SPECULARITY REMOVAL FOR SHAPE FROM SHADING

Frank Tong
Brian V. Funt

School of Computing Science
Simon Fraser University
Burnaby, British Columbia, CANADA V5A 1S6

ABSTRACT

Specularly reflecting surfaces confuse traditional shape-from-shading algorithms because the variation in image intensity within a specularity does not directly relate to the cosine of the incident angle, as it would for a simple Lambertian reflector. To overcome this problem, color is introduced and a method of removing the specular component of the intensity variation is proposed based on a dichromatic model of surface reflection. Unlike Shafer's method for specularity removal, which is restricted to uniformly colored surface patches, our algorithm uses information from several differently colored regions. The specular component due to interface reflection does not change across the regions even though the diffuse component due to body reflection does. In color space, the regions project to planes and the color of the specular component is found as the common intersection of these planes. Once the color of the specular component is known, it is removed from the original image. The resulting image preserves the relative intensity of the diffuse component so it can then be successfully input to a traditional shape-from-shading algorithm.

KEYWORDS:
specularity, surface orientation, shape from shading, Dichromatic Model, color space, Lambertian, interface, body and ambient reflection.

1. Introduction

Intensity shading provides many cues about surface orientation. Obtaining shape from shading information, however, is difficult because the intrinsic scene characteristics are encoded in a single intensity value which may have resulted from an infinite number of combinations of illumination, orientation and reflectance. While the encoding is unique, the decoding is not.

Although the general shape-from-shading problem is ambiguous, it becomes solvable when various constraints about the geometry and photometry of the reflecting surface are exploited. Woodham solves the problem by assuming a priori knowledge of the reflectance map and constraints on the surface curvature [15] [16]. Horn and Brooks derived a numerical scheme [6] [2] for computing local surface orientation from the variation in image intensity. To render the problem solvable without relying on a priori supply of reflectance map, the domain was restricted to scenes containing only smooth surfaces of uniform Lambertian reflectance, illuminated by a single point source at great distance.

While their methods will calculate the shape of uniformly colored Lambertian reflectors, due to the complications in modeling specular reflection component [14] [3], they cannot handle scenes containing specularly reflecting surfaces. Unfortunately, the majority of surfaces are non-Lambertian so specular highlights need to be taken into account for a more generally applicable shape-from-shading algorithm. Babu et. al., in order to determine the orientation of non-Lambertian surfaces, exploit the uniformity in image irradiance along the iso-brightness contours from specular reflection [1]. Their method, however, is insufficient as it is limited to planar surfaces only.

A typical specularity appears on the plastic, multi-colored beach ball shown in Figure 1-1. Just as one would expect, a plot of intensity as a function of position along the indicated line shows a sharp peak within the specularity.

Our approach to the problem of shape from shading for non-Lambertian surfaces is first to calculate the specular component using an algorithm different, but derived, from Shafer's [11] [8] [7]. Once the specular component is known, it can be removed to obtain an image containing only diffuse components. With the specular component no longer confusing the intensity data, the traditional shape-from-shading methods once again become applicable to the recovery of surface shape. Figure 1-2 shows the resulting diffuse-component image calculated by our algorithm. The intensity through the specular region now varies as would be expected of a Lambertian reflector.

Both the diffuse and the specular components are combined in the image intensity variation. Since the

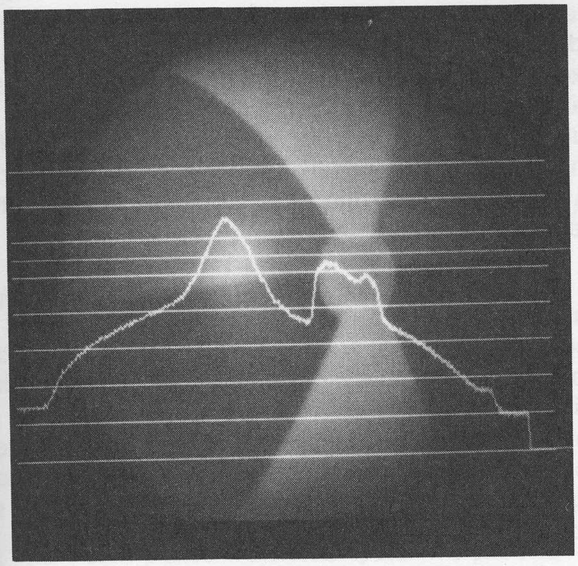

Figure 1-1:
Specular reflection is clearly visible in this black and white photograph of the color image used as one test of the algorithm described in the paper. The plastic beach ball is red, blue, yellow, green and white.

intensity profile is the only information available in a black-and-white picture, from a monochromatic image there is no way to separate the two features without assumptions about their relative spatial frequency. Nevertheless, assuming that the diffuse component is restricted to gradual change inevitably limits the class of surfaces to relatively flat ones. Similarly, assuming that the specular component corresponds to the high spatial frequency component restricts the world to smooth and shiny surfaces that reflect highlights concentrated around the specular direction. It is preferable to be able to uncover the specular component without resorting to restrictive assumptions about relative spatial frequency.

