

USING A COLOR REFLECTION MODEL TO SEPARATE HIGHLIGHTS FROM OBJECT COLOR

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

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

ABSTRACT Current methods for image segmentation are confused by artifacts such as highlights, because they are not based on any physical model of these phenomena. In this paper, we present an approach to color image understanding that accounts for color variations due to highlights and shading. Based on the physics of reflection by dielectric materials, such as plastic, we show that the color of every pixel from an object can be described as a linear combination of the object color and the highlight color. According to this model, all color pixels from one object form a planar cluster in the color space whose shape is determined by the object and highlight colors and by the object shape and illumination geometry. We present a method which exploits the color difference between object color and highlight color, as exhibited in the cluster shape, 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 that cannot detect or analyze highlights, such as stereo vision, motion analysis, shape from shading, and shape from highlights. We have applied this method to real images in a laboratory environment, and we show these results and discuss some of the pragmatic issues endemic to precision color imaging.

I. Introduction

When we look at an image, we can interpret what we see 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-aggregate properties to these scene characteristics. In the past, most work with color images has considered object color to be a constant property of an object. Color variation on an object was attributed to noise.¹ However, in real scenes, color variation depends to a much larger degree on the optical reflection properties of the scene, which cause the perception of object color, highlights, shadows and shading.^{2,3} As a consequence, color variation cannot be regarded to be of a merely statistical nature. On the contrary, it

exhibits characteristics that can be determined and used for color vision.

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 the color of every pixel from an object as a linear combination of the object color and the highlight color.⁴ According to our model, 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 present a method that exploits the color difference between object color and highlight color, as exhibited in the cluster shape, to separate the color of every pixel into a matte component and a highlight component. Our method generates two intrinsic images, one showing the scene without highlights, and the other one showing only the highlights. These intrinsic images can be a useful tool for a variety of algorithms in computer vision that cannot detect or analyze highlights, such as stereo vision, motion analysis, shape from shading and shape from highlights.^{2, 5, 6} We demonstrate the applicability of our method to real color images.

We begin by introducing the dichromatic reflection model which describes the interaction of light with opaque dielectric materials. Using this model, we describe the color variation on an object as a function of the reflection properties of the materials, object shapes and sensor characteristics. Next, we demonstrate how our reflection model can be used for the analysis of color images. We present methods to determine the illumination color and to detect and remove highlights from real color images. Finally, we discuss the implications and the future direction of our work.

II. The dichromatic reflection model

When we look at a glossy object, we usually see the reflected light as composed of two colors that typify the highlight areas and the matte object parts. The *dichromatic reflection model* describes this phenomenon for scene configurations in which a single light source of an arbitrary color illuminates opaque dielectric materials.⁴ As we will show later in this paper, this model accounts well for real color data under suitable conditions.

..mat.pre..

Figure 1: Light reflection at dielectric materials

Figure 1 illustrates the interaction of light with inhomogeneous, dielectric materials, such as plastic and paint. The material interface immediately reflects some percentage of the light, according to Fresnel's law of reflection. In general, the reflected light has approximately the same color as the light source. The remaining percentage of the incident light penetrates into the material body, which then scatters the light and absorbs it at some wavelengths, before it reemits the rest.^{7, 8, 9} The color of this light is determined by the illumination color and the reflection properties of the material. Common names for the reflection process at the interface are the terms *specular reflection*, *highlight* or *gloss*, whereas the reflection

process in the material body is generally called *diffuse reflection* or *matte color*. We refer to the two processes as *interface* and *body reflection* because those terms make a more precise, physical distinction between the reflection processes than the geometry-based terms specular and diffuse reflection. We will use this terminology and the corresponding terms highlights and matte object color throughout the paper.

Based on the above discussion, the dichromatic reflection model describes the spectral radiance $L(\lambda, i, e, g)^*$ of a point in the scene as a sum of an interface reflection component $L_i(\lambda, i, e, g)$ and a body reflection component $L_b(\lambda, i, e, g)$. The model assumes that the spectral properties of illumination and reflection on an object are independent of the orientation of the surface, which is a reasonable approximation.⁴ We thus decompose each of the two reflection components into a spectral *composition* $c(\lambda)$ that describes a color, and a *magnitude*, $m(i, e, g) \in [0, 1]$, that describes a geometric scale factor:

$$L(\lambda, i, e, g) = m_i(i, e, g) c_i(\lambda) + m_b(i, e, g) c_b(\lambda) \quad (1)$$

