Highlight Identification Using Chromatic Information

Ron Gershon, Allan D. Jepson*, and John K. Tsotsos*

Department of Computer Science
University of Toronto
Toronto, CANADA M5S 1A4

Abstract

This paper presents a method for identifying highlights in images through the use of chromatic information. An existing model of reflectance is shown to predict that a highlight will usually involve a gradual additive color shift from the color of the object to the color of the illuminant. This shift is observed in real images and forms the basis of our algorithm.

1 Introduction

The appearance of objects in the real world is related to the complex ways they reflect the light impinging on them from the light sources present. Different materials reflect light in various ways, but one thing common to most materials is the fact that they are neither totally matte nor totally mirror-like; it is a combination of these two modes of reflection that governs the appearance of most objects. In short, it is the sum of diffuse and specular reflections that determines the total reflection of light off surfaces. In this paper, we concern ourselves with one of the two reflections mentioned, namely, the specular reflection. This reflection is the one usually associated with highlights, or glossiness. The goals set in this paper are to: a) review a physical model underlying the structure of highlights; b) develop a statement relating this model to the actual appearance of highlights in images; and c) design an algorithm which will use this relationship and determine which image regions are probable highlights.

2 Why Highlights?

The motivation behind the quest for highlights originates in two opposing properties highlights possess. On one hand, highlights in images provide important information which can be used to determine shape and size of objects. Thrift and Lee [19] presented an analysis of how the observation of highlights on cylinders and spheres provides constraints on the size and location of these objects. Babu et al. [1] proposed methods of estimating the orientation of specularly reflecting planar surfaces from the shapes of contours of constant brightness. Ikeuchi [12] and Blake [4] use highlights found in pairs of stereo images to infer information about surface curvature. Buchanan [5] presents an algorithm for determining the surface normal of a plane from highlights in a single image, assuming knowledge of a relative distance of the light and viewer to the plane. He also proposes a method for determining the surface normal in a pair of images with no underlying distance assumptions.

On the other hand, highlights cause numerous problems for algorithms which require correspondence between consecutive images, such as stereo or motion algorithms. In an algorithm which generated a representation for a moving object of unknown form and size from a sequence of TV images, Nagel [14] pointed out that the parts of the moving object which were not significantly structured along the direction of motion may have been regions containing highlights. Thorpe [18] analyzed different interest operators which pick distinctive points to be tracked in sequences of stereo images for the navigation of robots. He found that one of the common problems to all the operators investigated is that they pick points which are easy to find in one image, but disappear in others. Such points are clearly viewpoint-dependent, and highlights are one of the causes for this behavior.

3 Previous Work

Ullman [20] investigated computational aspects underlying the detection of light sources by the human visual

* Also, Canadian Institute for Advanced Research

system. The domain in which Ullman carried out his study was that of achromatic Mondrians, so as to exclude colors and effects of shading and thus create an environment void of high level knowledge about light sources. The method proposed to detect light sources was based on two observations. The first was that the only measurement needed for the detection of light sources was the correct value of reflectance ratios in the scene. The second observation was that in many cases this real reflectance ratio can be computed by comparing intensity gradients. Therefore the method Ullman proposed was to take two adjacent areas and compute both their intensity ratio and their gradient ratio, and compare the two; if the ratios were not equal, then one of the areas was probably a light source. This resulted in an operator, termed as "S-Operator", which was implemented and tested on one-dimensional patterns of intensity. Some problems with the method and its assumptions are listed below:

- There is no actual specification of how the areas are to be computed, i.e., what criterion is used to distinguish one area from another.
- There are processes other than highlights (such as changes in pigments) which cause the gradient ratios to differ from intensity ratios.
- Beck (1964) pointed out that the first stage in the perception of glossy surfaces (which include light sources) was a contrast process that acts to determine a pattern of lightness in accordance with the distribution of reflected intensities. Therefore contrast cannot be discounted as a factor taking part in the perception of light sources.
- The extension of the method to chromatic images is not specified in the study and poses many questions. Although the goal was to examine the problem in a restricted domain, the ability to detect light sources is not confined to achromatic images and, in fact, color is a useful cue in the detection process.

Shafer [17] proposed a different approach to solving the problem of identifying highlights in images. The basis to Shafer's work is a simple mathematical model of reflectance, called the Dichromatic Reflection Model. This model states that the total radiance L of the reflected light is a sum of two independent parts: the radiance L_i of the reflected light at interface (the specular reflected intensity), and the radiance L_b of the light reflected from the surface body (the diffuse reflected intensity). Furthermore, the model states that these components of the reflected intensities can be decomposed into two parts: *composition*, a relative spectral power distribution c_i and c_b which depends only on wavelength, and *magnitude*, a geometric scale factor m_i and m_b which depends only on geometry. Formally, the model states:

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

where i, e, and g are the angles of incidence, emittance, and phase, respectively. With this model, an algorithm was proposed for computing the intrinsic images of m_i and m_b over a set of [R,G,B] values of pixels corresponding to a single surface in the image. The algorithm requires that a plane be fit to the pixel values, enclosed in a parallelogram such that the sides of the parallelogram are [R,G,B] values corresponding to c_i and c_b. The algorithm computes the magnitudes m_i and m_b. One problem with the model is:

- The algorithm proposed for the identification of the reflection components requires an *a priori* segmentation of the scene into regions which are known to lie on the same surface. Such a process of segmentation is not at all trivial, since part of the segmentation problem is to decide whether to include highlights within the region assigned to an object, or treat it as a separate object.