The specular component can be recognized, however, when color image data is introduced. Since specular reflection is a surface interface phenomenon, it bears little relationship to the surface color. Therefore, the specular component should be spectrally uniform across differently colored regions. Thus, in spite of its intensity variation being obscured in the monochromatic image, the spectral characteristics of the specular component can be calculated by comparing the data from different chromatic channels. Using the color of the specular component as a sort of filter, we can filter out all the specularity-colored components of the image leaving an image containing only diffuse components. Both Pellicano [9] and Gershon et. al. [5] address the problem of specularity detection, but neither attempts to remove the specularities once they are found.

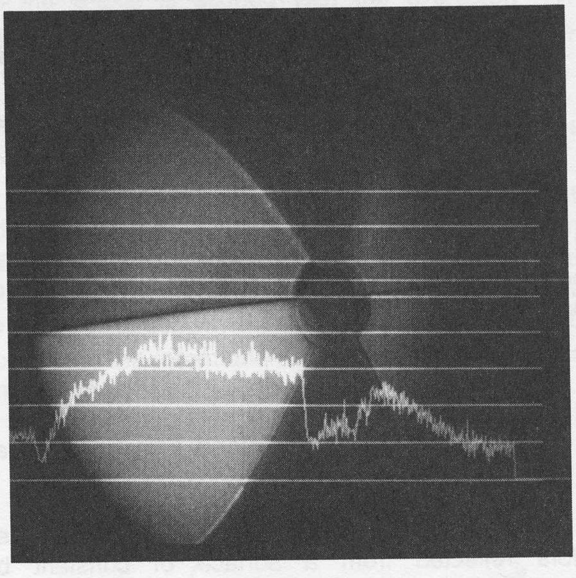

Figure 1-2:
The diffuse-component image of the beach ball shown in Figure 1-1.

We employ Shafer's Dichromatic Model [11] of color reflection. His model is applicable to a wide class of surfaces, namely those of dielectric materials. It describes color reflection as a linear process in which the reflected light is interpreted as a sum of independent reflection components, namely, the ambient, the diffuse and the specular components. Its color is a linear combination of the characteristic colors of the respective components.

Shafer [11] suggested an algorithm for computing specular and diffuse reflection component intrinsic images. This algorithm has since been implemented and tested [8] [7] on some images of real scenes in which all ambient background illumination was eliminated by the use of black curtains on the walls. In addition to not handling ambient illumination, one of the algorithm's main drawbacks is that it requires prior image segmentation into uniformly colored regions. Our new algorithm accounts for ambient illumination and avoids a priori segmentation. We use information from multiple regions of different color to increase the reliability of the specularity color that is calculated.

2. Shafer's Algorithm

Shafer [11] [12] addressed the problem of using color to separate the specular and diffuse image components. He described reflection as a linear process in which the reflected light is a sum of independent reflection components each corresponding to light of a specific color whose intensity is modulated by the imaging geometry. For a fixed geometry, Shafer arrived at the irradiance equation:

$$\mathbf{L}(\lambda) = m_i \mathbf{C}_i(\lambda) + m_b \mathbf{C}_b(\lambda) + \mathbf{C}_a(\lambda)$$

where:

$\mathbf{L}(\lambda)$ is the reflected light;

$\mathbf{C}_i(\lambda)$, $\mathbf{C}_b(\lambda)$ and $\mathbf{C}_a(\lambda)$ are the characteristic colors of the reflection components due to the interface reflection, body reflection and the ambient illumination;

m_i and m_b are the scaling factors due to the imaging geometry.

Applying the linearity of spectral projection, Shafer [11] [12] [10] mapped the light mixture into *RGB* coordinates. The irradiance equation encoded in color space parameters then reads:

$$\mathbf{c} = m_i\,\mathbf{c}_i + m_b\,\mathbf{c}_b + \mathbf{c}_a$$

In other words, the reflected color is a linear combination of the three color vectors: $\mathbf{c}_i$, $\mathbf{c}_b$ and $\mathbf{c}_a$.

