Anatomy of a Color Histogram

Carol L. Novak and Steven A. Shafer

School of Computer Science, Carnegie Mellon University, Pittsburgh, PA 15213

Abstract

One of the key tools in physics-based vision has been color histogram analysis. But to date histograms have only been used for pixel grouping, color analysis, and material type labeling. In this paper we present a new, quantitative model of histograms that yields a more complete description of scene properties.

In the mid-1980s it was recognized that the color variation for inhomogeneous surfaces may be modeled as a regular physical process with a planar distribution in color space. However, the colors do not fall randomly in a plane, but form clusters at specific points in color space. The location, dimensions, and orientation of these clusters directly relate to many scene properties. A full analysis of the histogram leads to a description of surface roughness and imaging geometry, as well as an improved estimate of illumination color and object color.

1. Introduction

Color histograms have long been used by the machine vision community in image understanding. Color is usually thought of as an important property of objects, and is often used for segmentation and classification. Unfortunately color is not uniform for all objects of a given class, nor even across a single object. Color variation has come to be expected in images, and vision researchers have been working on modeling this variation.

The earliest uses of color histograms modeled the histogram as a Gaussian cluster in color space [3]. For example, pixels that correspond to grass could be modeled as having a characteristic color of green with some possible deviation from this color. The color variation was modeled as a probability distribution, so that the further from the characteristic color, the less likely it was that a pixel was grass.

In 1984 Shafer showed that for dielectric materials with highlights, the color histogram associated with a single object forms a plane [8]. This plane is defined by two color vectors: a body reflection vector and a surface reflection vector. At the time the histogram was described as filling a parallelogram. The paper did not describe how to determine the two characteristic vectors out of the infinite set of vectors that could define the plane.

However, in 1987 Klinker and Gershon independently observed that the color histogram does not uniformly fill a parallelogram, but instead forms a T-shape or dog-leg in color space [7],[2]. They showed that it is composed of two linear clusters, one corresponding to pixels that exhibit mostly body reflection and one corresponding to pixels that exhibit mostly surface reflection. This T-shape made it possible to identify characteristic body reflection and illumination colors. In 1988, Healey showed that the number of dimensions occupied by the histogram may be used to distinguish metals from dielectrics [4].

In this paper, we show that the color histogram has an even closer relationship to scene properties than has been previously described. Color histograms have identifiable features that relate in a precise mathematical way to scene properties. Object color and illumination color are the most obvious properties that are related to color distribution, and their extraction has already been described [7],[5],[2]. We show that the histogram of color variation may be further exploited to relate its shape to surface roughness and imaging geometry. Furthermore, an understanding of these features allows us to make an improved estimate of illumination color and object color.

2. Color Histogram for a Single Object

When we talk about the color histogram, we mean a distribution of colors in the three-dimensional RGB space. For a typical imaging system with 8 bits for each color band, there are 256^3 "bins" into which a pixel may fall. In this paper, we only consider whether a bin is full or empty. We do not use a fourth dimension to display the number of pixels which have a particular RGB value. A fourth dimension would be difficult to visualize and also would be dependent on such things as object size and shape.

The histograms that we use in this paper were obtained either from synthetic images or by hand segmentation of real images. Obviously the end goal is to have a fully automated system that can perform its own segmentation, but

0-8186-2855-3/92 $3.00

that is beyond the scope of this paper. Klinker has shown that an automatic segmentation can be achieved for dielectric surfaces [6]; the analysis that we describe here could be performed on images automatically segmented in this way.

Figure 1 contains a sketch of a typical color histogram for a dielectric surface illuminated by a single light source. As labeled, the histogram has two linear clusters of pixels: the *body reflection cluster* and the *highlight cluster*. The first of these clusters extends from the black corner of the cube (point **a**) to the point of maximum body reflection (point **b**). The other cluster starts somewhere along the body reflection cluster (point **c**) and extends to the highlight maximum (point **d**).

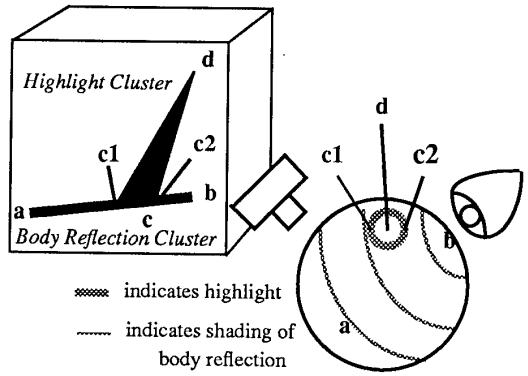

Figure 1: Histogram of single object

If the body reflection color is the same hue as the surface reflection color, the body reflection cluster and highlight cluster will be collinear. This is the case for white and gray objects since neither their body color nor surface color imparts any hue to reflected light. Therefore, objects of these colors cannot be analyzed by this type of method.