4 The Shift in Color

We have already seen in the previous section, in Equation (1), that the reflection of light striking an object is actually a sum of two components: specular and diffuse. This formulation is in agreement with other models of reflection, such as the Cook and Torrance model [6] or the Horn model [10]. This combination determines the appearance of the object everywhere, be it shadow, highlight, or any other conceivable region. The most intriguing aspect of this statement is the fact that specularity is an additive process; that is, the regions which appear glossy have the color of the illuminant added to the diffuse color of the surface. In other words, it is always the case that the transition from the diffuse to the specular regions involves a shift from one spectral distribution to another[1]. To illustrate, we use the Cook and Torrance model [6] to how when the weight, or magnitude, of the specular reflection becomes larger than that of the diffuse reflection. Suppose the geometry is such that the light

[1]This statement holds for nonmetallic objects. As for metals, there is hardly any shift in the spectral distribution, although the intensity of the highlight region is higher than that of the diffuse region.

source, the viewer, and the surface normal are coplanar. If we fix the viewer and the light source 120° apart, and thus $(g/2) = 60°$, we can change the surface normal and see how it affects the weights of the diffuse and specular reflectances. *Such is the case depicted in Figure 1, where* we changed the surface normal such that the angle of incidence varied between $i = 0$ and $i = 90$.[2] The results show

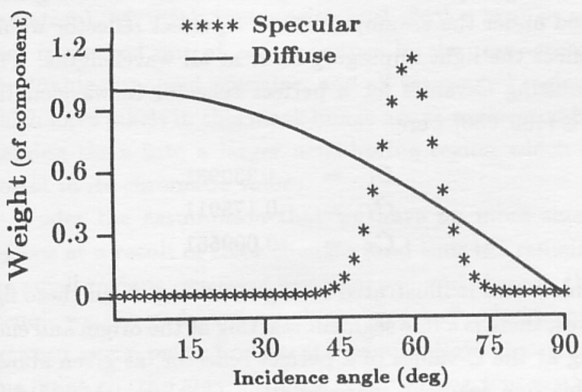

Figure 1: The diffuse and specular weights as a function of the incidence angle.

that as i approaches 60°, the specular weight becomes larger and more dominant. This shows that changes in the geometry imply changes in the relative weights of the two reflectance components. Thus, as the viewing and illumination geometry changes, there is a shift in the relative strengths of the diffuse and specular reflections, from a stronger diffuse reflection to a stronger specular reflection. Since this is in effect at all wavelengths, the implication on chromaticity will be a shift in spectral distribution from the diffuse spectral distribution to that of the light source.

Suppose that the images with which we deal are transformed into a color-constant space (*C-space*), one which is a linear transformation of the original [R,G,B] image into an image which discounts the illuminant; such a transformation is described in more detail elsewhere [8]. We therefore assume that the image has been transformed in such a way that no matter what illuminant is used in the scene, the C-values[3] of the pixels in the image will be similar for each distinct illuminant. This implies that if the highlights indeed reflect the spectral distribution of the light source (or close to it), then changing the light source (for the same scene) should result in the same C-*values for the highlight regions. Thus the shift in color* from the diffuse regions to the highlight regions should involve a transition from some C-values representing the (diffuse) spectral distribution of the object to C-values representing the perfect reflector (mirror). Moreover, this transition, barring any abrupt changes in the curvature or pigmentation of the object, should be a smooth transition, reflecting the monotonic increase in the magnitude of the specular reflection over the magnitude of the diffuse reflection. This transition, which is in fact the shift in color, is the phenomenon we wish to use as a cue towards the detection of highlights.

In order to verify that indeed the shift in color occurs in real images, some instances of highlights in images were sampled and viewed using two- or three-dimensional plots of C-space. In other words, the images were transformed into C-space and scatter plots of C-values were generated. The regions from which the C-values were picked were always adjacent highlight and diffuse areas. The resulting plots are presented in Figures 2 and 3; the original images are presented in Figures 4 and 5. In Figure 2, a

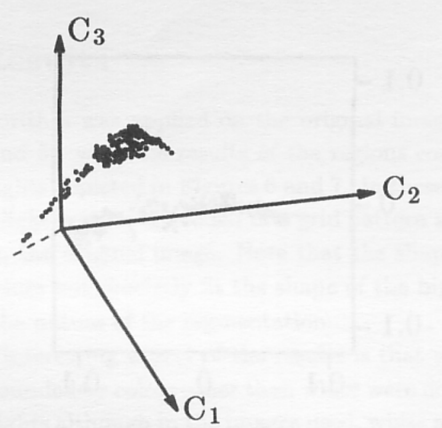

Figure 2: A 3-d scatter plot of a "dog-leg".

The dashed line represents the C-values of perfect reflectors.

three-dimensional scatter plot of C-space from the "jelly beans" image is presented, illustrating that around a highlight area, there is a "dog-leg" structure in the C-values. Figure 3 illustrates the same behavior through three 2-

[2] For the other parameters used in this example, see [8].

[3] The [R,G,B] values transformed into a three-dimensional color-constant space.