From this equation it can be seen that the set of colors reflected from a surface of constant body reflectance will lie on a parallelogram in color space. The parallelogram is defined by the two characteristic vectors $\mathbf{c}_i$ and $\mathbf{c}_b$, and is displaced from the origin by $\mathbf{c}_a$.

Based on this observation, Shafer suggested an algorithm for computing the intrinsic images, m_i and m_b, of an imaged surface for the scenes in which there is no ambient illumination. The algorithm involves four steps:

1. Histogram the pixel values in color space.
2. Fit a plane to these points, with the restriction that the plane must pass through the origin.
3. Fit a parallelogram to this plane having its lowest corner at the origin. The sides are $\mathbf{c}_i$ and $\mathbf{c}_b$ respectively.
4. At each pixel, express its color as a linear combination of $\mathbf{c}_i$ and $\mathbf{c}_b$. The coefficients of the combination are the values for m_i and m_b.

Although not included in the implementation [8] [7], it should be possible to extend the algorithm to cope with ambient illumination as well except that the parallelogram will no longer contain the origin.

Shafer's algorithm is very nice. However, in addition to the fact that most scenes contain more than one surface color, fitting the parallelogram to the set of points is not always easy. Shafer notes that the distribution of pixel values within the parallelogram must not be "pathological" if the correct parallelogram is to be found. What this means is that most of the pixel points must be close to either the $\mathbf{c}_i$ or the $\mathbf{c}_b$ axis. Very shiny surfaces are ideal because their highlights are sharp and well-defined. The reflection is dominated by the specular component within the highlight and by the diffuse component outside of it. However, in practice, many surfaces are rough and have broad highlights. Those highlights are neither sharp nor well-defined and this leads to a "pathological" distribution of pixel values. In addition, for enough points to lie close to the $\mathbf{c}_i$ and $\mathbf{c}_b$ axes, the surface must have significant orientation variation. However, it is common that only part of a surface is imaged, or very often, a surface will not have enough orientation variation to yield a sufficient number of pixels lying near the $\mathbf{c}_i$ and $\mathbf{c}_b$ axes to define them clearly.

Shafer's algorithm computes all four of m_i, m_b, $\mathbf{c}_i$ and $\mathbf{c}_b$. As our concern here is the removal of the specular component for successful shape from shading, we do not necessarily require all four of these values. With $\mathbf{c}_i$ alone, we can remove the specularities and obtain the diffuse component required for the shape-from-shading calculation. The next section describes an algorithm that takes advantage of the fact that only $\mathbf{c}_i$ is required.

3. An Algorithm for Specularity Removal

Regions of different reflectance reflecting light of different color will sweep out different planes in color space. The reflected colors, however, share a common specular component. The specular component arises from interface reflection, so it is not affected by the spectrally-biased body reflectance; and thus every region should reflect specularities of the same color. Since the specular component color is common to all the regions, it will similarly be common to all the planes swept out in color space. Therefore as shown in Figure 3-1, it can be calculated as a line parallel to the lines generated by pairwise intersection of the color planes formed by the different regions. As a result, instead of segmenting the image into different color regions and using each to compute its $\mathbf{c}_i$, $\mathbf{c}_b$, m_i, and m_b -- as Shafer's algorithm does -- we combine the information across regions to compute just $\mathbf{c}_i$, the color of the common specular component. Once $\mathbf{c}_i$ is obtained, it is used to filter out the specular component from the original image.

In the context of trying to establish the illuminant chromaticity from specularities for use in color constancy, D'Zmura et. al. [4] also observed that the illuminant color could be found as the intersection of color planes from different regions, but then abandoned it on the basis that it required that the scene already be segmented. As we show below, however, a global a priori segmentation of the scene is not required. Instead, our algorithm calculates the line that best approximates the common intersection of a large number of planes fit to small local neighborhoods.