2.1. The Body Reflection Cluster

The linear cluster that we call the body reflection cluster corresponds to pixels that exhibit mostly body reflection with very little surface reflection. If there is no ambient illumination in the scene, this cluster begins at the black point of the color cube (point **a**), corresponding to points on the surface whose normal is 90 degrees or more away from the direction of the illumination. Such points fall beyond the self-shadow line and so are completely dark. The point at the other extreme of the body reflection cluster (point **b**), corresponds to the largest amount of body reflection seen anywhere on the object. If we assume that the body reflection component is Lambertian, pixels at this point in the histogram correspond to points on the surface with normals pointing directly at the light source. The body reflection m_b will obey the relation

$$m_b = c_b \cos(\theta_i) \qquad (1)$$

where θ_i is the angle of illumination incidence and c_b is the body reflection color. Thus pixels located half-way along the body reflection cluster would correspond to surface points with normals $\cos^{-1}(1/2)$ or 60 degrees away from the illumination direction.

If the object exhibits all possible surface normals, the body reflection cluster will be full length and densely filled. If the object is composed of a small number of flat surfaces, there will be gaps in the body reflection cluster. For this paper we will assume that objects we are looking at have a broad, continuous distribution of surface normals.

The sketch shows that the body reflection cluster is long and narrow. This agrees with our simulations and also with real, high-quality images of clean surfaces. A vector fitted to this cluster (from point **a** to **b**) will point in the direction of the body reflection color which is the product of the object color and the illumination color. Automatic fitting of this vector has been successfully demonstrated [6]. Once the illumination color has been determined from analysis of the highlight, the object color alone may be calculated by dividing out the influence of the illumination, as proposed in some color constancy methods [1],[5],[9].

2.2. The Highlight Cluster

The cluster of pixels we refer to as the highlight cluster corresponds to pixels that show a non-negligible amount of surface reflection. This corresponds exactly to the area of the image that we would call the highlight. In the histogram, the highlight cluster starts where it intersects with the body reflection cluster (point **c**) and extends upwards from there to the brightest point of the highlight (point **d**). For many shiny objects, the highlight is so bright that the highlight cluster is clipped at the white point of the color cube where the highlight has saturated the camera [6].

In this presentation we use the Torrance-Sparrow model of scattering [10]. This models a surface as a collection of tiny facets, each of which may have a local surface normal that is different from the global surface normal. For shiny surfaces, most of the facets have an orientation very close to the global surface normal. Rougher surfaces have a greater number of facets that are tilted with respect to the global normal. The distribution of facet normals is modeled as Gaussian, with σ describing the standard deviation. The distribution is also assumed to be isotropic, with rotational symmetry about the surface normal. The facets are larger than the wavelength of visible light, but too small to be seen as texture. We will assume that the facet size is a constant for the surfaces we are interested in.

The equation that we use for scattering gives the amount of surface reflection m_s as

$$m_s = c_s \frac{F G \beta}{\sigma \cos(\theta_r)} \exp\left(-\frac{\theta_s^2}{2\sigma^2}\right) \qquad (2)$$

where θ_s is the off-specular angle and θ_r is the angle of reflectance. The surface reflection color is given as c_s; β is a constant that includes the facet size (a variable in the original Torrance-Sparrow model). F is the Fresnel coefficient that describes what percentage of the light is reflected at the interface; it is a function of geometry, wavelength, polarization state, and index of refraction of the material in question. G is an attenuation factor that depends upon geometry and which comes into play at grazing angles. (G is a complex function of incidence angle, reflectance angle, and off-specular angle and we will not reproduce it here; see [10] for details).

2.2.1. Direction of Highlight Cluster

Surface reflection occurs at the interface, so it is not affected by the body color of inhomogeneous objects. The Fresnel coefficient F is very weakly dependent upon incidence angle and wavelength of illumination [8]. Thus it is often assumed to be a constant for a given type of material. If we follow this assumption, the only term in equation (2) that relates to hue is the illumination color c_s. All other terms relate to the magnitude of surface reflection. Thus the direction of the surface reflection component is equal to the illumination color.