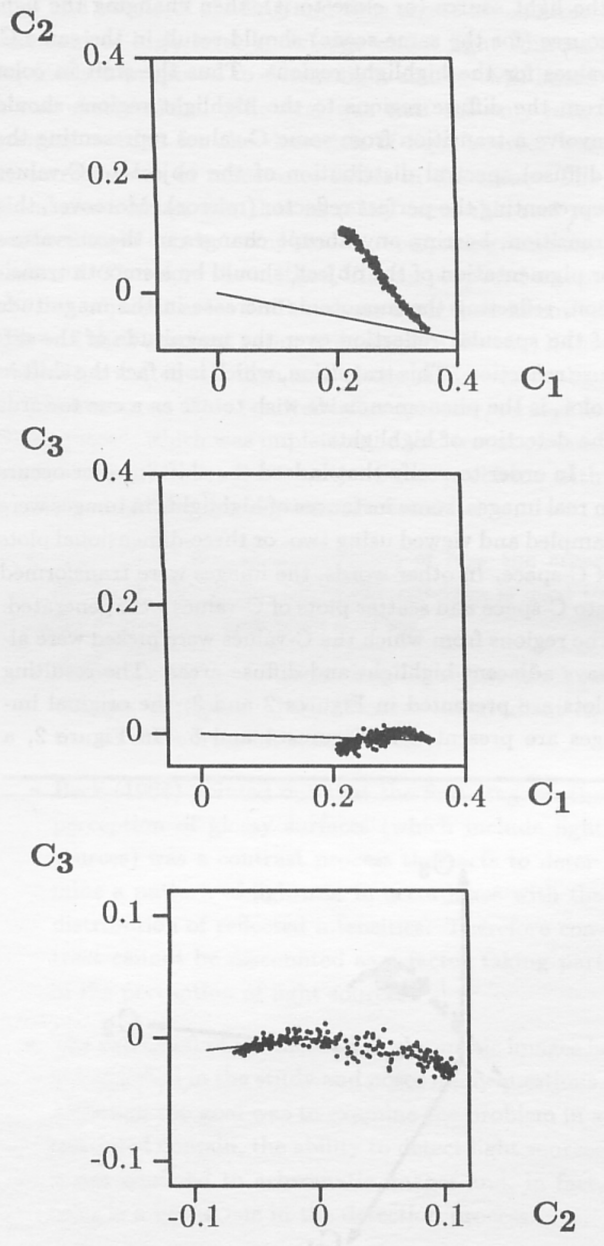

Figure 3: 2-d scatter plots of a "dog-leg" in C-space. The thin line represents the C-values of perfect reflectors.

d slices of the C-space for a highlight in the "peppers" image. As can be observed in the figures, there is indeed a smooth transition in the C-values in the form of a "dog-leg". Moreover, the transition always involves a shift of the color in a direction towards color of the illuminant. In terms of the C-space, this direction is actually pointing towards the values of a perfect reflector, since highlights are supposedly a realization of mirror-like surfaces. This observation was analyzed numerically in the following way. C-values of a perfect reflector were generated under the assumption that a perfect reflector would reflect the light impinging on it at all wavelengths. The resulting C-values for a perfect reflector in wavelengths $\lambda \in [400, 650]$ were:

$$\begin{aligned} C_1 &= 1.339931 \\ C_2 &= -0.175911 \\ C_3 &= -0.009551 \end{aligned}$$

This claim is illustrated in Figures 2 and 3. In these figures, there is a line segment starting at the origin and ending at the C-values of a perfect reflector, as given above. This line segment represents the C-values which reflect light equally at all wavelengths, for reflectances from 0 to 100%. We can see that one of the dog-leg points towards the end of this segment, especially for C_2 and C_3. C_1 may sometimes be smaller than its expected value as computed by the line segment because the highlight may not be very intense. This is a result of the fact that not all the light is specularly reflected, but rather some of it is absorbed or diffusely reflected.

5 The Algorithm

The algorithm designed to detect possible highlights is based on the observations made in the previous section. It is a two step algorithm: first, it requires the segmentation of the image into regions, and compilation of information about regions and neighbors, and secondly, it attempts to find a "dog-leg" relationship between every adjacent pair of regions. In this section we provide more details regarding each of the steps in the algorithm.

The segmentation is a split-and-merge spectrally based scheme [11]. It is important to note, though, that the segmentation is performed on the original [R,G,B] image and is based on chromatic differences. There is, however, one problem which must be addressed. Since the segmentation scheme is based on a split-and-merge technique, it creates many small regions which cannot be merged into bigger ones, particularly around areas where there is a change from one color to another, such as material changes or changes caused by highlights. Since our algorithm requires comparisons between adjacent regions, it is not desirable to have a collection of larger regions separated by a set of very small regions. This is essentially an instance of the "small region elimination" problem so typical of split-and-merge techniques [15]. We therefore propose the following solution: the areas in which such small regions may be formed are predicted through the use of a band-pass double-opponent operator, which signals a change in color (see [7,16] for the nature of these operators). A mask image with such "color change areas" is created out of zero-crossings in the response of the double-opponent operator, and all segmented regions which have pixels in this mask image are re-segmented by merging them into a larger neighboring region which is closest in its chromatic values.