In order to intersect the color planes, they must first be found. At each pixel, we pick a neighborhood ρ containing pixels which are likely to belong to a single region. The image of ρ projected into color space will thus tend to reside on a single plane. Furthermore, all the pixels in that dichromatic region will in principle

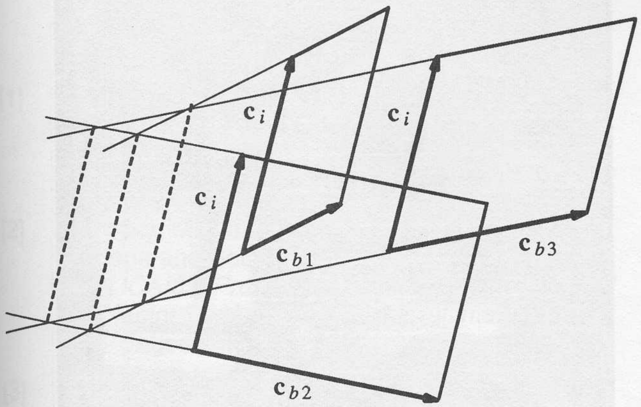

Figure 3-1:
The color planes in the *RGB* space. The pairwise intersections of the planes, as represented by the dashed lines, all lie parallel to $\mathbf{c}_i$.

project to the same plane. So, at each pixel in a dichromatic region, if we pick its neighbors and map them to color space, a plane fitted to them will approximate the global plane corresponding to the entire region [13].

The specularity color vector, $\mathbf{c}_i$, is defined by the common intersection of the planes obtained by the above process. However, instead of intersecting the planes directly, we find the line which is most parallel to all the planes. To find this line, the normal to each plane is calculated and then a least-squares fit is used to find the line closest to perpendicular to all the plane normals [13]. Thus, in order to get the intersection of the color planes, we do not need to do a general region segmentation.

Once $\mathbf{c}_i$ has been found, we can proceed to remove it from the image. As discussed before, the reflected color is considered to be a linear composition of three components.

$$\mathbf{c} = m_i\mathbf{c}_i + m_b\mathbf{c}_b + \mathbf{c}_a$$

By the linearity property of color composition, we can decompose the color terms in the equation into two orthogonal parts: one, $\mathbf{c}_i^{\|}$ is parallel to $\mathbf{c}_i$, the other, $\mathbf{c}_i^{\perp}$, is perpendicular to it.

$$\mathbf{c}^{\|} + \mathbf{c}^{\perp} = m_i\mathbf{c}_i + m_b(\mathbf{c}_b^{\|} + \mathbf{c}_b^{\perp}) + \mathbf{c}_a^{\|} + \mathbf{c}_a^{\perp}$$

This decomposition of $\mathbf{c}$ is shown in Figure 3-2.

A projection to a plane orthogonal to $\mathbf{c}_i$ eliminates all the $\mathbf{c}_i^{\|}$ terms, resulting in an image with the specular component removed.

$$\mathbf{c}^{\perp} = m_b\mathbf{c}_b^{\perp} + \mathbf{c}_a^{\perp}$$

We carry out the projection by rotating the *RGB* axes so that the *B*-axis is aligned with the $\mathbf{c}_i$ vector. Then projection onto the new *RG*-plane yields the $\mathbf{c}_i^{\perp}$ terms with the $\mathbf{c}_i^{\|}$ components eliminated [13].

Note that while the specular component has been removed so have $\mathbf{c}_a^{\|}$ and $\mathbf{c}_b^{\|}$. This means that the colors in the image formed from the $\mathbf{c}_i^{\perp}$ component of each pixel will not be the same as if the objects simply had not had specularities in the first place. Nonetheless, the $\mathbf{c}_i^{\perp}$ component does preserve the relative diffuse intensity variation, which is proportional to m_b, so it encodes all the information that is necessary for shape from shading.

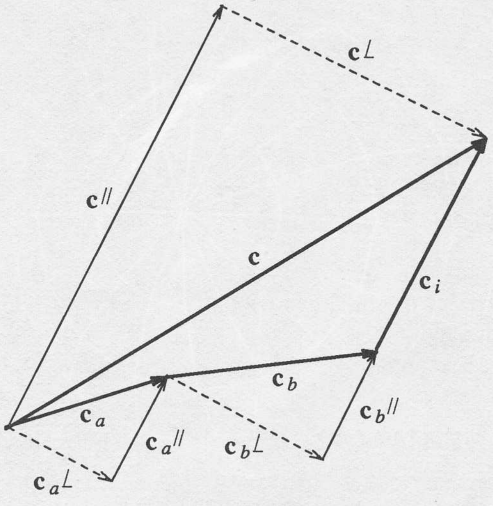

Figure 3-2:
The color components of $\mathbf{c}$. The dashed lines represent the $\mathbf{c}_i^{\perp}$ projection.

