Colored Interreflections and Shape Recovery

Shree K. Nayar and Yitao Gong

Department of Computer Science, Columbia University, New York, N.Y. 10027

Abstract

Interreflections cause recovery algorithms to extract erroneous estimates of surface shape and color. Due to the complex nature of the interreflection phenomenon, its effects are difficult to analyze in a multi-dimensional color space. Noting that the interreflections of one wavelength are in general unaffected by those of any other wavelength, we decompose the interreflection process into three independent processes corresponding to the three spectral bands of color images. Photometric stereo is applied independently applied to each spectral band and recovered (pseudo) shapes and color are related to the actual shape and color of the surface. An algorithm is proposed that recovers the actual shape and color of the surface from the pseudo estimates. The accuracy and robustness of the algorithm is demonstrated on a variety of colored and multi-colored surfaces.

1 Interreflections and Shape Recovery

Points in the scene, when illuminated, reflect light not only in the direction of the sensor but also between themselves. This is always true with the exception of scenes that consist of only a single convex surface, in which case, no two points on the surface are visible to one another. In general, however, scenes concavities and points in the scene reflect light between themselves. These *interreflections* can appreciably alter the appearance of the scene. Existing vision algorithms do not account for the effects of interreflections and hence often produce erroneous results.

Two separate problems associated with interreflections can be identified; the *forward* (graphics) problem and the *inverse* (vision) problem. Most of the previous work done in this area is related to the forward problem. The forward problem involves the prediction of image brightness values given the shape and reflectance of a scene. Horn [Horn 70] discussed the changes in image intensities due to interreflections caused by polyhedral surfaces that are Lambertian in reflectance. Koenderink and van Doorn [Koenderink 83] formalized the interreflection process for Lambertian surfaces of arbitrary shape and varying reflectance (albedo). They proposed a solution to the forward problem in terms of the eigenfunctions of the interreflection kernel. Cohen and Greenberg [Cohen 85] modeled the scene as a finite collection of Lambertian planar facets and proposed a radiosity solution to the forward problem and used it to render images for graphics. Later, Forsyth and Zisserman [Forsyth 89] used a similar numerical solution to the forward problem to compare predicted and measured image intensities.

More recently, Nayar et.al [Nayar 91] demonstrated the effects of interreflections on shape-from-intensity algorithms, such as, shape-from-shading [Horn 70], photometric stereo [Woodham 78]. These are algorithms that recover three-dimensional shape information from image intensities. All shape-from-intensity methods, are based on the assumption that points in the scene are illuminated only by the sources of light and not other points in the scene; interreflections are assumed *not* to exist. As a result, these methods produce erroneous results when applied to concave surfaces. Nayar et.al analyzed the incorrect shape, the pseudo shape, recovered by shape-from-intensity methods when applied to concave Lambertian surfaces. They established a relation between the actual shape and pseudo shape and developed an algorithm that recovers the actual surface from the pseudo shape and reflectance estimates.

In their analysis, Nayar et.al [Nayar 91] focused on gray surfaces; each point on the Lambertian surface was assumed to have a constant albedo value that is independent of the wavelength of incident light. In the case of colored surfaces, however, the reflectance of a surface point is dependent on the spectral distribution of the incident light. In the case of concave surfaces, the spectral distribution of light rays incident on a given surface point depend on the spectral characteristics of both the source as well as other points on the surface. Further, the light reflected by the surface point also depends on its own spectral reflectance properties. Interreflections in the case of colored surfaces, therefore, can cause the color of one surface region to bleed onto another. These effects have been noted by other researchers [Bajcsy 89], [Brill 89], [Novak 90], and [Drew 90]. The common approach is to analyze these interreflection effects in a multi-dimensional color space. Since, interreflections are influenced by several factors including the unknown shape and color of the surface, the analysis in color space is difficult.

Figure 1 shows the effects of colored interreflections on photometric stereo. The surface has three regions of different color. The actual color of the surface points are seen to cluster at three points in the color space. The photometric

stereo method is applied independently to the three bands of the color image. As seen from the figure, three different shapes (all incorrect) are recovered from the three bands. The recovered color is also erroneous; it is not constant within each constant-color region of the surface.

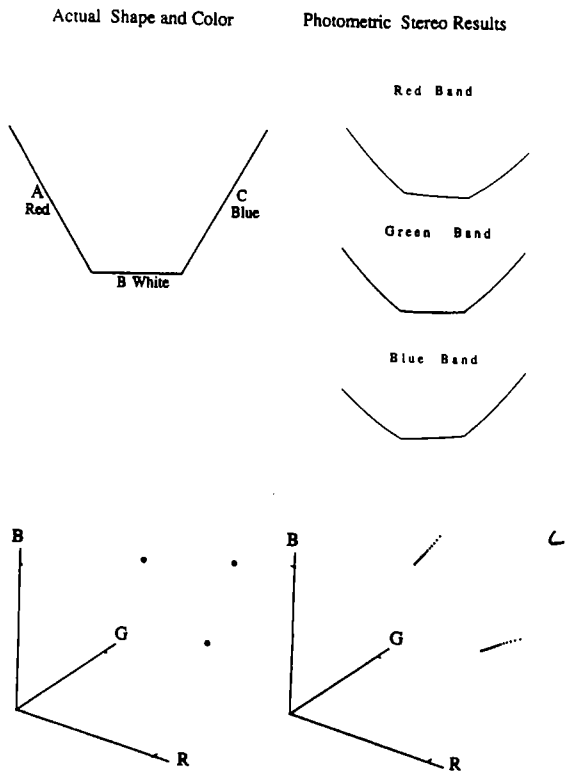

Figure 1: The actual shape and color of a surface and its shape and color computed by photometric stereo.

In this paper, we extend the results in [Nayar 91] to colored and multi-colored surfaces. Noting that the three bands of a color image correspond to different wavelengths of light, we decompose the interreflection process into three independent processes. We then analyze the pseudo shape and pseudo albedo function produced by photometric stereo in each band. Then, we establish a relation between the pseudo shape and color and the actual shape and color of a surface. Using this relation, an algorithm is given that recovers the actual shape and color of the surface from the pseudo estimates. The algorithm is tested on a variety of multi-colored surfaces. The results demonstrate the robustness and accuracy of the algorithm.

2 A Diffuse Interreflection Model

Our solution to the inverse interreflection problem is based on the solution to the forward problem, i.e. modeling interreflections for diffuse surfaces of known shape and reflectance. The interreflection model described here is based on the formulation proposed by Koenderink and van Doorn [Koenderink 83]. All surfaces in the scene are assumed to be **Lambertian**. We will shortly see that this assumption is necessary to obtain a closed form solution to the forward interreflection problem. The Lambertian surface can have any arbitrary shape and varying reflectance, i.e. albedo (ρ) may vary from one surface point to the next. Here, we will assume that the incident light is **monochromatic**, i.e the surface reflects and interreflects light rays of a single wavelength. Hence each surface point can be assumed to have a constant reflectance coefficient (albedo). Later we show that the interreflection model developed here for monochromatic light can be extended and used for colored surfaces.

Consider the concave surface shown in Figure 2. The surface is divided into m infinitesimal facets. Let $\mathbf{x}_i$ and dx_i represent the three-dimensional coordinates and the surface area of the i^{th} facet, respectively. The radiance (brightness) and albedo values of each facet are assumed to be constant over the entire facet and equal to the radiance and albedo values at the center point $\mathbf{x}_i$ of the facet, i.e. $L_i = L(\mathbf{x}_i)$ and $\rho_i = \rho(\mathbf{x}_i)$. Consider the two facets i and j. The radiance of the facet i due to the radiance of the facet j is determined using basic radiometric definitions [Nicodemus 77] as:

$$L_i = \frac{\rho_i}{\pi} K_{ij} L_j \tag{1}$$

where the factor K_{ij} is given by:

$$K_{ij} = \frac{[\mathbf{n}_i \cdot \mathbf{r}_{ij}]\,[\mathbf{n}_j \cdot \mathbf{r}_{ji}]}{[\mathbf{r}_{ij} \cdot \mathbf{r}_{ij}]^2} \cdot dx_j \tag{2}$$

K_{ij} is a function of the relative positions and the orientations of the two facets; it determines the interreflections between i and j from a purely geometrical perspective. It is referred to as the interreflection **kernel**.