Under the assumption that we have no more small regions as a result of color changes (and thus the remaining small regions represent small meaningful areas in the scene), we proceed to check for a "dog-leg" relationship between every pair of adjacent regions. Note that from this point on, the algorithm uses the C-value representation of the image, and not the original [R,G,B]. Thus the original [R,G,B] image was used but for the segmentation and as a basis for the transformation to C-space. In order to get the "legs" themselves, a least squares linear fit of the C-values in each region is computed, resulting in two lines in C-space, one for each region. If the two lines are close to being parallel then the two regions will not be considered any more. An approximation to the point representing the intersection of the two lines is computed. Furthermore, a check is made to verify that the lines indeed meet, by making sure that there are some pixels in both regions which have C-values around the point of intersection. Then the following test is performed: the region in which most the C_1 values are beyond the point of intersection is considered as a candidate for a highlight; the other region is considered diffuse. If both regions have most of their C_1 values beyond the points of intersection, the regions will not be considered any more, and a new pair will be picked. The reason behind this test is that C_1 represents in essence the luminance response,[4] and it is evident that if one of the two neighboring regions is a highlight, its luminance will be higher than the diffuse region. The final decision as to whether the candidate region will be considered a highlight requires one further test. This one checks whether the line approximating the C-values of the region points towards values of a perfect reflector. This is done by making a comparison of the imaginary line formed by connecting the point of intersection to the point representing the perfect reflector, with the line which was already fitted to the C-values of the candidate region.

To summarize, the algorithm is comprised of the following steps:

- Spectrally-based segmentation of the [R,G,B] image and aggregation of small regions into larger ones.
- Transformation of the [R,G,B] values into C-values.
- Least squares approximation of the C-values in each region to a line in 3-d.

For each pair of adjacent regions:

- Verification of whether the lines are parallel — if so, a new pair of regions is tested.
- Computation of the point of intersection between the two lines and verification that the two lines indeed meet.
- Detection of the region in which most of the C_1 values are beyond the point of intersection, and labeling of it as a candidate highlight. If both regions pass the test — a new pair of regions is tested.
- Verification of whether the line representing the candidate region is in the direction of a perfect reflector.

6 Results

The algorithm was applied on the original images (Figures 4 and 5), with the results of the regions considered as highlights depicted in Figures 6 and 7. In these figures, the highlight regions are coded in a grid pattern superimposed on the original image. Note that the shape of the regions does not precisely fit the shape of the highlights, due to the nature of the segmentation.

One interesting aspect of the results is that white areas surrounded by colors other than white were not picked as highlights although in the images used, white appeared to be the color of the light source. This can be seen in the "jelly beans" image, where there are a few locations where the background can be seen between some jelly beans. This is because there is no shift in color and thus no "dog-leg" in the C-values. Furthermore, the white areas are fairly "flat", i.e., there are hardly any shading effects, which implies that they are diffuse regions. To stretch this point even further, we "flattened" a highlight, (by assigning all the [R,G,B] values of the highlight region to the same number, which corresponds to white)

[4] See [8] for details.

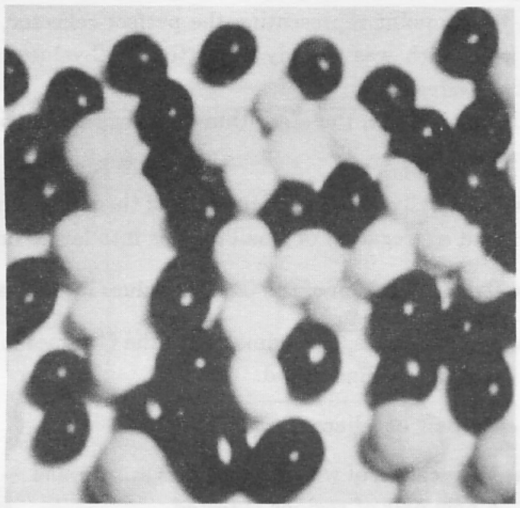

Figure 4: The original "jelly beans" image.

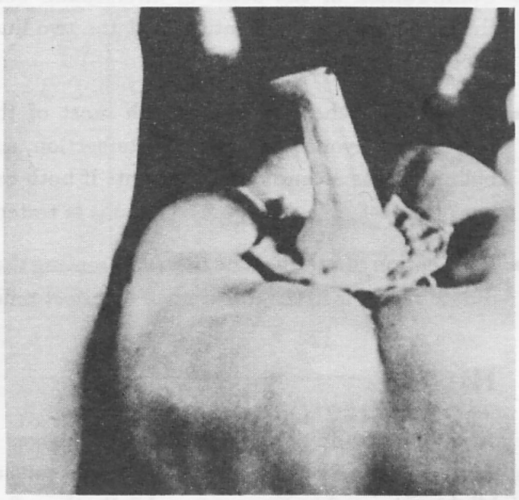

Figure 5: The original "peppers" image.

in the "peppers" image (the top left highlight on the pepper — see Figure 8), with the result that all the regions which were considered highlights previously still came out as highlights, except for the one which was subject to the change. The reason is twofold: first, the region is considered diffuse by the algorithm since the C-values in the "highlight" regions are all the same, meaning that the program cannot use least squares to fit a line. Furthermore, if this rule (about the inability to use least squares because of minimal distribution in the C-values) is taken out, the program still fails to find a "dog-leg" since the C-values of the two regions do not meet, assuming that a hypothetical point of intersection can be found. The output of the program is shown in Figure 9; note that the