The highlight cluster is usually long and narrow in shape and a vector can be fitted to it (from point **c** to **d**). Klinker argued that this vector will usually correspond closely to the surface reflection color [6]. This is true for smooth objects where the highlight has a small area, and for imaging geometries where the body reflection changes slowly over that area. In this case, the amount of body reflection at the base of the highlight cluster and the amount at the tip varies by a small amount.

On the other hand, if the object is optically rough and the highlight occurs on a part of the object where the cosine of the incidence angle changes more rapidly, then the amount of body reflection at the base of the highlight cluster may vary significantly from the amount at the tip. This has the effect of skewing the highlight cluster away from the direction of the illumination color. The estimate of the illumination color made from fitting a vector to this cluster will be somewhat inaccurate.

We can visualize this phenomena by projecting the histogram into the plane defined by the body reflection color and surface reflection color. We simulated dichromatic reflection for objects with the same body color but with different roughness values. Figure 2 shows a cross section of the histograms that result. The horizontal direction is the direction of increasing amounts of body reflection; all three histograms fall right on this line for pixels where there is no surface reflection. The vertical direction is defined by increasing amounts of surface reflection. For roughness values of 4 (or less), a vector fitted to the highlight cluster will point exactly in this direction. A vector fitted to the highlight cluster when σ equals 8 will deviate slightly from the vertical direction. In the extreme case, where σ equals 16, the vector will deviate dramatically.

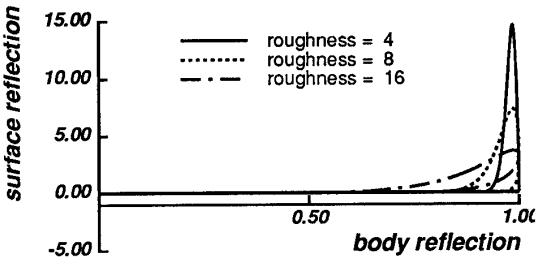

Figure 2: Dependence of highlight cluster direction on object roughness

Figure 3 shows an enlargement of the histogram of the roughest surface, with key points labeled. An estimate of illumination color from this highlight cluster will give the direction of the light color as $\overline{cd}$, whereas the correct illumination color is defined by the direction $\overline{c'd}$.

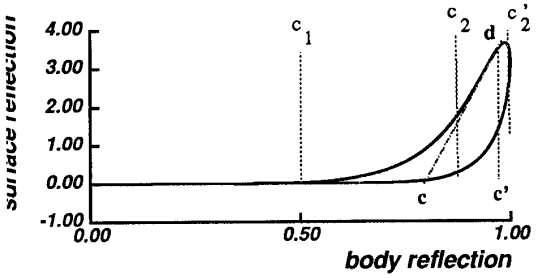

Figure 3: A skewed highlight

In the color histogram, the vectors describing the body reflection color and the illumination color are not generally perpendicular. The angle between them depends upon the hue difference between the two colors, which is not known in advance. If the vector fitted to the highlight cluster does not point exactly in the direction of increasing amounts of surface reflection, the estimate of illumination color will be off by some amount. This in turn will bias the estimate of the object color which is obtained by dividing the body reflection color by the illumination color.

Automatic fitting of a vector to highlight clusters to estimate illumination color was demonstrated in [6]. The vector fitted to the highlight cluster (from point **c** to point **d**) is a good first estimate of the illumination color, but we now know that it may be skewed. If we know the surface roughness and the imaging geometry, we can calculate the amount of skewing and compensate for it.

In the next few sections we will show how these factors may be estimated directly from the histogram itself.

2.2.2. Length of Highlight Cluster

When looking at highlights on a variety of surfaces, we quickly observe that highlights are brighter and sharper on some surfaces, while they are dimmer and more diffused on other surfaces. Very shiny surfaces exhibit only a tiny amount of scattering of the surface reflection, whereas very matte surfaces have a great deal of scattering. This scattering of surface reflection is a result of the optical roughness of the surface.

We see from equation (2) that the sharpness of the peak is determined by the standard deviation σ, and that the height of the peak is inversely proportional to σ. Intuitively this makes sense, since surface reflection scattered over a very small area will be more "concentrated." A smooth object will have a small standard deviation of facet slopes, σ, resulting in a long highlight cluster. A rough object will have a large σ, and so will exhibit a shorter cluster. Figure 4 shows a plot of the length of the highlight cluster vs. the object's roughness for simulated images where all other factors have been held constant.