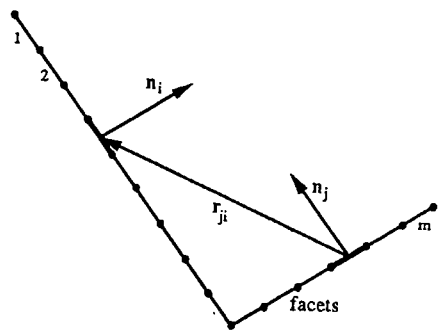

Figure 2: Modeling the surface as a collection of facets, each with its own radiance and albedo values.

Now, let us consider the entire surface shown in Figure 2. Assume the surface to be illuminated by a distant point source of light. Rays of light that impinge upon the surface are reflected between the facets. Since the albedo of

each facet is less than unity, the rays of light lose a fraction of their energy with each bounce. Eventually, after possibly an infinite number of bounces, the radiance at each point on the surface converges to a final steady-state value. Hence, the radiance of the facet i may be expressed as a sum of the radiance due to direct illumination from the source and the radiance due to the final radiance values of other facets on the surface:

$$L_i = L_{si} + \frac{\rho_i}{\pi} \sum_{j \neq i} L_j K_{ij} \quad (3)$$

where L_{si} is the radiance due to direct illumination from the source and the summation term corresponds to the radiance due to mutual illumination. This is the *interreflection equation*. Note that the radiance values L_i and L_j are assumed to be constants in the above equation. It is important to note that this assumption is valid only for Lambertian surfaces; the radiance of a Lambertian surface element is independent of the viewing direction.

The interreflection equation for the complete surface can be written using vector notation. We define the facet radiance vector as $\mathbf{L} = [\, L_1, L_2, ..., L_m \,]^T$ and the source contribution vector as $\mathbf{L_S} = [\, L_{s1}, L_{s2}, ..., L_{sm} \,]^T$. We also define the albedo matrix $\mathbf{P}$ and the kernel matrix $\mathbf{K}$ as:

$$\mathbf{P} = \frac{1}{\pi} \begin{bmatrix} \rho_1 & 0 & \ldots & 0 \\ 0 & \rho_2 & 0 & \ldots 0 \\ \ldots & \ldots & \ldots & \ldots \\ \ldots & \ldots & \ldots & \ldots \\ 0 & 0 & \ldots & \rho_m \end{bmatrix} \quad \mathbf{K} = \begin{bmatrix} 0 & K_{12} & \ldots & \ldots \\ K_{21} & 0 & \ldots & \ldots \\ \ldots & \ldots & 0 & \ldots \\ \ldots & \ldots & \ldots & 0 \ldots \\ \ldots & \ldots & \ldots & 0 \end{bmatrix} \quad (4)$$

Then, equation 3 may be written as:

$$\mathbf{L} = \mathbf{L_S} + \mathbf{PKL} \quad (5)$$

or:

$$\mathbf{L} = (\mathbf{I} - \mathbf{PK})^{-1} \mathbf{L_S} \quad (6)$$

where $\mathbf{I}$ is the identity matrix. Thus, we have obtained a non-iterative, closed-form solution to the forward interreflection problem. The kernel and albedo matrices are determined by the shape and reflectance of the surface, respectively. The source direction and intensity may be used to compute the source contribution vector $\mathbf{L_S}$. Then the radiance of the surface facets, $\mathbf{L}$, can be determined using the above equation.

3 The Pseudo Shape

From equation 6 it is clear that surface radiance values are affected by interreflections. This indicates that if a shape-from-intensity method is applied to a concave surface it is expected to produce erroneous estimates of shape. In order to generalize the inverse interreflection problem, we assume that the reflectance of the Lambertian surface is also *unknown* and may vary from point to point. Therefore, by the term shape-from-intensity, we mean local methods that extract both shape (orientation) and reflectance (albedo) information. Photometric stereo [Woodham 78] is an example of such a shape-from-intensity method. In the presence of interreflections, photometric stereo extracts erroneous shape as well as erroneous reflectance estimates. We refer to the extracted shape as the pseudo shape and the extracted reflectance as the pseudo reflectance of the surface. In this section, we again assume that the light sources used to illuminate the surface are monochromatic. Under this assumption, we investigate how the pseudo shape and reflectance are related to the actual shape and reflectance of the surface. In the next section, we extend these results to colored interreflections.

Once again, consider the surface comprised of m facets (Figure 2). The i^{th} facet may be mathematically represented as:

$$\mathbf{N}_i = \frac{\rho_i}{\pi} \mathbf{n}_i \quad (7)$$

where $\mathbf{n}_i = [n_{xi}, n_{yi}, n_{zi}]^T$ is the unit surface normal and ρ_i is the albedo value for the facet. Therefore, the term "facet" represents both local orientation as well as local reflectance information. The complete surface is then defined by the *facet matrix* $\mathbf{F} = [\, \mathbf{N}_1, \mathbf{N}_2, ..., \mathbf{N}_m \,]^T$. Consider, once again, the interreflection equation given by equation 6. Since the surface is Lambertian, the source contribution vector $\mathbf{L_S}$ may be determined from the facet matrix $\mathbf{F}$ and the source direction vector $\mathbf{s} = [s_x, s_y, s_z]^T$ as:

$$\mathbf{L_S} = \mathbf{F}\,\mathbf{s} \quad (8)$$

Hence, we obtain:

$$\mathbf{L} = (\mathbf{I} - \mathbf{PK})^{-1} \mathbf{F}\,\mathbf{s} \quad (9)$$

We define the matrix $\mathbf{F}_p$ as:

$$\mathbf{F}_p = (\mathbf{I} - \mathbf{PK})^{-1} \mathbf{F} \quad (10)$$

Note that $\mathbf{F}_p$ has the same dimensions as the facet matrix $\mathbf{F}$. In fact, in the absence of interreflections, $\mathbf{K}$ is a null matrix and $\mathbf{F}_p = \mathbf{F}$. In the presence of interreflections, $\mathbf{F}_p$ may be viewed as representing another Lambertian surface whose shape and reflectance differ from those of $\mathbf{F}$. Therefore, if photometric stereo is applied to the concave surface, the extracted shape and reflectance is $\mathbf{F}_p$ and not the actual shape and reflectance given by $\mathbf{F}$. We refer to $\mathbf{F}_p$ as the *pseudo facet matrix*; it represents the pseudo shape and pseudo reflectance that are extracted in the presence of interreflections.

In the case of photometric stereo, three different source directions, $\mathbf{s}_1$, $\mathbf{s}_2$, and $\mathbf{s}_3$, are used sequentially to illuminate the surface. The three resulting surface radiance vectors $\mathbf{L}_1$, $\mathbf{L}_2$, and $\mathbf{L}_3$ may be expressed as:

$$[\mathbf{L}_1, \mathbf{L}_2, \mathbf{L}_3] = \mathbf{F}_p . [\mathbf{s}_1, \mathbf{s}_2, \mathbf{s}_3] \quad (11)$$

The pseudo facet matrix is computed as:

$$\mathbf{F}_p = [\mathbf{L}_1, \mathbf{L}_2, \mathbf{L}_3].[\mathbf{s}_1, \mathbf{s}_2, \mathbf{s}_3]^{-1} \quad (12)$$

The i^{th} *pseudo facet* in $\mathbf{F}_p$ may be written as:

$$\mathbf{N}_{p_i} = \frac{\rho_{p_i}}{\pi} \mathbf{n}_{p_i} \quad (13)$$

where $\mathbf{n}_{p_i}$ and ρ_{p_i} are the pseudo surface normal and the pseudo albedo for the facet i and, in the presence of interreflections, differ from the actual surface normal and actual albedo of the facet.

