshop*, L.S. Bauman, ed., Los Angeles, CA, February 1987, pp. 599-613.
30. H.-C. Lee, "Method for computing the scene-illuminant chromaticity from specular highlights", *Journal of the Optical Society of America A (JOSA-A)*, Vol. 3, No. 10, October 1986, pp. 1694-1699.
31. 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)*, Vol. 3, No. 1, January 1986, pp. 29-33.

Reflection Vectors: Eight Plastic Objects under White Light		
	body reflection vector	surface reflection vector
dark red donut	(0.99,0.11,0.12)	(0.62,0.52,0.59)
yellow cup	(0.84,0.52,0.14)	(0.48,0.62,0.62)
green cup (right half)	(0.27,0.89,0.37)	(0.54,0.58,0.61)
red cup	(0.95,0.26,0.14)	(0.68,0.48,0.56)
yellow donut	(0.77,0.61,0.19)	(0.69,0.51,0.51)
bright red donut	(0.98,0.18,0.12)	(0.77,0.61,0.18)
green donut	(0.21,0.78,0.59)	(0.47,0.63,0.31)
green cup (left half)	(0.27,0.86,0.42)	(0.61,0.52,0.59)
blue donut	(0.00,0.25,0.97)	(0.49,0.63,0.60)
illumination vector: - computed by alg. - independent meas.		(0.57,0.58,0.57) (0.58,0.57,0.58)

Table 9-1: Body and surface reflection vectors of the eight plastic objects under white light

Reflection Vectors: Plastic Cups under White Light		
	body reflection vector	surface reflection vector
green cup	(0.22,0.91,0.37)	(0.46,0.51,0.72)
yellow cup	(0.81,0.57,0.15)	(0.56,0.56,0.61)
orange cup	(0.95,0.26,0.15)	(0.47,0.58,0.66)
illumination vector: - computed by alg. - independent meas.		(0.69,0.62,0.38) (0.58,0.57,0.58)

Table 9-2: Body and surface reflection vectors of the three plastic cups under white light

Reflection Vectors: Plastic Cups under Yellow Light		
	body reflection vector	surface reflection vector
green cup	(0.49,0.86,0.16)	(0.76,0.62,0.20)
yellow cup	(0.85,0.52,0.08)	(0.74,0.65,0.14)
orange cup	(0.97,0.24,0.08)	(0.73,0.66,0.18)
illumination vector: - computed by alg. - independent meas.		(0.81,0.56,0.19) (0.73,0.66,0.19)

Table 9-3: Body and surface reflection vectors of the three plastic cups under yellow light