When a color TV camera records an image of a scene, it generally uses three color primaries to represent the spectrum of the light that is reflected from the objects towards the camera. In the image formation process, the camera transforms the light of the incoming ray at pixel position (x, y) via tristimulus integration from an infinite light spectrum into a triple of color values, $\mathbf{C}(x, y) = [r, g, b]$. This process sums the amount of light at each wavelength weighted by the transmittance of the color filters and the responsivity of the camera. Because this is a linear transformation, and because the photometric angles, i , e , and g depend on x and y , the dichromatic reflection model can be applied to color pixel values. This allows us to describe the color pixel value $\mathbf{C}(x, y)$ as a linear combination of the vectors representing the colors of interface reflection \mathbf{C}_i and body reflection \mathbf{C}_b at the corresponding point in the scene:

$$\mathbf{C}(x, y) = m_i(i, e, g) \mathbf{C}_i + m_b(i, e, g) \mathbf{C}_b \quad (2)$$

This equation describes how the pixel color $\mathbf{C}(x, y)$ of an arbitrary pixel depends on the optical properties of the scene. We will now discuss the relationship between the colors of all pixels on an object.

III. Color variation and object shape

As a means to model the color variation over an entire object, we use a color histogram in the color space, which is the projection of the colors of all pixels from the object into the color space.

* i , e , and g describe the angles of light incidence and exitance and the phase angle with respect to the surface normal, λ is the wavelength parameter

The dichromatic reflection model assumes that there is a single light source in the scene, without ambient light or inter-reflection between objects. Under this assumption, the colors of all pixels from an object are linear combinations of the same interface and body reflection colors \mathbf{C}_i and \mathbf{C}_b . Color variation within an object area thus depends only on the geometric scale factors m_i and m_b while \mathbf{C}_i and \mathbf{C}_b are constant. Accordingly, \mathbf{C}_i and \mathbf{C}_b span a *dichromatic plane* in the color space, and the colors of all pixels from one object lie in this plane.

Within the dichromatic plane, the color pixels form a dense cluster. There exists a close relationship between the shape of such a color cluster and the geometric properties of body and interface reflection and the shape of the object. We can use this relationship to determine characteristic features of the color clusters. As an aid to intuition, we will assume that body reflection is approximately Lambertian and that interface reflection is describable by a function with a sharp peak around the angle of perfect mirror reflection. This is a simplified view of the reflection processes in real scenes. However, as we will demonstrate, it is sufficient for our analysis of color images.

..cya.pre..

Figure 2: The shape of the color cluster for a cylindrical object

Figure 2 shows a sketch of a shiny cylinder. The left part of the figure displays the magnitudes of the body and interface components as curves showing the loci of constant body or interface reflection. The right part of the figure shows the corresponding color cluster in the dichromatic plane. This represents the configuration observed in a histogram of pixel values in the color space. To relate the terminology of the dichromatic reflection model to the shape of the color clusters, we classify the color pixels as *matte pixels*, *highlight pixels* or *clipped color pixels*. The following paragraphs discuss the characteristic features of each of these classes.

Matte pixels are projections of points in the scene that exhibit only body reflection in the direction of the viewer. The color of such pixels is thus determined by the color of body reflection, scaled according to the geometrical relationship between the local surface normal of the object and the illumination direction. Consequently, the colors of the matte pixels form a *matte line* in the color space, in the direction of the body reflection vector \mathbf{C}_b .

Highlight pixels are projections of scene points that exhibit both body reflection and interface reflection in the viewing direction. The colors of all pixels in a highlight area that lie on a line of constant body reflection m_{bH} vary only in their respective amounts of interface reflection. The colors of these pixels thus form a straight *highlight line* in the color space which is parallel to the interface reflection vector \mathbf{C}_i . The line starts from the matte cluster at the position representing the body reflection component m_{bH} of the highlight pixels. Combined with the neighboring highlight pixels of slightly different amounts of body reflection, all highlight pixels form a highlight cluster in the color space that has the shape of a skewed wedge. If more than one highlight exists on an object, each of them describes a highlight line in the color space (see Figure 3).

The combined color cluster of matte and highlight pixels thus looks

..gea.pre..