We conclude this section by highlighting three important properties of the pseudo shape and reflectance:

- The pseudo shape and reflectance are *illumination invariant*. In equation 10, note that the albedo matrix $\mathbf{P}$, the kernel matrix $\mathbf{K}$, and the actual facet matrix $\mathbf{F}$ are all invariant to the direction and intensity of the illumination. As a result, the matrix $\mathbf{F}_p$ is also illumination invariant. It is independent of sources directions used by the shape-from-intensity method to illuminate the surface.

- The pseudo shape and reflectance are *unique*. From equation 10 we see that the pseudo facet matrix $\mathbf{F}_p$ is dependent on the actual facet matrix $\mathbf{F}$, the albedo matrix $\mathbf{P}$, and the kernel matrix $\mathbf{K}$. Note that $\mathbf{P}$ and $\mathbf{K}$ are in turn determined by $\mathbf{F}$. Hence, $\mathbf{F}_p$ is dependent only on $\mathbf{F}$. In other words, there exists only a single pseudo shape and pseudo reflectance corresponding to a given actual shape and reflectance.

- The pseudo shape tends to be *less concave* than the actual shape of the surface. A proof of this property is provided in [Nayar 90]. Figure 3 illustrates this property through a few examples of actual shapes and pseudo shapes All the surfaces are assumed to have a constant albedo value, $\rho = 0.95$. The pseudo shapes are computed using equation 10 and are seen to be less concave than the actual shapes. Figure 4 shows that the pseudo shape gets less concave as albedo increases.

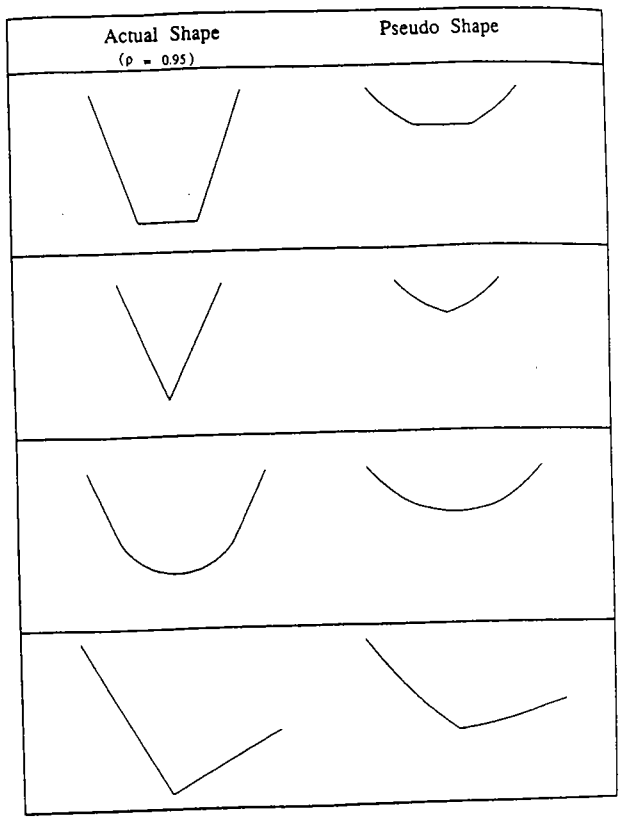

Figure 3: A few actual shapes (with $\rho = 0.95$) and their pseudo shapes.

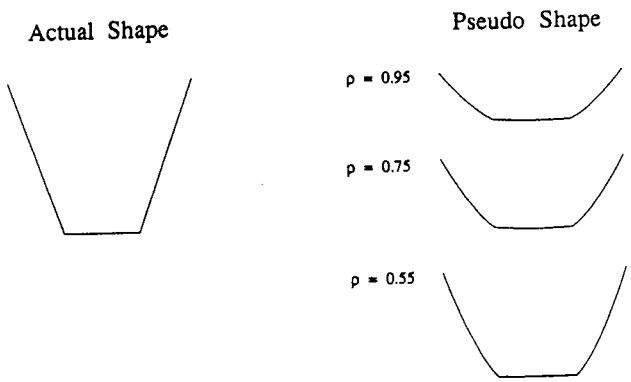

Figure 4: The difference between the actual and pseudo shapes increases with surface albedo.

4 Colored and Multi-Colored Surfaces

While developing the interreflection model (Section 2), we assumed a given surface facet has a constant albedo value. This assumption is only valid under either one of the following two conditions. (a) The surfaces are not colored but rather of different shades of gray and they reflect all wavelengths of incident light equally without attenuating some wavelengths more than others. (b) The incident light is monochromatic and therefore only light rays of a single wavelength are reflected and interreflected by the surface points. In the second case, since we are only concerned with a single wavelength, each point on the surface may be assumed to have a constant albedo value, namely, the albedo value for the given wavelength of incident light.

In the case of colored surfaces, the albedo of a facet would also depend on the spectral distribution of the incident light. In fact, with each bounce the spectral distribution of the reflected light would be altered depending on the spectral characteristics of the reflecting surfaces. Consider, for example, a concave surface that has two regions of different colors, say, red and blue. When the surface is illuminated, the red and blue regions interreflect incident light. If the surface is illuminated using white light, the blue region would reflect only blue light onto the red region. Therefore, a red surface point receives light rays of different spectral content; it receives white light from the source, blue light from the blue region, and red light from other red surface points. Further, the spectral distribution of light reflected by the red surface point depends not only on the spectral distribution of light received by it but also its own spectral characteristics (reflectance properties). As a result of the above effect, the color of a given surface region may be altered due to reflections from neighboring regions that have different colors. This phenomenon is often referred to as *color bleeding*. The reader is directed to [Bajcsy 89], [Brill 89], [Novak 90], and [Drew 90] for a more detailed discussion on this effect.

In the case of colored surfaces, therefore, our assumption that each surface point has a constant albedo value is no longer valid: The albedo of a point would depend on its color and the spectral distribution of the incident light, where incident light includes light received from other surface points. However, from condition (b) we know that, for a given wavelength, a surface point (irrespective of its color) can be assumed to have a unique albedo value. Therefore, pseudo shape and reflectance estimates for a multi-colored surface can be computed by using a narrow-band filter at the sensor end so that only light waves of a particular wavelength are detected by the sensor. Note that the pseudo shape and reflectance will in general be different for different wavelengths.

Typically, color images are obtained using three narrow-band filters at the sensor end. These three filters have narrow-band spectral responses in the **red**, **green**, and **blue** regions of the visible light spectrum. We will refer to the three images produced as the red, green, and blue bands of the color image. Since the three filters are narrow-band filters, the image produced by each filter represents reflections and interreflections of almost a single wavelength of light. Hence, a pseudo shape and pseudo albedo function for the surface can be computed using each band of the color image. The pseudo shapes and albedo functions for the three band are expected to be different since each point may have a different actual albedo for the three different wavelengths of light. Hence, the pseudo facet matrices for the three band may be written as:

$$\begin{aligned} \mathbf{F}_p{}^{\mathrm{R}} &= (\mathbf{I} - \mathbf{P}^{\mathrm{R}}\mathbf{K})^{-1}\mathbf{F}^{\mathrm{R}} \\ \mathbf{F}_p{}^{\mathrm{G}} &= (\mathbf{I} - \mathbf{P}^{\mathrm{G}}\mathbf{K})^{-1}\mathbf{F}^{\mathrm{G}} \\ \mathbf{F}_p{}^{\mathrm{B}} &= (\mathbf{I} - \mathbf{P}^{\mathrm{B}}\mathbf{K})^{-1}\mathbf{F}^{\mathrm{B}} \end{aligned} \qquad (14)$$

where $\mathbf{P}^{\mathrm{R}}$, $\mathbf{P}^{\mathrm{G}}$, and $\mathbf{P}^{\mathrm{B}}$ are the albedo matrices of the surface for three wavelengths determined by the three filters. Note that the interreflection kernel $\mathbf{K}$ is a purely geometrical quantity and hence remains the same for all wavelengths. $\mathbf{F}^{\mathrm{R}}$, $\mathbf{F}^{\mathrm{G}}$, and $\mathbf{F}^{\mathrm{B}}$ represent the actual shape and color of the surface.