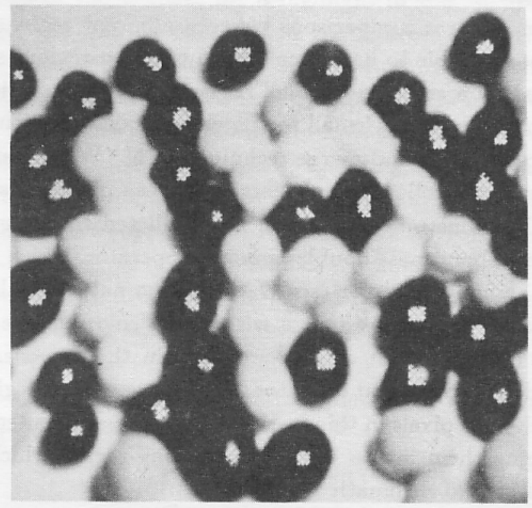

Figure 6: Highlights detected in the "jelly beans" image.

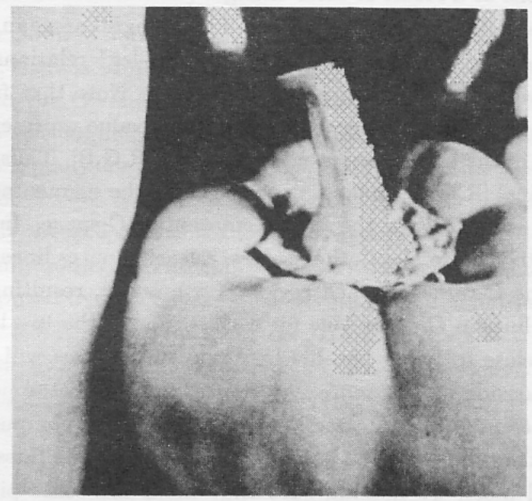

Figure 7: Highlights detected in the "peppers" image.

program ignored the flat highlight. The scatter plots of the C-values in the region around the flattened highlight are illustrated in Figure 10.

Another interesting result has to do with the uniformity of the distribution of light in surfaces, an important cue used by humans. Beck [2,3] conducted experiments on human perception of shadows, where the task was to point out whether two patches of an object looked as a shadow and lit area or not, both in the presence and absence of an obstructing patch between them. The results pointed towards the hypothesis that when a surface reflects a uniform distribution of intensities, it is perceived as one object with the changes (in intensities) attributed to orientation or illumination changes. We extended this

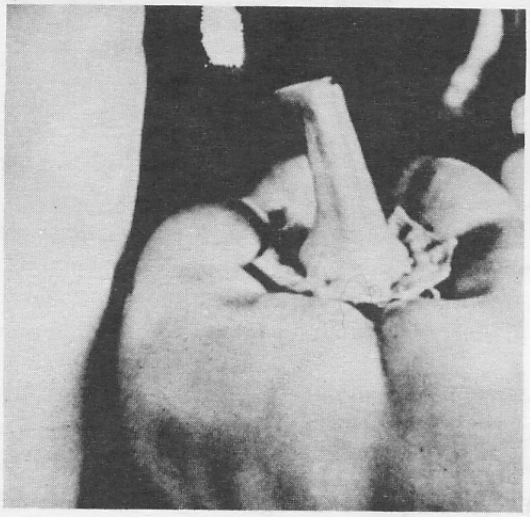

Figure 8: The "peppers" image with a flattened highlight.

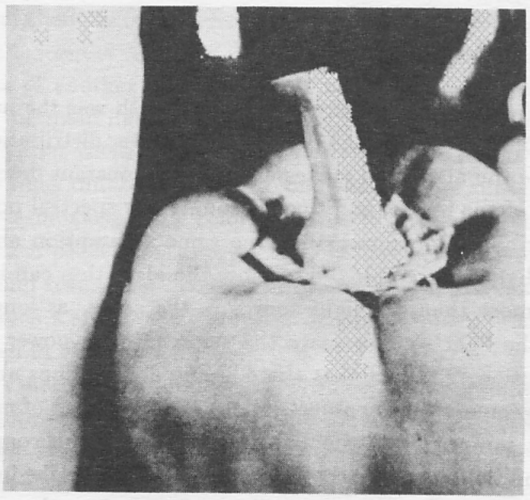

Figure 9: Highlights detected in the "flat highlight" image.

hypothesis to highlights, and created an obstructing patch (a yellow line) between the diffuse part of an object and its specular region (see Figure 11; the yellow line appears as the white stripe surrounding the left highlight in the darker (red) pepper). The algorithm was run on this image, and as a result of the segmentation, the yellow patch appeared as one segment located between the diffuse region highlight. Since the C-values of this segment were different from those of the diffuse region, and also from those of the highlight, as is illustrated by Figure 12, the program ignored this highlight in its final output (Figure 13).