Figure 3: Color cluster shapes under changing illumination geometry

like a skewed T or comb. The exact shape of the cluster depends 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 very close to the surface normal. According to Lambert's law, the body reflection component is maximal at such object points and thus, the starting point m_{bH} for the highlight line is at the tip of the matte line. The wider g becomes at a highlight, the smaller is the amount of underlying body reflection and, thus, the greater is the distance between the tip of the matte line and the starting point m_{bH} of the highlight line (see Figure 3). At the same time, small positional variations on the object (as for neighboring pixels) become less influential on the incidence and exitance angles i and e . The highlight thus becomes dimmer and spreads out over a larger area, covering a larger range of values of underlying body reflection. As a result, the highlight line grows wider, exhibiting more strongly the shape of a wedge.

Clipped color pixels are highlight pixels at which the light reflection exceeds the dynamic range of the camera. Depending on the color of the object, the dynamic range may be exceeded at some points in one color band but not in the other two, and the highlight cluster then bends near the wall of the color cube that describes the limit of sensitivity of that color band. A second color band may saturate at some brighter object points, causing the color cluster to bend again at the edge of the limiting walls of both color bands. At the innermost points of the highlight the dynamic range of the camera may even be exceeded in all three color bands, thus looking white, even though the color of illumination may not be white.

..clu.pre..

Figure 4: Shape of color clusters in color space

Summarizing this discussion, the general shape of a color cluster is displayed in Figure 4. Although the color cluster of a specific object may not fill the entire parallelogram, it reliably exhibits some features that can thus be used and searched for by algorithms. Figure 4 displays these features as bold lines. Color clusters generally provide a *matte line* on which all matte pixels of an object area lie. Depending on the object shape and the illumination geometry, color clusters have some number of highlight lines. We use the *brightest highlight line* as a representative of all these lines. Further lines connected to the highlight lines are *clipped color lines*. An algorithm can analyze a color cluster by searching for these lines in the color space. They determine the general shape of the parallelogram and the orientation of the dichromatic plane.

IV. Analysis of real color images

There exist some inherent assumptions in the dichromatic reflection model.⁴ However, as we will illustrate, the model accounts well for real color data under suitable conditions. We also present methods for color image analysis that use the shape of the color clusters.

A) Color clusters from real color images

We have taken a series of color images in the Calibrated Imaging Laboratory at Carnegie-Mellon. The scene consists of an orange, a green and a yellow plastic cup under white or yellow illumination. Black curtains on the walls were used to eliminate ambient light in the visible spectrum of the scene. We use a spectral linearization method to compensate for the non-linear response of our camera to image brightness. Figure 5a shows the linearized image of the orange, the yellow and the green cup under yellow illumination.

- a) linearized image (cups under yellow light)
- b) color histogram (cups under yellow light)
- c) fitted dichromatic planes (cups under yellow light)

Figure 5: Color clusters and dichromatic planes of cups under yellow light**

We use a graphical display program which takes a color picture as input and displays the colors of all pixels of selected object areas as a color histogram in the color space. Figure 5b displays the color histogram of the pixels from the marked areas of Figure 5a in the color space. The color space is shown as a color cube, with each dimension of the cube representing the intensity scale of one of the three color primaries. Its origin is at the black corner of the cube.

The color clusters of the cups each lie approximately in dichromatic planes. Within the planes, they form matte and highlight lines, thus demonstrating that real color data follows the theory of the dichromatic reflection model. The color clusters also have clipping lines, due to the bright intensity of the yellow illumination, as reflected from the object in the middle of the highlights. Accordingly, the colors in the middle of the highlight areas look white, whereas the pixels closer to the highlight boundaries exhibit yellow interface reflection (in addition to the underlying amounts of the respective body reflection colors).

B) Determining the color of illumination

If several glossy objects of different color are illuminated by the same light source, each object produces a dichromatic plane. Because all of these dichromatic planes contain the same interface reflection vector \mathbf{C}_i , they intersect along a single line which is the color of the illumination. This fact is demonstrated in Figure 5c. Color constancy algorithms^{11, 12} can use this fact to remove the influence of the illumination color from the body reflection component, thus "normalizing" the image to a standard white illumination.

C) Detecting and removing highlights

We have developed and implemented an algorithm that uses the shape of the color clusters to detect and remove highlights from color images. The program projects the pixels of selected image areas into

**A set of color photographs is available from the authors.¹⁰

the color space and fits a dichromatic plane to the color data from each image area. The program then searches within each dichromatic plane for the matte line, the brightest highlight line and lines of clipped colors. These lines are extracted from the dichromatic plane by using a recursive line splitting algorithm,¹³ and classified as the *matte vector*, the *highlight vector* and *clipped color vectors*. The program assumes that the line starting closest to the black corner is the matte vector. The next one connected to it is classified as the highlight vector. The remaining lines are assumed to describe clipped color data.