The three pseudo shapes and the pseudo color of the surface may be estimated using photometric stereo. Let the three sources used by photometric stereo have directions s_1, s_2, and s_3. We assume that all three sources emit white light and hence have the same radiant intensity for the three wavelengths. Then, the pseudo facet matrices for the different color bands are determined as:

$$\begin{aligned} \mathbf{F}_p{}^{\mathrm{R}} &= [\mathbf{L}_1{}^{\mathrm{R}}, \mathbf{L}_2{}^{\mathrm{R}}, \mathbf{L}_3{}^{\mathrm{R}}].[s_1, s_2, s_3]^{-1} \\ \mathbf{F}_p{}^{\mathrm{G}} &= [\mathbf{L}_1{}^{\mathrm{G}}, \mathbf{L}_2{}^{\mathrm{G}}, \mathbf{L}_3{}^{\mathrm{G}}].[s_1, s_2, s_3]^{-1} \\ \mathbf{F}_p{}^{\mathrm{B}} &= [\mathbf{L}_1{}^{\mathrm{B}}, \mathbf{L}_2{}^{\mathrm{B}}, \mathbf{L}_3{}^{\mathrm{B}}].[s_1, s_2, s_3]^{-1} \end{aligned} \qquad (15)$$

The above equations indicate that the same surface produces three different pseudo shapes and albedo functions. By noting that the three bands of a color image correspond to different wavelengths of light, we have been able to decompose the interreflection process into three independent processes. Having done this we are in a position to analyze the interreflections in each band independent of the other bands.

5 Recovering Actual Shape and Color

In [Nayar 91] an algorithm is developed that uses equation. 10 to iteratively recover the actual shape and reflectance of a gray surface from its pseudo shape and reflectance. Here, we briefly describe the algorithm for gray surfaces and then extend it to colored surfaces. At first, photometric stereo is applied to the scene. If the scene consists of a single convex surface, the extracted pseudo shape and reflectance are simply the actual ones. However, if the scene consists of concavities, the pseudo shape and reflectance differ from the actual ones. As we showed in Section 3, the pseudo shape is a shallower (less concave) version of the actual shape. Hence, the algorithm uses the pseudo shape and reflectance as conservative initial estimates of the actual shape and reflectance, to compute initial estimates for the albedo matrix $\mathbf{P}$ and the kernel matrix $\mathbf{K}$. The computed $\mathbf{P}$, $\mathbf{K}$, and the pseudo facets $\mathbf{F}_p$ are then inserted in equation 10 to obtain the next estimate of the actual facets. This estimate of the surface is expected to be more concave than the previous one and is used in the next iteration to obtain an even better estimate. The algorithm may hence be written as:

$$\begin{aligned} \mathbf{F}^{k+1} &= (\mathbf{I} - \mathbf{P}^k\mathbf{K}^k)\mathbf{F}_p \\ \text{where } \mathbf{F}^0 &= \mathbf{F}_p \end{aligned} \qquad (16)$$

In the above equation, $\mathbf{P}^k = \mathbf{P}(\mathbf{F}^k)$ and $\mathbf{K}^k = \mathbf{K}(\mathbf{F}^k)$. Note that each estimate of $\mathbf{F}$ provides estimates of *both* shape and reflectance. With each iteration, more accurate estimates of shape and reflectance are obtained and the result finally converges to the actual shape and reflectance.

In implementing the algorithm, the surface is assumed to be continuous. The interreflection kernel depends not only on the orientations of individual facets but also on their relative positions. Therefore, a depth map of the scene must be reconstructed (by integration) from the orientation map computed in each iteration of the algorithm. The continuity assumption is necessary to ensure integrability of the orientation maps.

Several experiments were conducted on real surfaces and the results indicate that the algorithm is robust and accurate [Nayar 90]. The convergence properties of the algorithm are discussed in detail in [Nayar 90].

We now extend the algorithm to recover true shapes and colors of colored surfaces. First, equation 15 is used to estimate the three pseudo shapes and the pseudo color of the surface using photometric stereo. Then, the actual shape and color of the surface are iteratively recovered from the pseudo estimates. Here, the recovery algorithm given by equation 16 is applied to each of the three pseudo facet matrices $\mathbf{F}_p{}^{\mathrm{R}}$, $\mathbf{F}_p{}^{\mathrm{G}}$, and $\mathbf{F}_p{}^{\mathrm{B}}$. The recovery algorithm for colored surfaces may be written as:

$$\mathbf{F}^{\mathrm{R}\,k+1} = (\mathbf{I} - \mathbf{P}^{\mathrm{R}\,k}\mathbf{K}^k)\,\mathbf{F}_p{}^{\mathrm{R}} \quad (17)$$

$$\mathbf{F}^{\mathrm{G}\,k+1} = (\mathbf{I} - \mathbf{P}^{\mathrm{G}\,k}\mathbf{K}^k)\,\mathbf{F}_p{}^{\mathrm{G}} \quad (18)$$

$$\mathbf{F}^{\mathrm{B}\,k+1} = (\mathbf{I} - \mathbf{P}^{\mathrm{B}\,k}\mathbf{K}^k)\,\mathbf{F}_p{}^{\mathrm{B}} \quad (19)$$

where:

$$\mathbf{F}^{\mathrm{R}\,0} = \mathbf{F}_p{}^{\mathrm{R}} \quad (20)$$

$$\mathbf{F}^{\mathrm{G}\,0} = \mathbf{F}_p{}^{\mathrm{G}}$$

$$\mathbf{F}^{\mathrm{B}\,0} = \mathbf{F}_p{}^{\mathrm{B}}$$

Note that the recovery algorithm is applied independently to the three pseudo facet matrices $\mathbf{F}_p{}^{\mathrm{R}}$, $\mathbf{F}_p{}^{\mathrm{G}}$, and $\mathbf{F}_p{}^{\mathrm{B}}$ to recover the facet matrices $\mathbf{F}^{\mathrm{R}}$, $\mathbf{F}^{\mathrm{G}}$, and $\mathbf{F}^{\mathrm{B}}$, respectively. Therefore, we obtain three estimates of the actual shape of the surface, one obtained from each one of the three pseudo facet matrices. The three shape estimates are identical and correspond to the true shape of the colored surface. Therefore, the shape of the surface may be recovered from any one of the three pseudo facet matrices. However, in order to recover the actual color at each point of the surface, we need to recover all three facet matrices.

The above algorithm is applicable to multi-colored surfaces, i.e. surfaces that are comprised of regions of different colors. Once again, the surface is assumed to be continuous to ensure integrability of orientation maps produced in each iteration. We also assume that each surface point has a non-zero albedo value for all three wavelengths of the color image. This assumption is not a severe one as most surfaces encountered in practice have non-zero albedo values for all wavelengths in the visible spectrum.

6 Results

The recovery algorithm has been tested on a variety of gray surfaces and the results indicate that the algorithm is robust and accurate in recovering the actual shape and reflectance of gray surfaces [Nayar 90]. Here, we present the simulation results of applying the algorithm to colored surfaces. The surfaces are three-dimensional with translational symmetry in one direction. The interreflection kernel for the transalational symmetry case was derived by Forsyth and Zissermann [Forsyth 89] and is given in Appendix A.1.

The example shown in Figure 5 is a bucket shaped surface that has three faces with different colors, namely, red, white, and blue. The red face of the surface reflects more red light than any other wavelength in the visible spectrum. Note that the red region does reflect green and blue light but in less amounts relative to red light. We assume that the color images are obtained using three narrow-band filters that pass red light, green light, and blue light. Each color image therefore has three bands, namely, the red band, green band, and the blue band. The actual color of the surface is shown by plotting the three albedo functions (along the surface crossection) for the three wavelengths (red, green, and blue). Within each of the three regions (A, B, and C) on the surface, the albedo functions are constant since all surface points within a region have the same color.