Although color was introduced in order to provide additional information about the specular component of the image intensity variation, it also adds an extra dimension to the complexity of the analysis. In color images -- or for that matter black-and-white images of scenes containing multi-colored objects -- the reflectance edges that arise will confuse the standard shape-from-shading algorithms. Reflectance edges cause intensity discontinuities that are unrelated to changes in surface orientation.

Like Horn et. al. [6], we approach shape from shading as a variational problem. In solving for the local surface orientation, we assume we have smooth surfaces, constant reflectance within a region of a single color, and invariant illumination. Since reflectance edges do not represent a change in surface orientation the usual surface smoothness constraint is sufficient to propagate surface orientation information across them [13].

4. Results

Color pictures were taken using a CCD camera with an infra-red cutoff filter and color separation Wratten filters numbers 25, 58 and 47B. Scenes were illuminated by a single tungsten bulb and uniform background illumination from flourescent lighting.

A color version of Figure 1-1 was used to calculate the surface orientations of the local planes. Figure 4-1 shows the surface orientations of these planes represented as points on a Gaussian sphere. That they fall nicely on a geodesic of the sphere shows that they share a common intersection line. Once this intersection was calculated, the c_i^{ll} component was removed from the color version of Figure 1-1 to produce Figure 1-2.

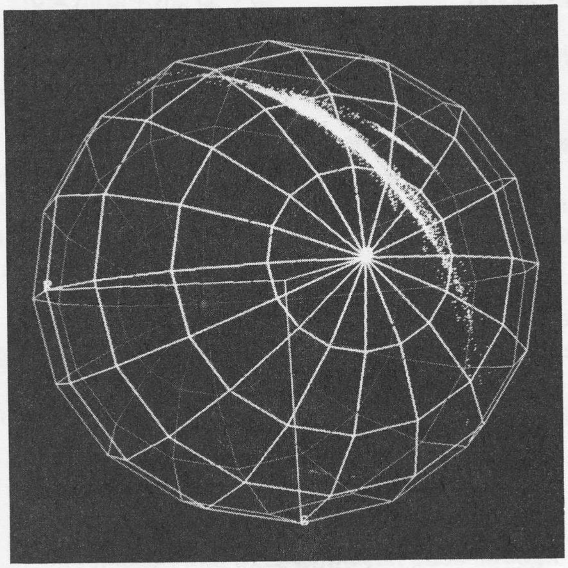

Figure 4-1:
The gaussian sphere on which the local plane normals calculated from the color version of Figure 1-1 are plotted.

Similar results were obtained for the scene of Figure 4-2 containing a blue leather purse, a red plastic cup (standing) and a yellow plastic cup (lying). A color version of Figure 4-2 was processed producing Figure 4-3.

The intensity plots in both Figure 1-2 and Figure 4-3 show a considerable degree of noise. This results from the fact that most of the dynamic range in the input images is used up by the specularities. When they are removed, the remaining signal-to-noise ratio is low. The resulting image intensities have been scaled so that the images do not appear too dark and this of course amplifies the noise.

5. Conclusion

We have addressed the problem of determining shape information from intensity variation in images of scenes containing non-Lambertian reflecting surfaces. The specular reflection component that arises can be removed from the images by use of the dichromatic model of reflection. Because it arises from interface reflection, the specular component is common to surfaces of different color. When the pixel values are plotted in color space, those from the same region lie on a plane in color space. The intersection of the color planes defined by different regions yields the color of the specular component of the reflected light. Traditional shape-from-shading algorithms can then be applied to the component of the reflected light having a color vector perpendicular to the specular color since it preserves the relative intensity of the body reflection.

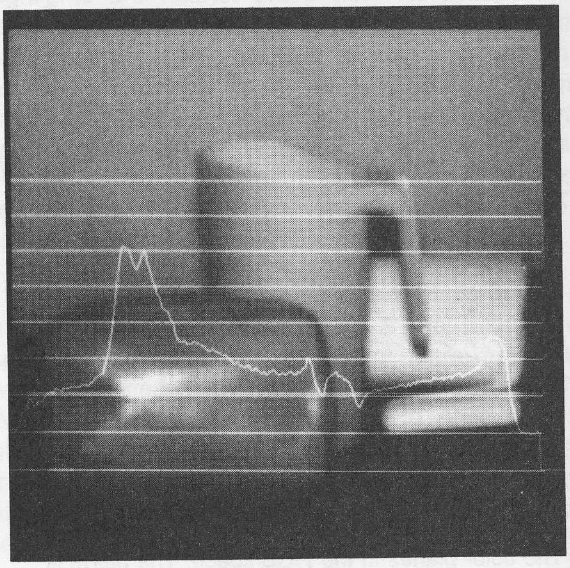

Figure 4-2:
The intensity map of a scene of a blue leather purse, a red plastic cup (standing) and a yellow plastic cup (lying).