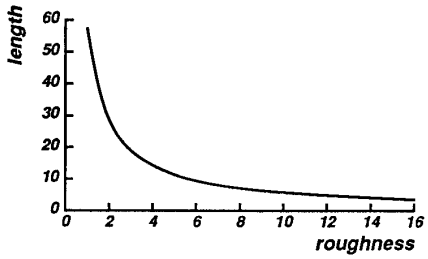

Figure 4: Highlight cluster length vs. roughness

Suppose that we want to estimate roughness values for a variety of different color objects. The body reflection will be different for each of the surfaces, but the amounts of surface reflection are not affected by this. The length of the highlight cluster (the distance from **c** to **d**) indicates the degree of roughness, regardless of body color.

According to equation (2), we need to know the Fresnel coefficient for the material we are looking at in order to relate roughness to highlight cluster length. However, it turns out that for a wide range of plastics and paints, the indices of refraction are very nearly identical. Henceforth we will assume that materials have an index of refraction of 1.5, corresponding to 4.0% Fresnel reflectance.

Figure 2 showed that if the surface is very rough ($\sigma = 16$), the highlight cluster will be skewed from the direction of the illumination color. This makes the cluster longer than it would otherwise be. The amount of surface reflection at the brightest point is the vertical displacement in the graph (the distance from point **c**′ to point **d** in Figure 3) rather than the distance along the highlight cluster (distance from point **c** to point **d**). For the smoother surfaces shown in the graph, the length of the highlight cluster is virtually identical to the amount of surface reflection. For very rough surfaces, we need to be aware of the effect of highlight skewing.

The graph in Figure 4 was calculated for the imaging geometry where the light source and camera are separated by zero degrees. However, equation (2) predicts that the imaging geometry will have an effect upon highlight magnitude, as indicated by the $\cos(\theta_r)$ term in the denominator and the attenuation term G in the numerator. G corrects for masking or shadowing of facets at grazing angles. As θ_r increases to 90 degrees, the attenuation term G goes to zero, so the equation does not make the improbable prediction of surface reflection becoming infinite at reflectance angles equal to 90 degrees. Figure 5 shows how the length of the highlight cluster changes as the camera and light source are separated by different angles with respect to the object. It demonstrates that the length changes slowly as the imaging geometry changes. If we do not know the imaging geometry, it would be impossible to make fine distinctions of roughness, although it would still be possible to tell a very smooth object from a fairly rough one. Furthermore, in section 2.2.4. we will show how the imaging geometry may be estimated.

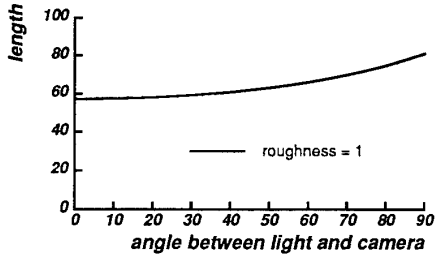

Figure 5: Cluster length vs. imaging geometry

The result of this section is that the roughness of surfaces made from similar materials can be directly compared by examining their color histograms. The method does not require us to fit surface orientations to image data. It can be done from a single color image of the object. At worst (when the materials are different or the imaging geometry is not held constant) we can make a crude estimate of roughness. At best, we can make a fairly accurate estimate.

Once we have the length of the highlight cluster, we can make a good estimate of surface roughness. Unfortunately, we cannot always obtain the true length. In many real images the highlight is so shiny that it saturates the camera. This means that the highlight cluster is clipped, so that its true length is unknown. We will be able to distinguish rough objects from smooth ones, but we would be unable to distinguish *fairly* smooth surfaces from *very* smooth

ones by this method. Fortunately, the histogram encodes roughness in another feature.

2.2.3. Width of Highlight Cluster

Another difference between histograms for smooth and rough surfaces is the width of the highlight cluster where it meets the body reflection cluster (the distance from point c_1 to point c_2 in Figure 1). The highlight cluster will be wider for rougher surfaces, and narrower for smoother surfaces. We see from equation (2) that for rougher objects with a larger standard deviation of facet angles σ, the surface reflection is scattered more widely, over a larger number of reflectance angles.

