International Journal of Computer Vision, 6:1, 5-24 (1991)

Color Constancy from Mutual Reflection

BRIAN V. FUNT, MARK S. DREW AND JIAN HO
School of Computing Science, Simon Fraser University, Vancouver, British Columbia, Canada V5A 1S6

Abstract
Mutual reflection occurs when light reflected from one surface illuminates a second surface. In this situation, the color of one or both surfaces can be modified by a *color-bleeding* effect. In this article we examine how sensor values (e.g., RGB values) are modified in the mutual reflection region and show that a good approximation of the surface spectral reflectance function for each surface can be recovered by using the extra information from mutual reflection. Thus color constancy results from an examination of mutual reflection. Use is made of finite dimensional linear models for ambient illumination and for surface spectral reflectance. If m and n are the number of basis functions required to model illumination and surface spectral reflectance respectively, then we find that the number of different sensor classes p must satisfy the condition $p \geq (2\,n + m)/3$. If we use three basis functions to model illumination and three basis functions to model surface spectral reflectance, then only three classes of sensors are required to carry out the algorithm. Results are presented showing a small increase in error over the error inherent in the underlying finite dimension models.

1 Introduction

Two illuminated surfaces of different reflectance can appear to have their colors bleed into one another in regions where light reflected from one surface falls onto the other surface. This is a mutual reflection effect. Figure 1 shows two surfaces that form an edge, and the resulting interreflection effect. In this article we examine how the lights reflected separately from each surface combine to form RGB values in the mutual reflection region. Using vector models for the illumination spectral power distribution and for surface reflectance functions we show that by looking at the extra information coming from measurements of the interreflection RGB values—in addition to sensor measurements from each surface separately—it is possible to recover the surface reflectance properties of both surfaces. In this way we use mutual reflection to obtain "color constancy."

The term *color constancy* refers to the ability of humans and some animals to perceive object colors as approximately constant, independent of changing illumination (see, e.g., Beck 1972). There is evidence that several simultaneous mechanisms contribute to color constancy in humans (Blackwell & Buchsbaum 1988). Nevertheless, we have found in (Ho, et al., 1990), that it is the case that an algorithm can be built to recover reflectance from the color signal, in which illumination and surface characteristics are confounded, without regard to colored surrounds, selective adaptation, memory of colors, etc., provided that the complete color signal is known. Most recent attempts to accomplish this end have made use of finite dimensional linear models in which illumination and surface spectral reflectance are approximated by a weighted sum of a few basis functions of wavelength (Brainard et al. 1989; Brill & West 1986; Buchsbaum 1980; D'Zmura & Lennie 1986; Gershon et al. 1987; Maloney 1985; Maloney & Wandell 1986; Wandell 1987; Yuille 1987).

In the innovative work of Maloney and Wandell 1985; 1986; 1987, several special conditions have to be true for reflectance to be recoverable from the color signal entering the camera. Some of these conditions are that the illumination must be constant over a given segment of the image, and that a sufficiency of different color information must be available over the image. More importantly, in this method the number of sensor classes (e.g., 3 for RGB sensors) must be at least as great as the number of basis functions modeling the surface spectral reflectance, plus 1. This means that for a usable