Each color pixel of the image is then broken up into its reflection components. In order to remove interface reflection, the program projects the color of every pixel onto the respective dichromatic plane. If the projected color is close to a clipped color vector, it is replaced by the color at the end of the highlight vector. The algorithm then projects the color of every pixel along the highlight vector onto the matte line. The result is the intrinsic *matte image* of the scene. Conversely, the program generates the intrinsic *highlight image* of a scene by projecting the color of every pixel along the matte vector onto a line that is parallel to the highlight vector but goes through the origin of the color space (i.e.: onto the interface reflection vector \mathbf{C}_i).

The program has been applied to the pictures of an orange, a yellow or a green plastic cup under white or yellow light. Figure 6 shows the results we obtained from running the algorithm on an image of the orange cup under white light. Figures 6a and 6b show the image of the cup and the color cluster that is generated by projecting the pixels from the marked object area into the color space. Figures 6c and 6d display the resulting intrinsic images of our algorithm. Figure 6c shows only the highlight, and Figure 6d shows the object without the highlight. The results of applying the algorithm to the other pictures are similar (see figure 7 for yellow illumination on the orange cup).

- a) linearized image (orange cup under white light)
- b) color histogram (orange cup under white light)
- c) intrinsic highlight image (orange cup under white light)
- d) intrinsic matte image (orange cup under white light)

Figure 6: Intrinsic images of the orange cup under white illumination

..photo..

Figure 7: Intrinsic images of the orange cup under yellow illumination

A qualitative inspection of the intrinsic images reveals that our algorithm is able to separate the highlights from the body reflection component. Note that the highlight image displays both the highlight from the middle of the cup and the small amount of gloss that is reflected from the handle of the cup. However, the current algorithm does not yet detect the interface reflection color reliably. The reason for this is the following: the scene of Figure 6 was illuminated by white light. In the color cube, the white light corresponds to the diagonal direction, from the black corner to the white corner. Since the highlight line of this image starts with an already high amount of

mainly red body reflection, the corresponding diagonal direction for the interface reflection component is hard to detect and, as a consequence, the clipping vector could not be distinguished from the highlight vector in this image. For this reason, the computed highlight vector misses the red component and the highlight image looks cyan. The algorithm can be improved by using the intersection of several dichromatic planes as an indication of the highlight vector (see figure 5).

The potential utility of the intrinsic images stems from the fact that both the matte and the highlight image have simpler geometric properties than does intensity in a monochrome image, that represents a weighted sum of the two. A subsequent process to infer object shapes from the geometric properties of highlights or shading can thus be performed more easily on the intrinsic images than on the original image. Furthermore, the intrinsic matte image is relatively insensitive to changes in viewpoint, as occurring in an image sequence or with stereo image pairs. Accordingly, the detected image features (lines or interesting points) from matte images are more likely to be at a fixed position on the object than image features taken from the original image, some of which might have been caused or disturbed by highlights. This is important for stereo vision and motion analysis.^{5, 6}

V. 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 the limited dynamic range 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 two intrinsic images are promising for improving the results of many other computer vision algorithms that cannot detect or analyze highlights. To demonstrate this, we plan to combine our approach with a photometric stereo system to guide a robot arm.¹⁴

The key point leading to the success of this work is our modeling of highlights as a linear combination of both body and interface reflection. In contrast, previous work on highlight detection in images has generally assumed that the color of the pixels within a highlight 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 the current binary methods based on thresholding intensity.

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 (material) form a single cluster in the color space. A color cluster thus groups pixels into areas of

constant material properties, independently of illumination geometry influences, such as shading or highlights. This finding can be used by color segmentation algorithms 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.