In the color histogram, a noticeable amount of surface reflection results in pixels that are displaced from the body cluster in the direction of the illumination color. If we take any highlight pixel and project along the surface color vector onto the body reflection vector, we can tell how much body reflection is present in that pixel. If we consider all the pixels in the highlight area of the image and look at how much body reflection is in each of them, we will obtain some range of body reflection magnitudes. If the surface is very smooth with a sharp, bright highlight, that range will be small. However if we consider a rougher object with a more diffused highlight, the range of body reflection magnitudes will be larger since the highlight is spread over a larger number of surface normals.

This property is independent of object size or shape. It simply shows the variation in surface normals over the area of the highlight. We do not have to fit a surface shape to the image to know how much scattering the object exhibits. It is encoded right there in the histogram.

We simulated objects with different roughness values and identified those pixels that showed surface reflection. We calculated the body reflection for each of these points and computed the variation. This variation was divided by the overall length of the body reflection vector to yield a fraction (the length of $\overline{c_1c_2}$ divided by the length of $\overline{ab}$ in Figure 1). A fraction of 0.5 would mean that the highlight cluster's base extended across half the length of the body reflection cluster. Figure 6 shows how the highlight cluster width varies with the surface roughness.

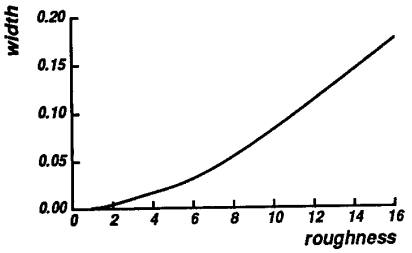

Figure 6: Highlight cluster width vs. roughness

Moreover this measurement is not hampered by camera saturation. It does not depend on how *much* surface reflection is at any given point, only whether there is *any* at all. Camera saturation is an inherent property of any real camera, since the dynamic range is limited. Since highlights are very often clipped in real images, it is very useful to have a feature that is not affected by saturation.

In the case of very rough surfaces, the highlight cluster may be skewed from the direction of the illumination color, as shown in Figure 3. The width of this cluster where it meets the body reflection cluster *underestimates* the variation of body reflectance values. The amount of body reflectance is correctly determined by projecting all pixels along the illumination color vector (straight down for this graph). The correct measure of highlight width is the variation in body reflectance (the distance from c_1 to c_2' rather than the distance from c_1 to c_2). However, the highlight and body reflection colors are not generally perpendicular in most histograms, so we cannot just measure the horizontal extent of the highlight cluster. For the smoother surfaces shown in Figure 2 (when $\sigma \leq 8$), the highlight cluster is centered above its base, so the width of the base and the horizontal extent of the cluster are virtually identical. For these types of surfaces, the width at the base is a very good approximation of the overall width. However, for very rough surfaces, the width may be significantly affected by highlight skewing.

Although the width of the highlight cluster at its base does not depend upon the object's size and shape, it does depend upon the imaging geometry. To see why this is so, imagine a highlight that spreads 15 degrees in every direction from its maximum. If the camera and light source are separated by 30 degrees, the perfect specular angle will be at 15 degrees with respect to the illumination direction. The highlight will spread over points with surface normals ranging from 0 degrees to 30 degrees. (For ease of explanation, we will ignore the influence of the $1/\cos(\theta_r)$ term.) The amount of body reflection at these points will vary from cos(0) = 1.0 to cos(30) = .87. If the camera and light source are separated by 90 degrees, the perfect specular angle will be at 45 degrees with the highlight spreading from 30 degrees to 60 degrees. Then the amount of body reflection will vary from cos(30) = .87 to cos(60) = .50.

Figure 7 shows how the width of the highlight cluster varies with roughness for a variety of imaging geometries. The angle label is a measure of the angle (in degrees) that separates the light source and camera (with respect to the object). We assume that the object is small enough and far away enough from the camera that this angle is the same for all points on the surface.

For the case of highlight cluster width, the measurement is very sensitive to different viewing geometries, so the angular separation of the camera and light source must be

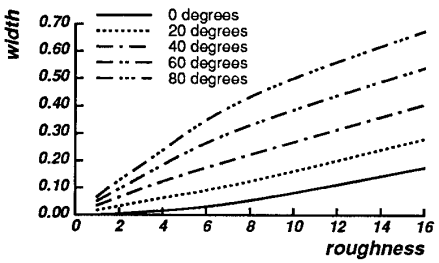

Figure 7: Width vs. roughness (with geometry)

known or estimated somehow. Thus it is particularly fortuitous that such an estimate can be made right from the histogram itself, as we will now describe.