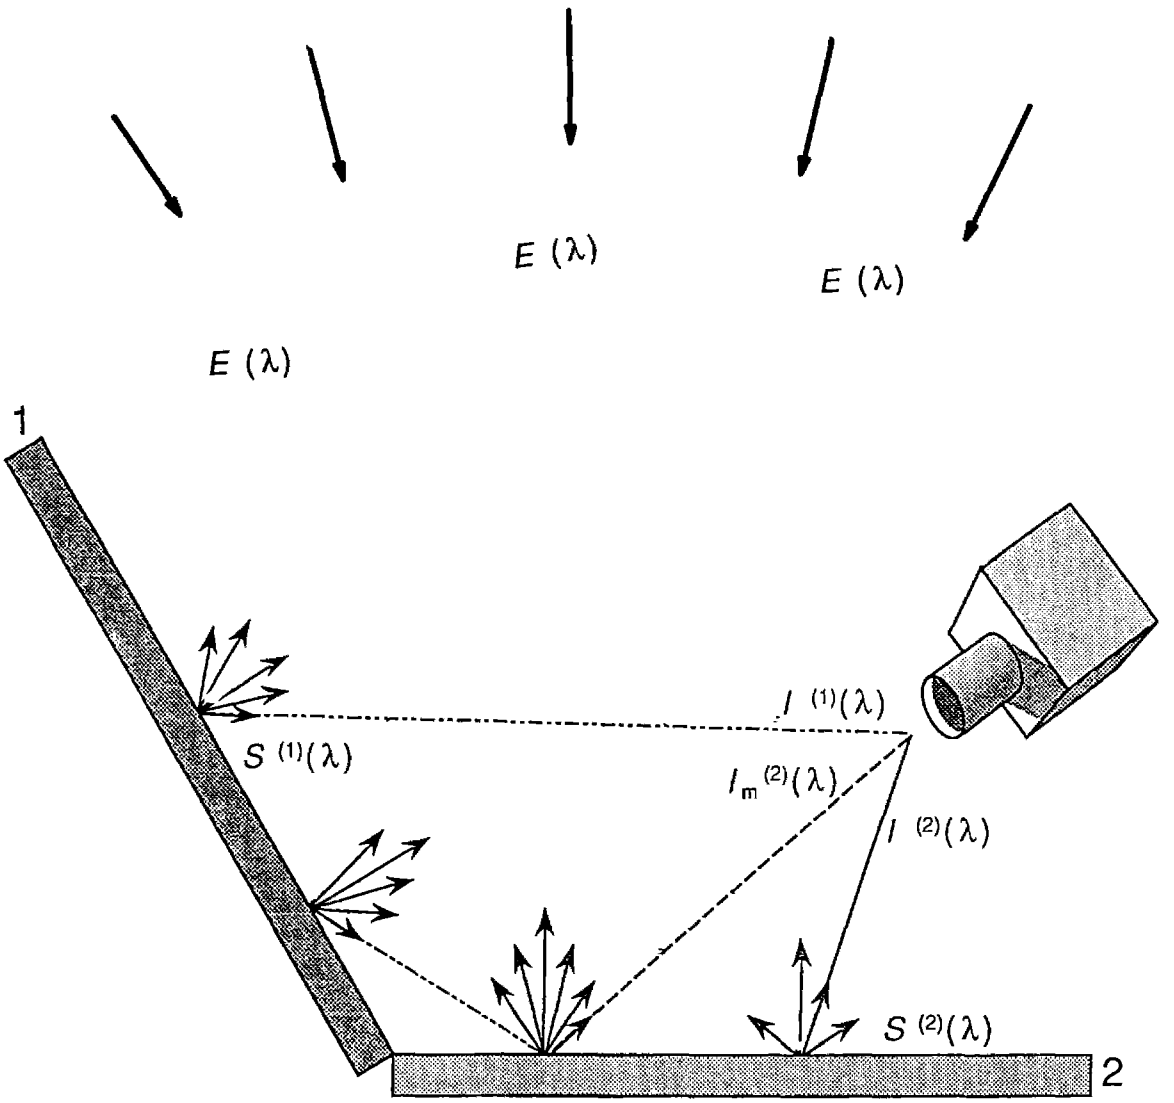

Fig. 1. Isotropic illumination with spectral power distribution $E(\lambda)$ impinges on surfaces 1 and 2, which have surface reflectance functions $S^{(1)}(\lambda)$ and $S^{(2)}(\lambda)$ respectively. Far from the edge, surface 1 reflects color signal $I^{(1)}(\lambda)$ toward the camera; and, similarly, surface 2 reflects $I^{(2)}(\lambda)$. Close to the edge, $E(\lambda)$ is augmented by light reflected from the other surface. The color signals reflected from the surfaces near the edge are $I_m^{(1)}(\lambda)$ and $I_m^{(2)}(\lambda)$.

dimensionality for the set of basis functions modeling reflectance, say 3 or better (Maloney 1986), one must somehow develop a "fourth sensor class" to provide enough information to allow a solution that disentangles illumination and surface spectral reflectance. Alternatively, one can continue using the usual 3 sensor classes, but at the expense of having only a dimensionality of 2 with which to model surface reflectance.