Photometric stereo is used to compute a pseudo shape for each band of the color image. The three pseudo shapes are shown in Figure 5. Note that the pseudo shapes produced by the three bands are different. The pseudo shape and albedo produced by the green band is symmetric since the actual shape and the actual albdeo function for green light are symmetric. On the other hand, the pseudo shapes and albedo functions for the red and blue bands are asymmetric since the actual albedo function of the surface for red and blue light are asymmetric. The pseudo albedo values are often greater that the actual albedo values, sometimes exceeding unity. This results from the fact that interreflections always increase the radiance of a surface point.

At the bottom of Figure 5 we show the results of applying the recovery algorithm given by equation 17. For all three bands the algorithm successfully recovers the actual shape and albedo of the surface. Note that the algorithm must be applied to all three bands to recover the true color of the surface. To recover the shape, however, the algorithm need be applied only to any one of the three band.

Figures 6 and 7 show results obtained for other multi-colored surfaces. In all the results included here and numerous other unreported results, the algorithm convergence (with near zero error) in about 8 iterations of the algorithm.

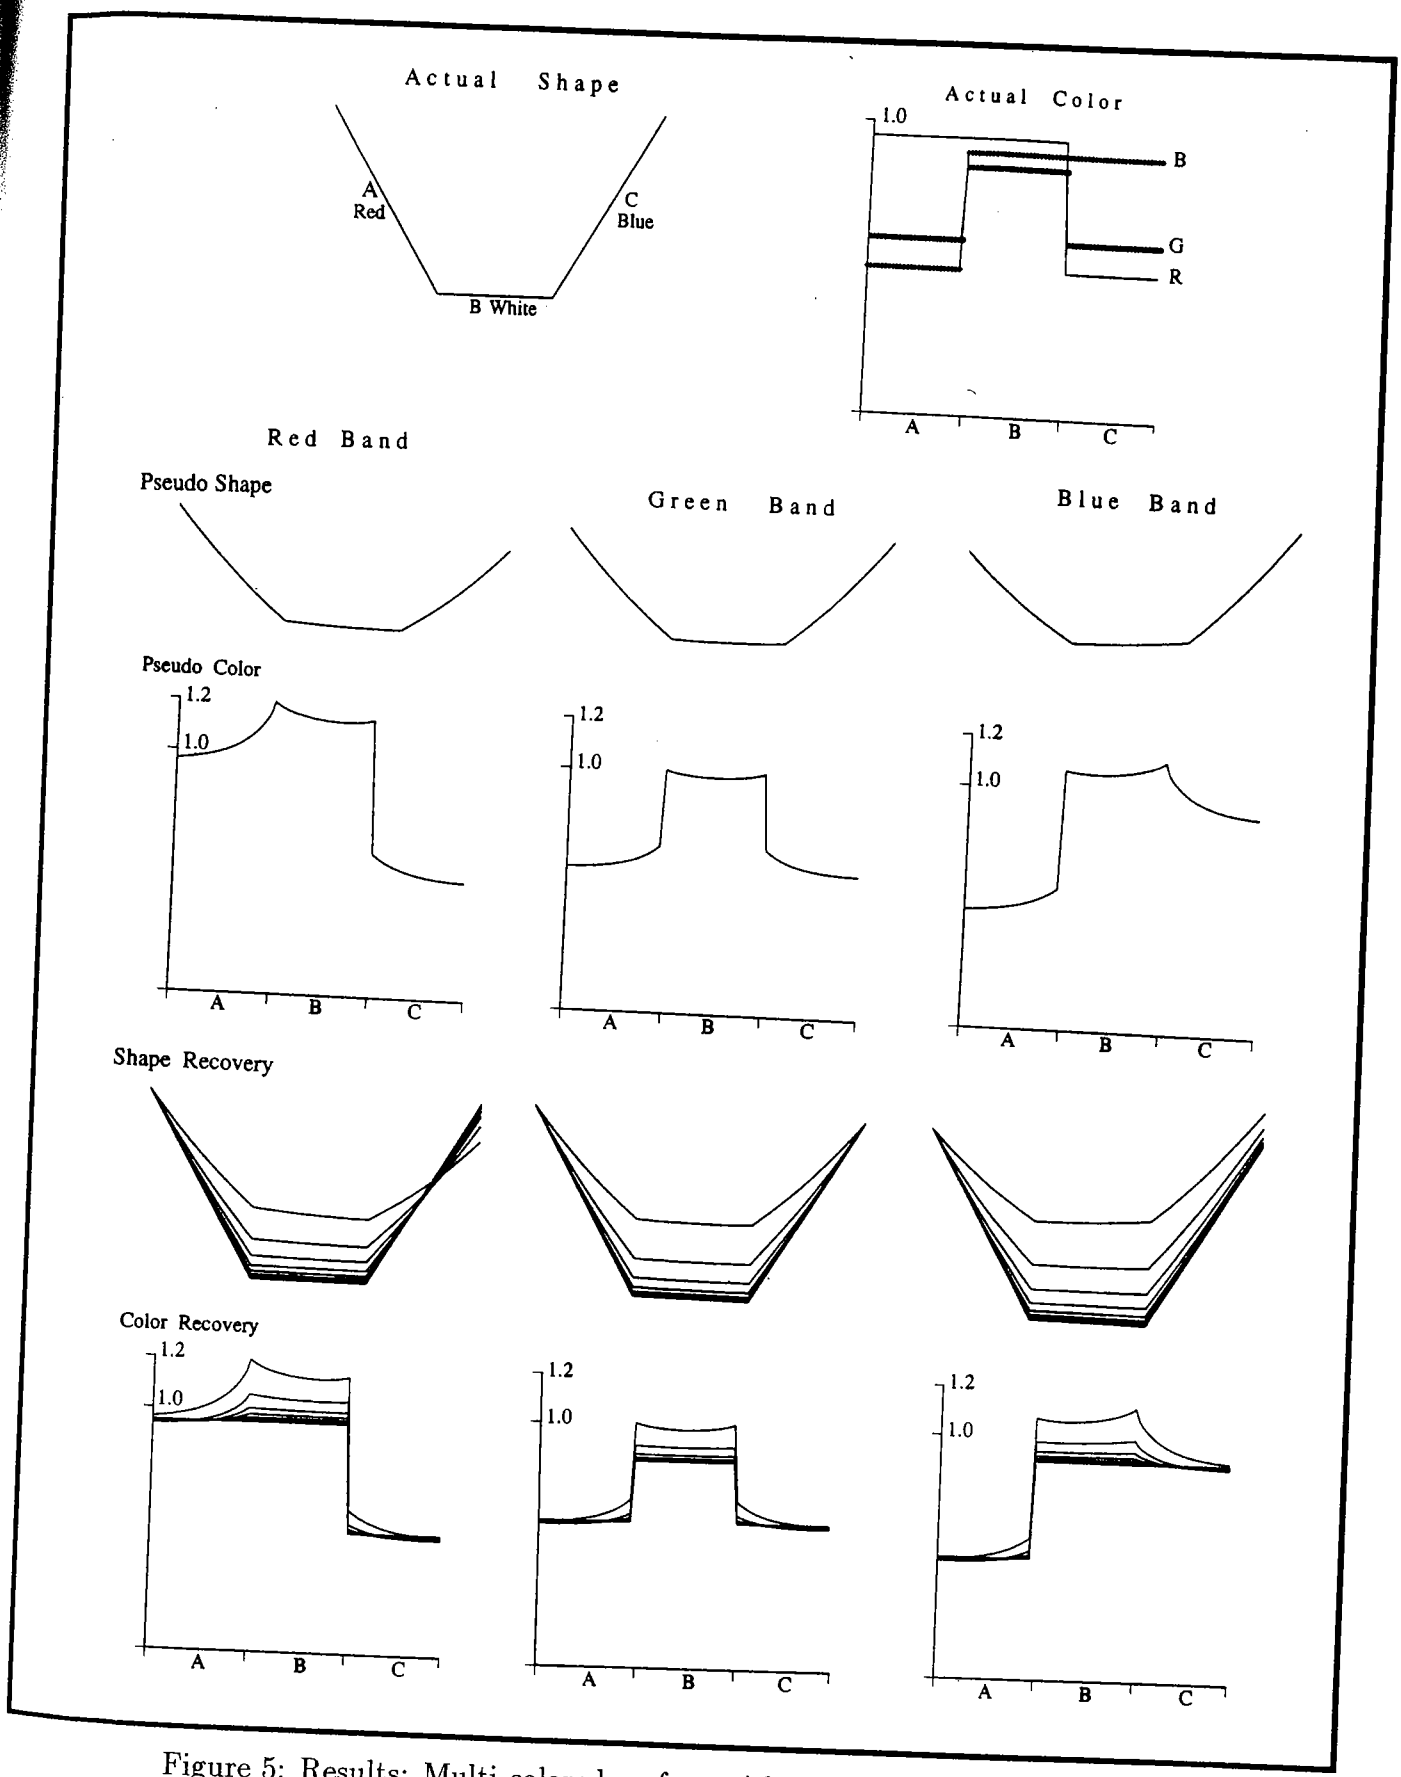

Figure 5: Results: Multi-colored surface with a bucket shaped cross-section.