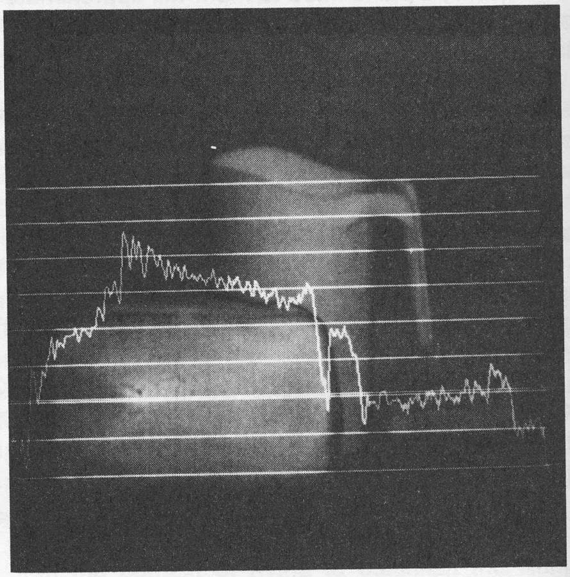

Figure 4-3:
The diffuse-component image of Figure 4-2.

Acknowledgement

This work was supported by NSERC grant 4322.

References

[1] Babu, M.D.R., Lee, C.H. and Rosenfeld, A. Determining Plane Orientation from Specular Reflectance. *Pattern Recognition* 18(1):53-62, 1985.

[2] Brooks, M.J. and Horn, B.K.P. Shape and Source from Shading. In *IJCAI85, Proceedings of the 9th International Joint Conference on Artificial Intelligence*, pages 932-936. 1985.

[3] Cook, R.L. and Torrance, K.E. A Reflectance Model for Computer Graphics. *ACM Transactions on Graphics* 1(1):7-24, 1982.

[4] D'Zmura, M. and Lennie, P. Mechanisms of Color Constancy. *Journal of Optical Society of America A* 3(10):1662-1672, 1986.

[5] Gershon, R., Jepson, A.D. and Tsotsos, J.K. Highlight Identification Using Chromatic Information. In *Proceedings of the First International Conference on Computer Vision*, pages 161-170. London, England, 1987.

[6] Horn, B.K.P. and Brooks, M.J. The Variational Approach to Shape from Shading. *Computer Vision, Graphics and Image Processing* 33:174-208, 1986.

[7] Klinker, G.J., Shafer, S.A. and Kanade, T. The Measurement of Highlights in Color Images. *International Journal of Computer Vision* (to appear),

[8] Klinker, G.J., Shafer, S.A. and Kanade, T. Using A Color Reflection Model to Separate Highlights From Object Color. In *Proceedings of the First International Conference on Computer Vision*, pages 145-150. London, England, 1987.

[9] Pellicano, P.N.E. Detection of Specularities in Color Images Using Local Operators. In *Proceedings of the Graphics Interface / Vision Interface 1986*, pages 370-374. Vancouver, B.C., 1986.

[10] Shafer, S.A. Describing Light Mixtures through Linear Algebra. *Journal of Optical Society of America* 72(2):299-300, 1982.

[11] Shafer, S.A. *Using Color to Separate Reflection Components.* Technical Report, Computing Science Department, University of Rochester, 1984.

[12] Shafer, S.A. *Optical Phenomena In Computer Vision.* Technical Report TR-135, Computing Science Department, University of Rochester, 1984.

[13] Tong, Frank. *Specularity Removal for Shape from Shading.* Technical Report TR-88-1, School of Computing Science, Simon Fraser University, 1988.

[14] Torrance, K.E. and Sparrow, E.M. Theory for Off-specular Reflection from Roughened Surfaces. *Journal of Optical Society of America* (9):1105-1114, 1967.

[15] Woodham, R.J. A Cooperative Algorithm for Determining Surface Orientation from a Single View. In *IJCAI77, Proceedings of the 5th International Joint Conference on Artificial Intelligence*, pages 635-641. 1977.

[16] Woodham, R.J. Relating Properties of Surface Curvature to Image Intensity. In *IJCAI79, Proceedings of the 6th International Joint Conference on Artificial Intelligence*, pages 971-977. 1979.