Maloney and Wandell's method essentially involves writing down a set of equations relating the sensor values ρ_k, $k = 1...p$, to the finite dimensional linear model weights, ϵ_i, $i = 1...m$ for illumination, and σ_j, $j = 1...n$ for surfaces, and then coupling the sets of equations at different pixels by using the assumption of constant illumination. Taking into account enough pixels, which must have sufficiently different sensor values, brings the number of equations equal to the overall number of unknowns, allowing a solution for ϵ_i and σ_j.

In Ho et al. (1990) it was shown that the need for additional sensor values can essentially be fulfilled by requiring that a finer sampling of the color signal

spectrum be introduced as additional input. In that work, it was shown that from detailed knowledge of the spectral power distribution of the color signal, a complete disambiguation of the reflectance and illumination from the color signal is possible, and in fact is possible independently at each pixel and without the need for constant illumination or any special requirement on the richness of color variation over image segments. This method (Ho et al. 1990) relies on a statistical determination of the finite dimensional model weights using a least-squares method, and has the sole requirement that a special mathematical relationship holds among the basis vectors, viz. that the set of functions generated by forming products of the basis functions for illumination and surface reflectance forms a linearly independent set. The main benefit of this method is that one can increase the dimensionality of the surface reflectance space without having to increase the number of ρ_k values beyond the 3 that correspond to the cones in the human visual system.

In Funt & Ho (1989) and Gershon et al. (1986, 1987) it was shown that the required spectral information for the color signal is derivable from the chromatic aberration effect at edges between adjacent color regions. In fact, what is found via chromatic aberration is the *difference* between color signals from each side of the edge. Coupled with the method of Ho et al. (1990) and given this spectral information plus RGB sensor values from each side of the edge, the values of ϵ_i and σ_j can be determined and color constancy achieved.

In this article, we take a different approach, although the basic motivation is the same—we seek more information than just the ρ_k values in order to circumvent the "fourth sensor problem." Here, we consider the mutual reflection at a 3-dimensional edge, and the resulting bleeding of colors at the 2-dimensional image edge, without regard to chromatic aberration. Such interreflections generally occur near the edges of objects and the boundaries where objects occlude one another (Horn 1986).

We can take measurements of ρ_k values from pixels where mutual reflection effects are present and compare them to values where such effects must be absent or small. By using such measurements on each side of an image edge, relatively far from the edge and also relatively close to the edge, we capture sufficient information to determine ϵ_i and σ_j separately on each side of the edge, and hence determine the true reflectance for each color patch. This is more than enough information to provide color constancy. In Gershon et al. (1986) it was shown that mutual reflection can help to distinguish shadow boundaries from material changes. Here we show that mutual reflection has another positive aspect—it is an effect that can be exploited to achieve color constancy. As well, once constant color descriptors are found, the mutual reflection component can be effectively removed from the image. Such an image filter could be useful for shape-from-shading, etc. (Forsyth & Zisserman 1989, 1990).

The only constraint on the formulation turns out to be that the number p of sensor classes must be greater than or equal to $(2n + m)/3$, where m and n are the number of basis functions required to model illumination and surface spectral reflectance respectively. This constraint is satisfied by values $p = m = n = 3$ and hence we can keep to only 3 sensor classes while still allowing $n = 3$ for describing reflectances. These values of p, m, and n are shown to recover both surface spectral reflectance and illumination spectral curves well within acceptable errors.