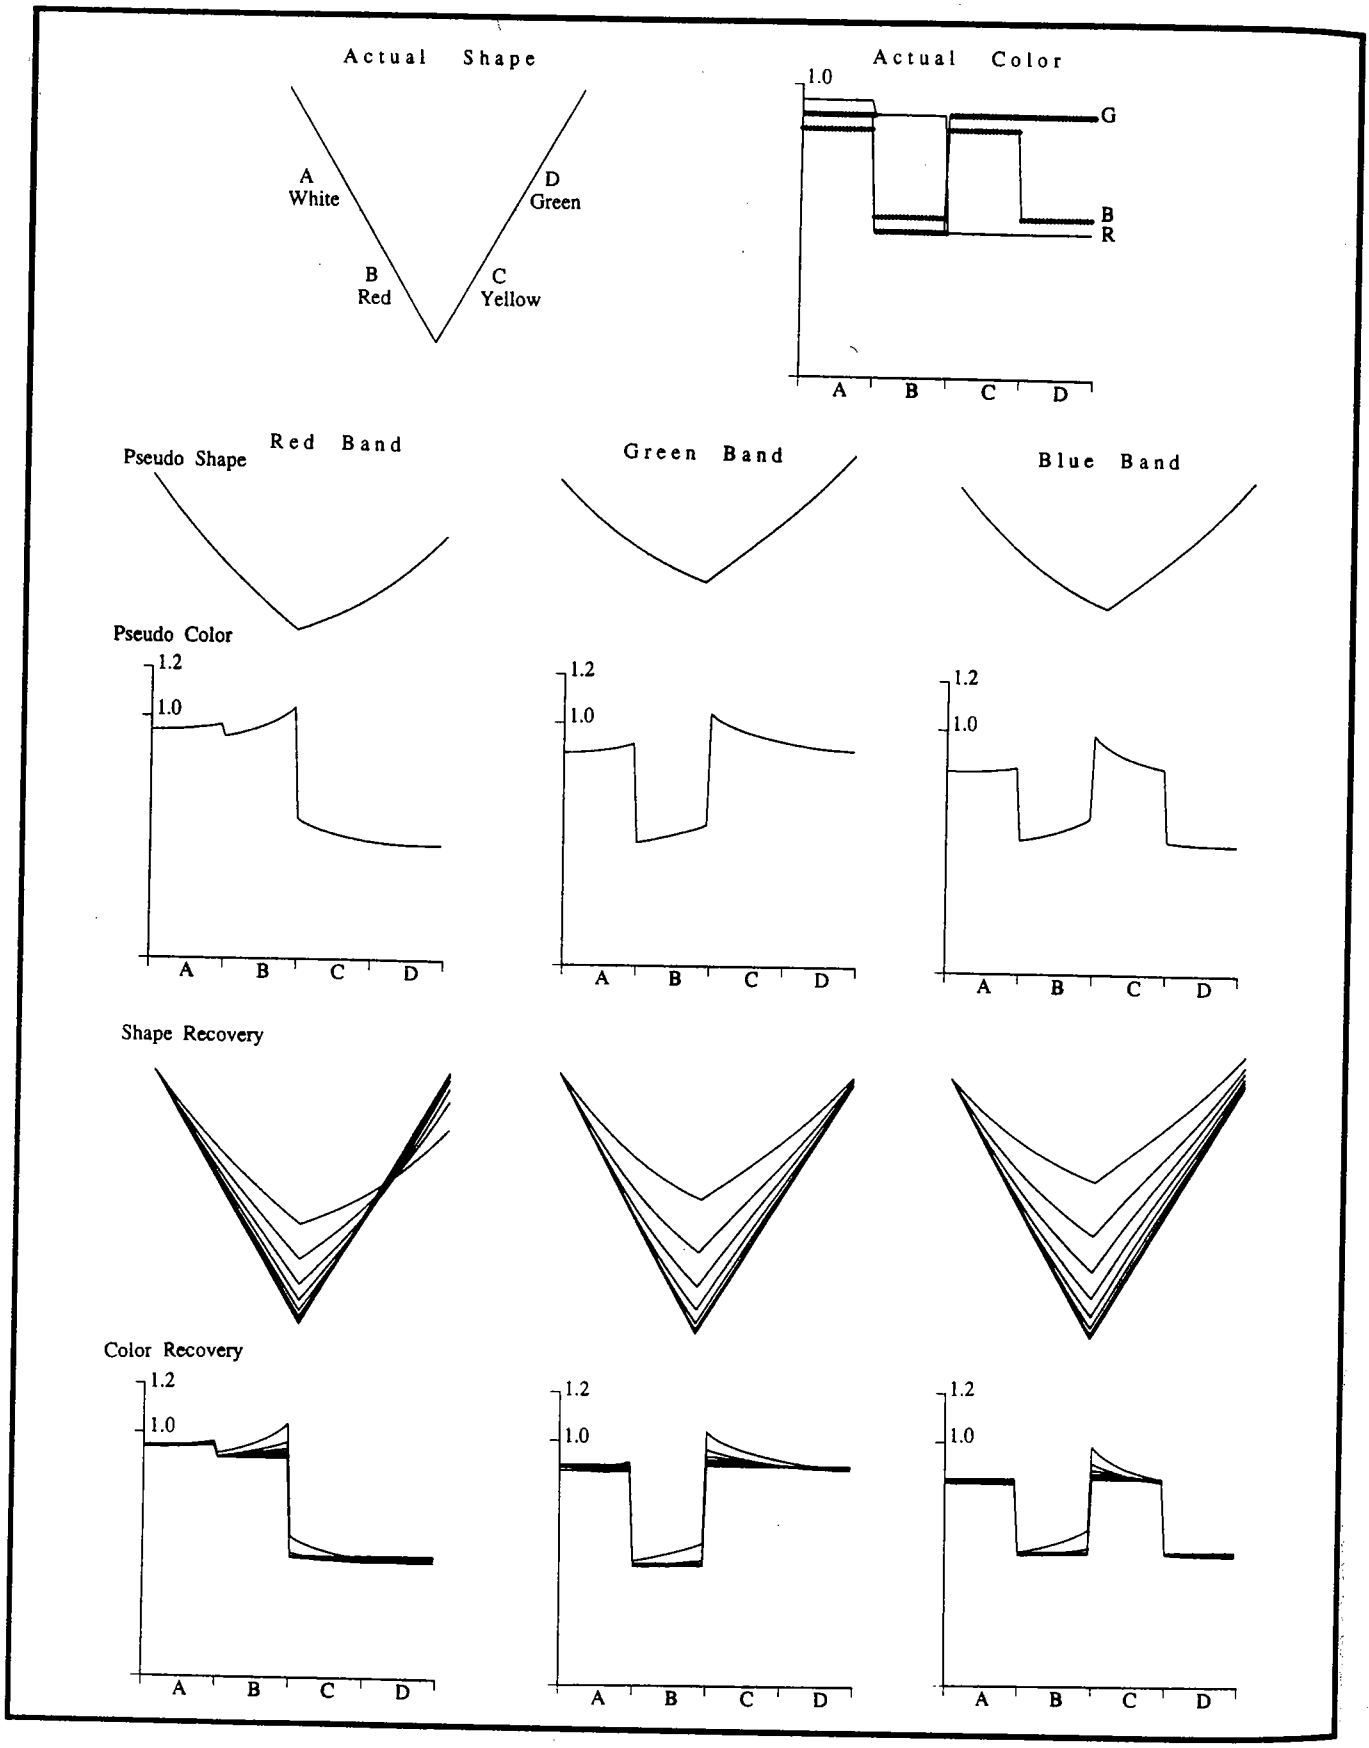

Figure 6: Results: Multi-colored surface with a V-groove cross-section.