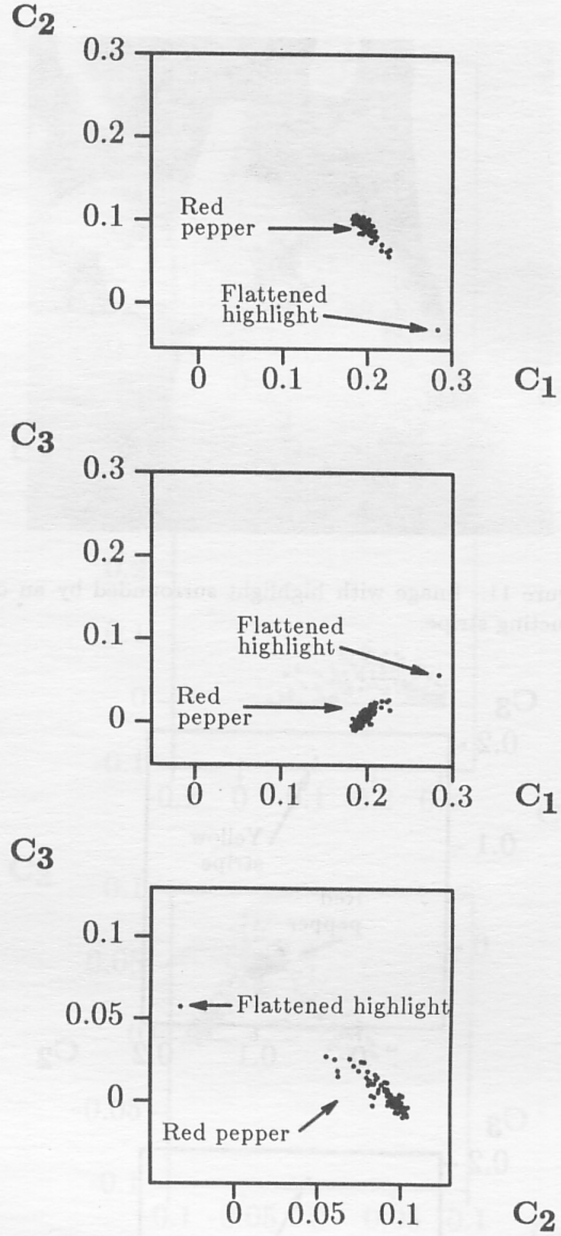

Figure 10: Scatter plots of the flat highlight and pepper.

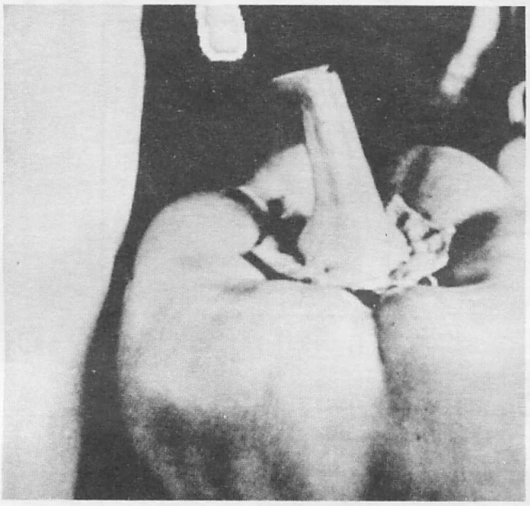

Figure 11: Image with highlight surrounded by an obstructing stripe.

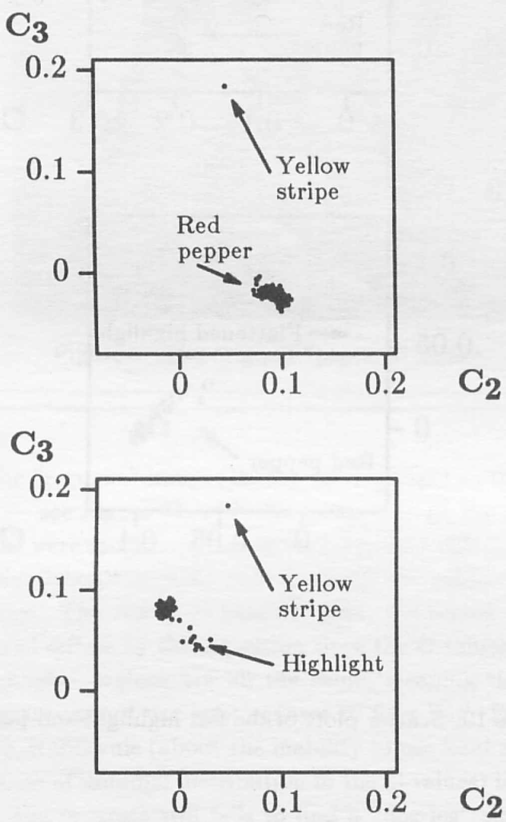

Figure 12: Scatter plot of the highlight, stripe, and pepper.

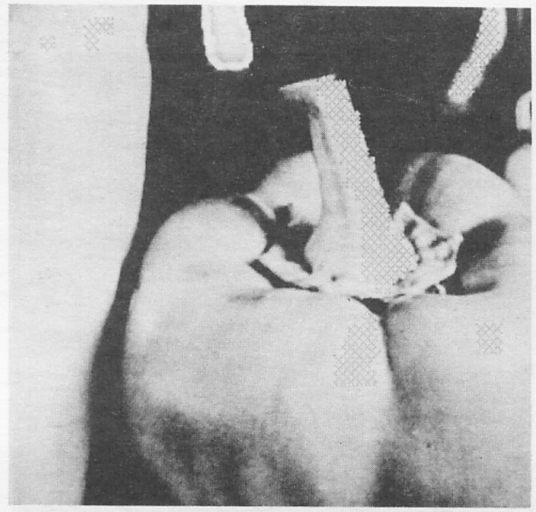

Figure 13: Output of the algorithm for the above image.