Here, we address only a small part of a large problem. We assume a preliminary edge-finding and segmentation algorithm that would identify those edges where mutual reflection might be present. Given such identification of appropriate sites our analysis concentrates on extracting any available reflectance information from mutual reflection effects.

In section 2 we show how ρ_k measurements on each side of a color edge, and inside the mutual reflection region, are related to the values of ϵ_i and σ_j for each side of the edge separately, assuming a finite dimensional linear model for describing natural lighting and reflectances. The resulting model is quite complex, but some reasonable assumptions reduce the equations to solvable form. We show that a reciprocity relation exists between the mutual reflection measurements on each side of an edge, so that only ρ_k values in the mutual reflection zone on a single side of the edge are required.

In section 3 an algorithm is presented for solving for the reflectance weights σ_j on each side of the edge separately, as well as the illumination weights ϵ_i, which are assumed to not change across an edge. We develop a uniqueness proof that guarantees convergence of the algorithm.

In section 4 we set out the geometrical significance of the mutual reflection configuration factor used in the analysis.

In section 5 we examine the accuracy of the solution method by carrying out simulations. The results are

seen to be very good in most cases, with accuracies of the resulting reflectance spectra not degraded much from the accuracy of the underlying finite dimensional model.

2 Finite-dimensional Models Applied to Mutual Reflection

2.1 Mutual Reflection Equations

We assume that illumination and surface spectral reflectance are each modeled to an acceptable degree of accuracy by finite dimensional linear models. Depending on the dimensionality of the set of basis functions, not all illumination spectra and surface reflectances will be well modeled; generally, the higher the dimensionality of the set of basis functions the better will be the approximation (see, e.g., Maloney 1986). However, we shall show below using simulations that whether or not the accuracy of the underlying finite dimensional model is good, the accuracy of the solution for color constancy from mutual reflection measurements is not greatly reduced.

For reflectance basis functions we use those determined by Cohen via a statistical principal component analysis of 150 Munsell chips randomly selected from a total of 433 chips (Cohen 1964); see also Parkkinen et al. (1989). For illumination, a similar analysis was carried out by Judd et al. (1964) for 622 typical daylight samples. We demonstrate below that it is possible to use $n = m = 3$ basis functions for surfaces and illumination and still have only $p = 3$ sensor classes.

Most naturally occurring daylights and many surface reflectances can be approximated reasonably well using these basis functions, so that by determining a good approximation of the weights ϵ_i and σ_j, we are in fact deriving an approximation of the entire illumination or reflectance spectrum. Hence what we require from a solution is actually more stringent than simply *color* constancy, in that we seek to recover more information than just chromaticity values. And, in fact, the measure of accuracy we shall use is the error sum of squares over the entire visible spectrum of an approximate reflectance spectrum compared with an actual one.

Consider the situation depicted in figure 1. *Isotropic* illumination $E(\lambda)$, assumed constant across an edge, impinges on two surfaces, 1 and 2. Assume surface 1 has spectral reflectance function $S^{(1)}(\lambda)$ and surface 2 has reflectance $S^{(2)}(\lambda)$, all surfaces being ideal diffuse (i.e., Lambertian) reflectors. Then depending on the geometry, interreflection may be present as shown in figure 1. In fact, the two surfaces need not meet at an edge for a mutual reflection effect to come into play; and we do not require the surfaces to be flat. We show ray geometry in figure 1 to clarify the interreflection process. However, ray geometry is actually unnecessary since perfect diffuse reflectors obliterate directional information: reflections are really in all directions into the hemisphere above each surface. We assume here that illumination is uniform on each surface and hence exclude shadow situations.