2.2.4. Intersection of Clusters

When we first introduced the diagram in Figure 1, we described the highlight cluster as beginning "somewhere" along the body reflection cluster. Klinker derived the "50% heuristic" which stated that for a large range of viewing geometries, the highlight cluster would start somewhere in the upper 50% of the body reflection cluster [6]. Now we will show how to pinpoint the location.

The distance along the body reflection cluster where the two clusters meet (the length of $\overline{ac}$ divided by the length of $\overline{ab}$) shows the amount of body reflectance at those points on the surface that are highlighted. Assuming that body reflection is Lambertian, we know from equation (1) that the amount of body reflection is proportional to the cosine of the incidence angle θ_i. If the two clusters meet at the maximum point on the body reflection cluster, it means the highlight occurs at those points that have the maximum amount of body reflection, which is precisely those points with surface normals pointing directly at the light source. If the two clusters meet halfway along the body reflection cluster, the highlight must occur at points with surface normal pointing $\cos^{-1}(1/2)$ or 60 degrees away from the illumination direction.

Assuming that the body reflection is Lambertian, it does not depend in any way upon the angle from which it is viewed. Thus the body reflection does not tell us anything about the camera direction. However, the surface reflection is dependent upon both the illumination and camera directions. If we ignore for the moment the $1/\cos(\theta_r)$ term in equation (2), we see that the maximum amount of surface reflection will occur at those points on the surface where the angle of incidence equals the angle of reflection. Thus if the highlight occurs at a point where the surface normal faces 10 degrees away from the light source direction, the light source and camera must be 20 degrees apart with respect to that point on the surface.

Figure 8 graphically illustrates this phenomenon. Once again the histograms have been projected into the plane defined by the body reflection and surface reflection colors. This time, the amount of roughness has been held constant at $\sigma = 1$ while the angular separation of the light source and camera has been varied from 0 to 80 degrees. This graph shows how the meeting point decreases as the angle separating the camera and light source increases. Incidentally it also shows how the length and width of the cluster is affected by imaging geometry as described in sections 2.2.2. and 2.2.3.

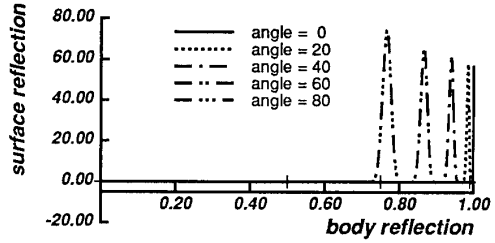

Figure 8: Illustration of how highlight cluster varies with changing imaging geometry

It does not matter whether the object has one highlight or many. The highlight or highlights will always occur at points with the same surface normal for a given imaging geometry, and that is what determines the meeting point of the clusters. Figure 9 shows what happens when we graph meeting point vs. imaging geometry. The horizontal axis shows the angular separation of the light source and the camera with respect to the object. A negative number means that the light source is to the left of the camera, whereas a positive number means the light source is to the right of the camera. The vertical axis is the ratio $|\overline{ac}|/|\overline{ab}|$. As can be seen from the graph, the meeting point is never less than 1/2 of the way along the body reflection vector, showing that the 50% heuristic is a sound one.

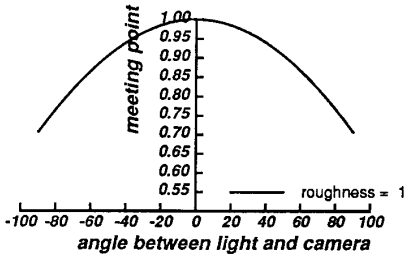

Figure 9: Meeting point vs. imaging geometry

We have indicated that the intersection point of the two clusters shows the amount of body reflectance at the brightest point in the highlight. This is true for cases where the highlight cluster extends in approximately the same direction as the illumination color. However if the highlight cluster is significantly skewed away from the direction of the illumination color, the intersection point will give an incorrect estimate of the imaging geometry. This is the case

for the roughest surface (σ = 16) shown in Figure 2. All three histograms correspond to an angular separation of 20 degrees, and indeed the highest point of each highlight cluster occurs at the same horizontal coordinate (which indicates the amount of body reflection). However, for the rough surface shown in Figure 3, the correct intersection ratio is given by $|\overline{ac'}|/|\overline{ab}|$ rather than $|\overline{ac}|/|\overline{ab}|$. This skewing of the highlight cluster means that the meeting point may be significantly altered for very rough surfaces.