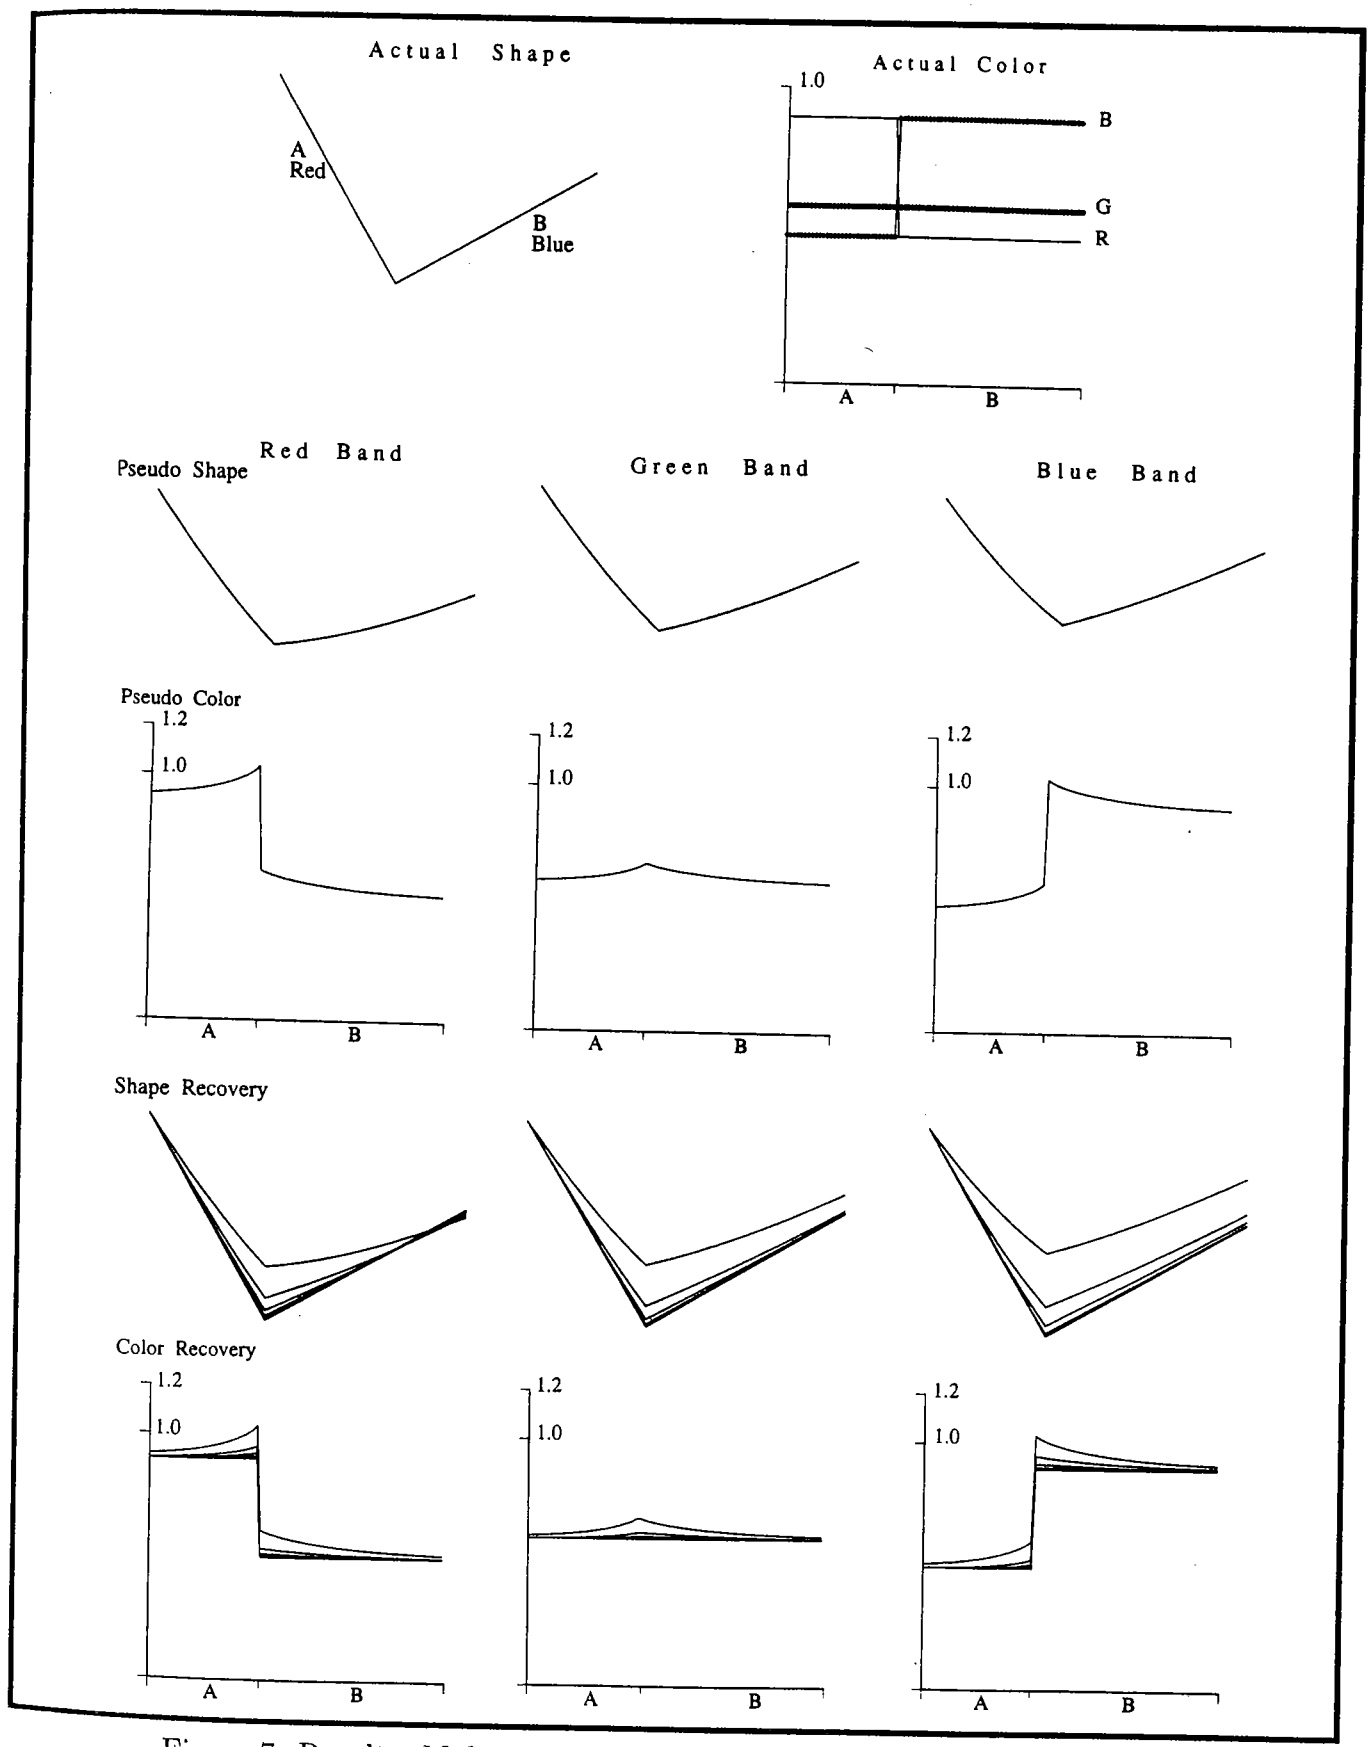

Figure 7: Results: Multi-colored surface with a titled V-groove cross-section.