Expanding the illumination function in terms of m basis functions, the spectral power distribution can be approximated by the sum

$$E(\lambda) = \sum_{i=1}^{m} \epsilon_i E_i(\lambda)$$

with weights ϵ_i. Similarly, for surface 1 the reflectance can be written

$$S^{(1)}(\lambda) = \sum_{j=1}^{n} \sigma_j^{(1)} S_j(\lambda)$$

where reflectance is modeled by n basis functions with n weights $\sigma_j^{(1)}$. Surface 2 has reflectance

$$S^{(2)}(\lambda) = \sum_{j=1}^{n} \sigma_j^{(2)} S_j(\lambda)$$

The color signal $I(\lambda)$ is the light reflected by the surface, and is given by the product of E and S. On the surface 1 side, at a point relatively far from the edge, the color signal is simply

$$I^{(1)}(\lambda) = E(\lambda)\, S^{(1)}(\lambda)$$

Similarly, on the surface 2 side we have

$$I^{(2)}(\lambda) = E(\lambda)\, S^{(2)}(\lambda)$$

Substituting the basis function decomposition, we have

$$I^{(1)}(\lambda) = \sum_{i=1}^{m}\sum_{j=1}^{n} \epsilon_i \sigma_j^{(1)} E_i(\lambda) S_j(\lambda)$$

$$I^{(2)}(\lambda) = \sum_{i=1}^{m}\sum_{j=1}^{n} \epsilon_i \sigma_j^{(2)} E_i(\lambda) S_j(\lambda)$$

Now consider the mutual reflection region in which part of $I^{(1)}(\lambda)$, the light reflected from surface 1, impinges onto surface 2. Similarly, part of $I^{(2)}(\lambda)$ impinges onto surface 1. The amount of $I^{(1)}(\lambda)$ that is intercepted by surface 2 is dependent on both the configuration of the surfaces and on the shape of each surface, since not every reflected ray is intercepted (consider a small triangle of surface 1 forming an edge with a large plane of surface 2). In computer graphics diffuse interreflections are modeled using *configuration factors* that give the fraction of light from one surface that reaches another surface (Goral et al. 1984). We discuss these in section 4.

Since there is no a priori reason to assume that one surface is more reflective than the other (although for clarity in figure 1 only a single interreflection is shown), we consider the fraction α_{12} of $I(\lambda)$ from 1 that strikes 2, and the similarly defined α_{21}. Denote by $I_m^{(2)}$ the color signal from a spot on surface 2 in the mutual reflection region, so that it includes a contribution from light reflected by surface 1. It consists of two parts. The first part is $I^{(2)}(\lambda)$, the reflection of $E(\lambda)$ from surface 2. This we assume is the same as from a spot relatively far from the edge, outside the mutual reflection region. The second part is due to the light reflected from surface 1. Since the latter is presumed to come from a spot on surface 1 that is near the edge, the second part is a contribution from $I_m^{(1)}$ (cf., Goral 1984). Hence the color signals coming from each side of the edge in the mutual reflection region are

$$\begin{aligned} I_m^{(1)}(\lambda) &= I^{(1)}(\lambda) + \alpha_{21} I_m^{(2)}(\lambda) S^{(1)}(\lambda) \\ I_m^{(2)}(\lambda) &= I^{(2)}(\lambda) + \alpha_{12} I_m^{(1)}(\lambda) S^{(2)}(\lambda) \end{aligned} \tag{1}$$

where α_{12} is not assumed to be the same as α_{21}, and we do not include dependency of the surface spectral reflectance function shape on angle of incidence or angle of line of sight (although it could change by an overall multiplicative constant).

Solving the above set of simultaneous equations, we can write each $I_m(\lambda)$ in terms of $I^{(1)}(\lambda)$ and $I^{(2)}(\lambda)$ from each surface separately and the geometrical factors α:

$$\begin{aligned} I_m^{(1)}(\lambda) &= \frac{I^{(1)}(\lambda) + \alpha_{21} I^{(2)}(\lambda) S^{(1)}(\lambda)}{1 - \alpha_{12}\alpha_{21} S^{(1)}(\lambda) S^{(2)}(\lambda)} \\ I_m^{(2)}(\lambda) &= \frac{I^{(2)}(\lambda) + \alpha_{12} I^{(1)}(\lambda) S^{(2)}(\lambda)}{1 - \alpha_{12}\alpha_{21} S^{(1)}(\lambda) S^{(2)}(\lambda)} \end{aligned} \tag{2}$$