Our approach to color image understanding by analyzing light reflection in color histograms may also be applicable to more complex illumination conditions, including ambient light, light sources in several colors, shadow casting and inter-reflection between objects. If a constant amount of ambient light illuminates the scene from all directions, a constant color vector describing its color must be added to the reflection model.¹⁵ This induces a translational color change in the reflected light at all object points, and the color cluster moves accordingly in the color space.⁴ Inter-reflection between objects is illustrated in figure 8. In this case, the color of the illuminating light at various places on an object changes due to the changing amount and color of light reflected from neighboring objects towards the object. The same occurs if several light sources in different colors exist in the scene. The changing illumination color affects the colors of both body and interface reflection from the objects. The color clusters in the color space reflect these dependencies as additional segments, attached to the previously described cluster parts. As a result of the changing body reflection color, the color clusters of neighboring objects have additional segments that are directed towards each other. These segments are not necessarily linear. When combined, the body reflection segments from clusters of neighboring objects exhibit the shape of a W, as demonstrated in the upper right box of Figure 8.

blankspace (4.0 inches)

Figure 8: Inter-reflection between the rings of a toy

We are also considering extensions of the model to account for non-dielectric materials such as metals. When light hits the surface of a metal, it is immediately absorbed at some wavelengths, according to the absorption properties of the metal, and the remaining light is immediately reflected back into the surrounding medium. Thus, light reflection from metals is governed by only one process which takes place at the object interface. In this case, the reflection model thus consists of a single term, and the light reflected from the object forms a linear cluster in the color space. The direction of the line depends on the absorption properties of the metal. The distribution of points on the line is governed by the laws of mirror reflection and scattering on rough surfaces.¹⁶ Similar findings have been described by Healey and Binford.¹⁷ However, their approach relies on spatial as well as spectral properties of light reflection whereas ours is purely spectral, describing the reflection properties along the entire object globally in a histogram.

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 signal-based methods for image segmentation.

Acknowledgements

We would like to thank Ruth Johnston-Feller for her comments and suggestions and Georg Klinker and Richard Szeliski for commenting previous drafts of this paper.

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 or the Defense Advanced Research Projects Agency or the US Government.

REFERENCES

1. Y. Ohta, T. Kanade, and T. Sakai, "Color Information for Region Segmentation", *Computer Graphics and Image Processing*, Vol. 13 1980, pp. 222-231.
2. B.K.P. Horn, "Understanding Image Intensities", *Artificial Intelligence*, Vol. 8 No. 11 1977, pp. 201-231.
3. J.M. Rubin and W.A. Richards, "Color Vision and Image Intensities: When are Changes Material ?", *Biological Cybernetics*, Vol. 45 1982, pp. 215-226.
4. 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, University of Rochester, NY, April 1984
5. 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.
6. 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
7. R.S. Hunter, *The Measurement of Appearance*, John Wiley and Sons, New York, 1975.
8. D.B. Judd and G. Wyszecki, *Color in Business, Science and Industry*, John Wiley and Sons, New York, 1975.
9. S.J. Williamson and H.Z. Cummins, *Light and Color in Nature and Art*, John Wiley and Sons, New York, 1983.
10. G. Klinker, S. Shafer, and T. Kanade, "Photo Set 1: Using a Color Reflection Model to Separate Highlights from Object Color", Computer Science Department, Carnegie-Mellon, Schenley Park, Pittsburgh PA 15213, USA, available upon request from the authors

11. M. D'Zmura and P. Lennie, "Mechanisms of color constancy", *Journal of the Optical Society of America A (JOSA-A)*, Vol. 3No. 10October 1986, pp. 1662-1672.
12. H.-C. Lee, "Method for computing the scene-illuminant chromaticity from specular highlights", *Journal of the Optical Society of America A (JOSA-A)*, Vol. 3No. 10October 1986, pp. 1694-1699.
13. T. Pavlidis, *Structural Pattern Recognition*, Springer Verlag, Berlin, Heidelberg, New York, 1977.
14. K. Ikeuchi, "Generating an Interpretation Tree From a CAD Model to Represent Object Configurations For Bin-Picking Tasks", Tech. report CMU-CS-86-144, Department of Computer Science, Carnegie-Mellon University, August 1986.
15. R.L. Cook and K.E. Torrance, "A Reflectance Model for Computer Graphics", *Computer Graphics*, Vol. 15No. 3August 1981, pp. 307-316.
16. K.E. Torrance, and E.M. Sparrow, "Theory for Off-Specular Reflection from Roughened Surfaces", *Journal of the Optical Society of America*, Vol. 57September 1967, pp. 1105-1114.
17. G. Healey and T.O. Binford, "Local Shape from Specularity", *DARPA-Image Understanding (IUS) workshop*, L.S. Bauman, ed., DARPA, Los Angeles, CA, February 1987, pp. 874-887.