7 Conclusion

We have presented an algorithm for recovering the shape and reflectance of colored surfaces in the presence of diffuse interreflections. The surfaces may be of arbitrary but continuous shape, and with possibly varying and unknown color.

We showed that for a Lambertian surface the shape and color extracted by shape-from-intensity methods are incorrect in the presence of interreflections. Interreflections can cause the color of one region on a surface to bleed onto a neighboring region. In general these effects are difficult to analyze in a multi-dimensional color space. Noting that the three different bands of a color image correspond to different wavelengths of light, we decomposed the interreflection process into three independent processes. We then showed that the pseudo shape and color extracted using photometric stereo, though incorrect, can be related to the actual shape and color of the surface. The pseudo estimates are shown to have interesting invariance properties; for any given Lambertian surface there exists a single pseudo shape and pseudo color that are invariant to the source directions used to recover them.

Using the properties of the pseudo surface, we developed an algorithm that recovers the actual shape and color of the surface from the pseudo estimates. The algorithm was applied to a variety of multi-colored test surfaces and was shown to be robust and accurate. Motivated by the results presented in this paper, we are currently conducting experiments on real multi-colored surfaces.

A Appendix

A.1 Kernel for Translational Symmetry Case

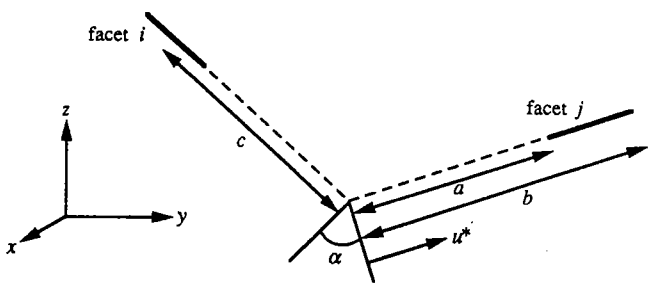

Figure 8: Cross-sectional view of two planar facets that are infinite in the x direction.

Forsyth and Zisserman [Forsyth 89] have derived the interreflection kernel for the special case of two planar facets that have translational symmetry in a single direction. Figure 8 shows a cross-sectional view of two such facets that are infinite in the x direction. The kernel K_{ij} is derived [Forsyth 89] by integrating along the x and y directions, the contribution of all points on facet j to the radiance of a point on the facet i:

$$K_{ij} = -\frac{\pi}{2}\left[\frac{c + u^* cos\alpha}{\left(c^2 + 2\,c\,u^* cos\alpha + u^{*2}\right)^{1/2}}\right]_{u^*=a}^{u^*=b} \quad (21)$$

where α is the angle between the surface normal vectors of the two facets, and the parameter u^* represents the cross-sectional length of the facet j. Since both facets are infinite in length, the same kernel is valid for all points on the facet i. Therefore, under the translational symmetry assumption, the kernel need only be evaluated for points along the cross-section of the surface. Note that the above kernel is valid only for surfaces that are infinite in the direction of symmetry. However, the kernel serves as a good approximation [Forsyth 89] for points that lie around the middle of a surface that is long though finite in the direction of symmetry.

Acknowledgements

The authors are grateful to Ushir Shah for his help in preparing the final manuscript.

References

[Bajcsy 89] R. Bajcsy, S. W. Lee, and A. Leonardis "Image Segmentation with Detection of Highlights and Interreflections Using Color," GRASP LAB 182 MS-CIS-89-39, University of Pennsylvania, Dept. of Computer and Info. Science, June 1989.

[Blinn 77] J. F. Blinn, "Models of Light Reflection for Computer Synthesized Pictures," *ACM Computer Graphics*, Vol. 11, No. 2, pp. 192-198, July, 1977.

[Brill 89] M. H. Brill, "Object-based segmentation and color recognition in multispectral images," *SPIE-SPSE Meeting*, January, 1989, Paper 1076-11.

[Cohen 85] M. F. Cohen and D. P. Greenberg, "The Hemicube: A Radiosity Solution for Complex Environments," *SIGGRAPH 1985*, 19, pp. 31-40, 1985.

[Drew 90] M. S. Drew and B. V. Funt, "Calculating Surface Reflectance using a Single-Bounce Model of Mutual Reflection," *Proc. Third International Conference on Computer Vision*, pp. 394-399, December, 1990.

[Forsyth 89] D. Forsyth and A. Zisserman, "Mutual Illumination," *Proc. of CVPR*, pp. 466-473, 1989.

[Forsyth 90] D. Forsyth and A. Zisserman, "Shape from shading in the light of mutual illumination," *Image and Vision Computing*, Vol. 8, No. 1, pp. 42-49, February 1990.

[Gilchrist 79] A. L. Gilchrist, "The Perception of Surface Blacks and Whites," *Scientific American*, bf 240, pp. 112-124, 1979.

[Horn 70] B. K. P. Horn, "Shape from Shading: A Method for Obtaining the Shape of a Smooth Opaque Object from One View," MIT Project MAC Internal Report TR-79 and MIT AI Laboratory Technical Report 232, November, 1970.

[Horn 77] B. K. P. Horn, "Image intensity understanding," *Artificial Intelligence*, Vol. 8, No. 2, pp. 201-231, 1977.

[Horn 89] B. K. P. Horn and M. J. Brooks, editors, *Shape from Shading,* MIT Press, 1989.

[Jacquez 55] J. A. Jacquez and H. F. Kuppenheim, "Theory of the integrating sphere," *Journal of Optical Society of America*, Vol. 45, pp. 460-470, 1955.

[Koenderink 83] J. J. Koenderink and A. J. van Doorn, "Geometrical modes as a general method to treat diffuse interreflections in radiometry," *Journal of Optical Society of America*, Vol. 73, No. 6, pp. 843-850, June, 1983.

[Nayar 90] S. K. Nayar, *Shape Recovery using Physical Models of Reflection and Interreflection,* Ph.D. Dissertation, Department of Electrical and Computer Engineering, Carnegie Mellon University, December, 1990.

[Nayar 91] S. K. Nayar, K. Ikeuchi, T. Kanade, "Shape from Interreflections," *International Journal of Computer Vision,* Vol. 6, No. 3, pp. 173-195, August, 1991.

[Nicodemus 77] F. E. Nicodemus, J. C. Richmond, J. J. Hsia, I. W. Ginsberg, and T. Limperis, "Geometrical Considerations and Nomenclature for Reflectance," NBS Monograph 160, National Bureau of Standards, October 1977.

[Novak 90] C. Novak, "Ph.D Thesis Proposal," Department of Computer Science, Carnegie Mellon University, August 1990.

[Woodham 78] R. J. Woodham, "Photometric stereo: A reflectance map technique for determining surface orientation from image intensity," *Proc. SPIE*, Vol. 155, pp. 136-143, 1978.