2.2 Simplified Equations

The above set of equations would be quite difficult to solve numerically. However, since both α and $S(\lambda)$ are less than 1, in most cases not much accuracy is lost by rewriting the equations to first order in $\alpha \cdot S$. This amounts to ignoring more than one bounce of the color signal between the surfaces. With typical values of S 0.01–0.3 and typical values of α of ≤ 0.5 this approximation might be off by at most a few percent. Of course, for highly reflective materials (0.3 represents the 92 percentile in the Krinov catalog (Krinov 1947) and for geometric configurations with high values of α, the approximation would be farther from the full model.

To first order, we have

$$\begin{aligned} I_m^{(1)}(\lambda) &= I^{(1)}(\lambda) + \alpha_{21} I^{(2)}(\lambda) S^{(1)}(\lambda) \\ I_m^{(2)}(\lambda) &= I^{(2)}(\lambda) + \alpha_{12} I^{(1)}(\lambda) S^{(2)}(\lambda) \end{aligned} \tag{1a}$$

as the set to solve.

To see how these equations can be further reduced, denote the terms involving α by

$$\begin{aligned} \bar{I}_m^{(1)}(\lambda) &\equiv I_m^{(1)}(\lambda) - I^{(1)}(\lambda) \\ \bar{I}_m^{(2)}(\lambda) &\equiv I_m^{(2)}(\lambda) - I^{(2)}(\lambda) \end{aligned}$$

That is,

$$\bar{I}_m^{(1)}(\lambda) = \alpha_{21} E(\lambda) S^{(2)}(\lambda) S^{(1)}(\lambda)$$

and also

$$\bar{I}_m^{(2)}(\lambda) = \alpha_{12} E(\lambda) S^{(1)}(\lambda) S^{(2)}(\lambda)$$

Therefore, we have a *reciprocity relation*

$$\bar{I}_m^{(1)}(\lambda) = \frac{\alpha_{21}}{\alpha_{12}} \bar{I}_m^{(2)}(\lambda)$$

Since all the signals $I(\lambda)$ are in principle measurable, this relation states that one α can be derived from the other. To obtain a solution that takes into account noise and errors in measurement, samples could be measured at N wavelengths so that the above relation would then represent N equations in 1 unknown. A statistical best estimate results from the least-squares solution for the ratio of the α's. Therefore, the number of unknowns to be solved for is reduced by one. Below, we show how it is possible to use only one of the pair of equations for the $I_m(\lambda)$'s to complete a sufficient set to solve for the remaining unknowns.

To simplify further, consider reducing all the equations involved to equations relating *sensor values* by multiplying by sensor sensitivity functions $R_k(\lambda)$ for $k = 1 \ldots p$ sensor classes and integrating over wavelength (e.g., $p = 3$ for RGB signals). Then following Maloney & Wandell (1986) and defining the precalculated tensor

$$g_{ijk} \equiv \int E_i(\lambda) S_j(\lambda) R_k(\lambda)\, d\lambda$$

the equations for $I^{(1)}(\lambda)$ and $I^{(2)}(\lambda)$ become the following:

$$\begin{aligned} \rho_k^{(1)} &\equiv \int I^{(1)}(\lambda) R_k(\lambda)\, d\lambda \\ &= \sum_{i=1}^{m}\sum_{j=1}^{n} \epsilon_i \sigma_j^{(1)} g_{ijk} \end{aligned} \tag{3}$$

and

$$\begin{aligned} \rho_k^{(2)} &\equiv \int I^{(2)}(\lambda) R_k(\lambda)\, d\lambda \\ &= \sum_{i=1}^{m}\sum_{j=1}^{n} \epsilon_i \sigma_j^{(2)} g_{ijk} \end{aligned} \tag{4}$$

with $k = 1 \ldots p$.