7 Discussion

One assumption which was not dealt with was the number of light sources and their spectral power distributions. Since the algorithm makes use of color-constant descriptors of the image, and thus discounts *one* spectral power distribution of illuminants, the same assumption about the light source holds. Therefore the algorithm can handle more than one light source in the scene, as long as *all* the light sources share the same spectral power distribution. The fact that there are no assumptions about the geometry of the objects allows the treatment of many light sources, since the algorithm is based on chromatic properties only. The scatter plot resulting from the inclusion of more light sources will not change since there will be no change in the chromaticity of the light sources, and consequently in the spectral distributions of the diffuse and specular regions.

An interesting observation can be made regarding diffuse regions and their corresponding C-values. It is expected that the line approximating to the C-values of a diffuse region should point to the origin of the C-space. The reason for such a behavior is that a diffuse region which exhibits some curvature will have values ranging from no color at all (if the diffuse region is totally in the shadow) to the fully saturated (pure) spectral distribution of the object. This range, in theory, is expected to lie on a line connecting the origin to the desaturated color, although there is no requirement that the object be in a shadow at all. However, in practice we have observed that because of noise, the least squares fit to a line yields lines which do not originate exactly from the origin, and thus there was no attempt to make any use of this observation. On the other hand, objects which do not exhibit considerable curvature (objects "flat" with respect to the light source) are expected to have very little variation in their C-values. This phenomenon can actually be observed in some images, and allows the inclusion of another rule in the algorithm, namely, the labeling of a region as diffuse if the distribution of all the C-values is very small.

The motivation behind the design of the algorithm was to be able to find changes from diffuse to specular regions in images. We have already seen that such changes involve a color shift from the color of the diffuse region to a color similar to that of the illuminant. Therefore the algorithm was designed to detect color shifts between adjacent regions in the image. Interestingly enough, there is another case in which a color shift may take place, and as such, may suggest another use for the algorithm. Such a case is for non-ideal shadows. As described elsewhere [9], non-ideal shadows occur when there is another "light source" in the form of reflection off an object in the scene onto an area of another object which is not directly illuminated by the main light source. In such cases, we have already seen that there is a shift in the color of the object from the weak diffuse reflection to one which is the sum of the weak diffuse and the color cast onto it by another object. Such a process is, as with the highlights, additive, and should result in a "dog-leg" in the C-space. An example of such a shift, taken from the original peppers image, is given in Figure 14.

Finally, there remains one problem with the technique employed in this paper, namely, highlights on materials which share the same spectral distribution as the light source. A simple example can be a highlight caused by a white illuminant on a white object. The problem, of course, is that there will be no color shift, since the diffuse and specular reflections share the same spectral power distribution. This results in the fact that there will be no "dog-leg" in C-space since the only difference between the two regions will be the strength of the values of C_1. This problem remains open at this stage.

8 Summary

This paper discussed the problem of identifying probable highlights in real images. The solution we have proposed is based on the physics of interaction between light and materials, as formulated in models of optics and reflection. Based on these models, a prediction of the behavior of highlights, or the specular reflection, was made, and

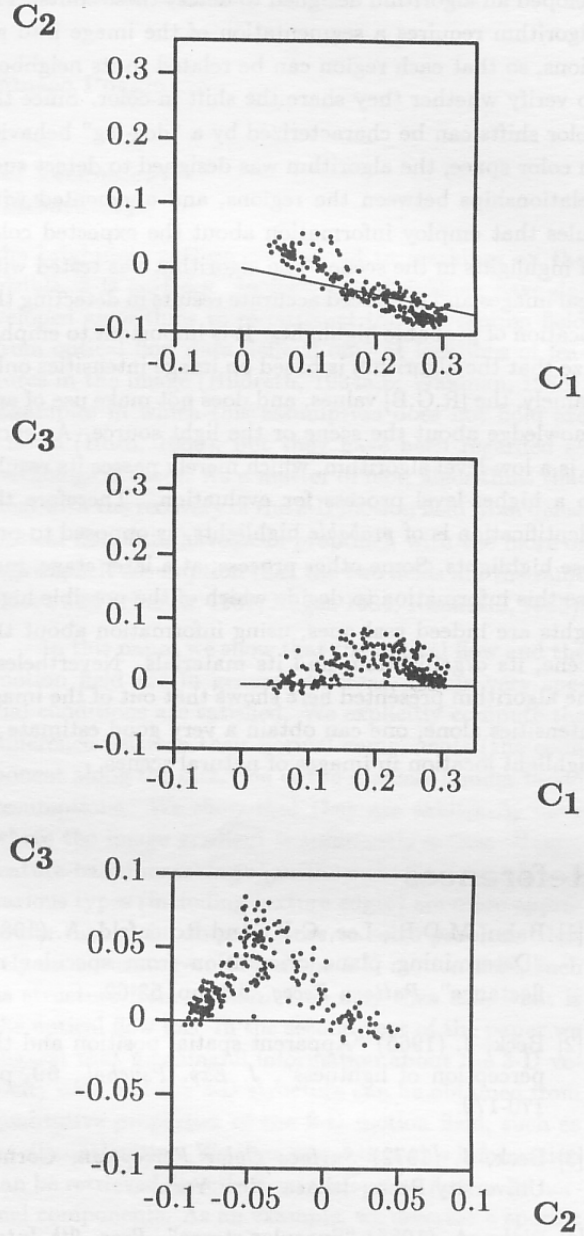

Figure 14: A "dog-leg" relationship in a non-ideal shadow.