The $1/\cos(\theta_r)$ term in equation (2) means that the maximum amount of surface reflection will not always occur precisely at the perfect specular angle. This is particularly true of rougher surfaces where the highlight is spread over a wide range of reflectance angles so that $1/\cos(\theta_r)$ varies significantly. This causes the "off-specular peaks" described in [10]. The result is that the meeting point is very slightly dependent upon the surface roughness.

In Figure 9 we see that there is an ambiguity about 0 degrees. We cannot tell from the histogram whether a light source is 45 degrees to the left or to the right of the camera. More generally, the light source could be anywhere on a 45° circle around the camera position. This method for determining the angle separating the light and camera by examining the cluster meeting point relies implicitly on the amount of body reflection at certain points. However, it does not require us to fit surface orientations to image data, nor does it require knowledge of the albedo of the surface. Moreover, it is done from a single image.

In sections 2.2.2. and 2.2.3. we claimed that we could determine the roughness of the object from the length and width of the highlight cluster, but that there was some dependence on imaging geometry. Now we are claiming that we can determine the imaging geometry from the meeting point of the two clusters, but that there is some dependence upon roughness. Furthermore, in section 2.2.1. we noted that the direction of the highlight cluster could be skewed away from the direction of the illumination color, depending upon the amount of roughness and the imaging geometry. Now we see that the estimates of cluster length, width, and intersection point depend to a certain extent upon knowing the true direction of the illumination color.

Obviously these factors are all interdependent. Moreover because of possible highlight skewing, our initial estimates of cluster length, width, and intersection point may be slightly or even significantly affected. Therefore we propose to solve for the amount of roughness and angular separation of the light and camera simultaneously, based on the initial estimates. Such a solution could be obtained by generating a lookup table for a large range of roughness values and imaging geometry.

3. Conclusions

The color histogram of an image is a rich source of information, but it has not been fully exploited in the past. The variation of color was once modeled as a random process with a Gaussian distribution. More recently, the model has become more structured, showing that the distribution of colors on an inhomogeneous object will form a plane in RGB space, and that the pixels fall within a parallelogram defined by the body color and illumination color. Experiments then showed that the parallelogram is not uniformly filled; rather, the color histogram will show two distinct clusters that meet in RGB space.

We have gone further to show that the direction, length, width, and intersection of these two clusters may by characterized by numerical measurements. The resulting measurements are not just of academic interest, but relate directly to scene properties, including both object color and roughness, as well as illumination color and position.

4. Acknowledgments

This research was sponsored by the Avionics Lab, Wright Research and Development Center, Aeronautical Systems Division (AFSC), U. S. Air Force, Wright-Patterson AFB, OH 45433-6543 under Contract F33615-90-C-1465, Arpa Order No. 7597; and by the Air Force Office of Scientific Research under Contract F49620-86-C-0127. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the U.S. Government.

References

[1] D'Zmura, M. and Lennie, P. "Mechanisms of color constancy." *J. Opt. Soc. Am.*, 3(10) 1986, 1622-1672.

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

[3] Haralick, R. M. and Kelly, G. L. "Pattern recognition with measurement space and spatial clustering for multiple images." *Proc IEEE*, 57, 1969, 654-665.

[4] Healey, G. "A color reflectance model and its use for segmentation." *Intnl. Conf. on Comp. Vision*, IEEE, 1987, 460-466.

[5] Healey, G. and Binford, T. O. "The role and use of color in a general vision system." *Proc. of the ARPA Image Understanding Workshop*, 1987, 599-613.

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

[7] Klinker, G. J. and Shafer, S. A. and Kanade, T. "Using a color reflection model to separate highlights from object color." *Intnl. Conf. on Comp. Vision*, IEEE, 1987, 145-150.

[8] Shafer, S. A. "Optical phenomena in computer vision." *Proc.* Canadian Soc. Comp. Studies of Intel., 1984, 572-577.

[9] Tominaga, S. and Wandell, B. A. "Standard surface-reflectance model and illuminant estimation", *J. Opt. Soc. Am.* 6(4), 1989, 576-584.

[10] Torrance, K. and Sparrow, E. "Theory for off-specular reflection from roughened surfaces." *J. Opt. Soc. Am.* (57), 1967.