The scatter plot is taken from the same area which was considered non-ideal shadow in the original "peppers" image. Note that although the chromaticity of one of the legs (the shadow) points to the direction of a perfect reflector, careful examination of this leg reveals that its C_1 values are decreasing. Therefore it can not be considered a highlight, but perhaps a shadow.

was observed in real images. Based on the observed phenomenon, namely, the shift in the color of objects from the diffuse reflection to the the specular one, we have developed an algorithm designed to detect these shifts. The algorithm requires a segmentation of the image into regions, so that each region can be related to its neighbors to verify whether they share the shift in color. Since the color shifts can be characterized by a "dog-leg" behavior in color space, the algorithm was designed to detect such relationships between the regions, and augmented with rules that employ information about the expected color of highlights in the scene. The algorithm was tested with real images and produced accurate results in detecting the location of probable highlights. It is important to emphasize that the algorithm is based on image intensities only, namely, the [R,G,B] values, and does not make use of any knowledge about the scene or the light source. As such, it is a low-level algorithm, which merely passes its results to a higher-level process for evaluation. Therefore the identification is of *probable* highlights, as opposed to precise highlights. Some other process, at a later stage, may use this information to decide which of the possible highlights are indeed real ones, using information about the scene, its organization, and its materials. Nevertheless, the algorithm presented here shows that out of the image intensities alone, one can obtain a very good estimate of highlight location in images of natural scenes.

References

[1] Babu, M.D.R., Lee, C-H., and Rosenfeld, A. (1985) "Determining plane orientation from specular reflectance", *Pattern Recog.*, **18**, pp. 53-62.

[2] Beck, J. (1965) "Apparent spatial position and the perception of lightness", *J. Exp. Psychol.*, **69**, pp. 170-179.

[3] Beck, J. (1972) *Surface Color Perception*, Cornell University Press: Ithaca, New York.

[4] Blake, A. (1985) "Specular stereo", *Proc. 9th International Joint Conference on Artificial Intelligence*, Los Angeles, California, pp. 973-976.

[5] Buchanan, C.S. (1986) *Determining surface orientation from specular highlights*, M.Sc. Thesis, Dept. of Computer Science — University of Toronto.

[6] Cook, R.L., and Torrance, K.E. (1981) "A reflectance model for computer graphics", Proc. SIGGRAPH 1981, *Computer Graphics*, **15(3)**, pp. 307-316.

[7] Gershon, R. (1985) "Empirical results with a model of color vision", *Proc. of IEEE Conf. on Computational Vision and Pattern Recognition*, San Francisco, California, pp. 302-305.

[8] Gershon, R. (1987) *The Use of Color in Computational Vision*, Ph.D. thesis, Dept. of Computer Science — University of Toronto, forthcoming.

[9] Gershon, R., Jepson, A.D., and Tsotsos, J.K. (1986) "Ambient illumination and the determination of material changes", *J. Opt. Soc. Am. A*, **3**, pp. 1700-1707.

[10] Horn, B.K.P. (1977) "Understanding image intensities", *Artificial Intelligence*, **8**, pp. 201-231.

[11] Horowitz, S.L., and Pavlidis, T. (1974) "Picture segmentation by a directed split-and-merge procedure", *Proc. 2nd International Joint Conference on Pattern Recognition*, Copenhagen, Denmark, pp. 424-433.

[12] Ikeuchi, K. (1981) "Determining surface orientations of specular surfaces by using the photometric stereo method", *IEEE Trans. Pattern Anal. Machine Intell.*, **PAMI-3**, pp. 661-669.

[13] Judd, D.B., and Wyszecki, G. (1975) *Color in Business, Science and Industry*, 3rd Edition, John Wiley and Sons: New York.

[14] Nagel, H-H. (1978) "Formation of an object concept by analysis of systematic time variations in the optically perceptible environment", *Computer Graphics and Image Processing*, **7**, pp. 149-194.

[15] Pavlidis, T. (1986) "Comments on "Low level segmentation: An expert system" ", *IEEE Trans. Pattern Anal. Machine Intell.*, **PAMI-8**, pp. 675-676.

[16] Rubin, J.M, and Richards, W.A. (1984) "Color vision and image intensities: Representing material changes", AI Memo 764, MIT AI Lab.

[17] Shafer, S.A. (1984) "Using color to separate reflection components", Technical Report TR-136, Computer Science Dept. — University of Rochester.

[18] Thorpe, C.E. (1984) "An analysis of interest operators for FIDO", *Proc. of IEEE Workshop on Computer Vision: Representation and Control*, Annapolis, Maryland, pp. 135-140.

[19] Thrift, P., and Lee, C. (1983) "Using highlights to constrain object size and location", *IEEE Trans. Syst., Man, Cybern.*, **SMC-13**, pp. 426-431.

[20] Ullman, S. (1976) "On visual detection of light sources", *Biol. Cybern.*, **21**, pp. 205-212.