For the mutual reflection color signal, we first substitute the basis function expansions yielding

$$\begin{aligned} I_m^{(2)}(\lambda) &= \sum_{i=1}^{m}\sum_{j=1}^{n} \epsilon_i \sigma_j^{(2)} E_i(\lambda) S_j(\lambda) \\ &\quad + \alpha \sum_{i=1}^{m}\sum_{j=1}^{n}\sum_{j'=1}^{n} \epsilon_i \sigma_j^{(1)} \sigma_{j'}^{(2)} E_i(\lambda) S_j(\lambda) S_{j'}(\lambda) \end{aligned}$$

where we have abbreviated α_{12} simply as α. Now multiplying by $R_k(\lambda)$ and integrating, we find the final equation in our set of equations to solve,

$$\begin{aligned} \rho_k^{\mathrm{m}} &\equiv \int I^{(2)}(\lambda) R_k(\lambda)\, d\lambda + \alpha \int I^{(1)}(\lambda) S^{(2)}(\lambda) R_k(\lambda)\, d\lambda \\ &= \rho_k^{(2)} + \alpha \sum_{i=1}^{m}\sum_{j=1}^{n}\sum_{j'=1}^{n} \epsilon_i \sigma_j^{(1)} \sigma_{j'}^{(2)} h_{ijj'k} \end{aligned} \tag{5}$$

for $k = 1 \ldots p$, where

$$h_{ijj'k} \equiv \int E_i(\lambda) S_j(\lambda) S_{j'}(\lambda) R_k(\lambda)\, d\lambda$$

and we have abbreviated $\rho_k^{\mathrm{m}(2)}$ as ρ_k^{m}.

From equation (5) we can see that it is indeed possible to restrict our attention to measurements of $\rho_k^{(1)}$, $\rho_k^{(2)}$, and $\rho_k^{\mathrm{m}} \equiv \rho_k^{\mathrm{m}(2)}$ on *one* side of the edge, since from the similar equation for $\rho_k^{\mathrm{m}(1)}$ we can see that if a solution is found for α_{12} then α_{21} can be determined via a least-square solution involving the $\rho_k^{\mathrm{m}(1)}$, $\rho_k^{\mathrm{m}(2)}$, $\rho_k^{(1)}$, and $\rho_k^{(2)}$. So that if a solution for ϵ_i, $\sigma_j^{(1)}$, $\sigma_j^{(2)}$, and α_{12} can be found via measurements on one side of an edge in the image then a solution for α_{21} is also available, given measurements $\rho_k^{\mathrm{m}(1)}$ in the mutual reflection region on the other side.

Using measurements $\rho_k^{(1)}$, $\rho_k^{(2)}$, and ρ_k^{m}, we are faced with solving the set of equations (3), (4), and (5) for the ϵ_i, for the reflectance weights of both surfaces $\sigma_j^{(1)}$ and $\sigma_j^{(2)}$, and also for the geometrical factor α.

2.3 Assumptions

It is important to underline the assumptions inherent in the pair of equations (1) and equations (3), (4), and (5).

a. We assume isotropic diffuse illumination, i.e., a sky model as opposed to a sun model or a sun-sky combination.
b. The pair of equations (1) encompasses only two surfaces. However, the set of equations we actually solve, (3), (4), and (5), can be generalized to multiple surfaces.
c. We assume that the surfaces are convex—they do not see themselves.
d. We assume that each surface has a surface reflectance function that is independent of position, in contradistinction to the more complex model of Koenderink & van Doorn (1983).
e. In equation (1) we assume that I_{m} is the same everywhere in the mutual reflection zone. This amounts to assuming that the mutual reflection is between two semi-infinite convex surfaces. With two such surfaces, any point on either surface sees the same fraction of the hemisphere above it filled by the opposite surface (see Horn 1977). Each point, therefore, receives an equal amount of illumination due to mutual reflection from the opposite surface. This is a strong assumption; however, we show below that we can later eliminate it.
f. We also assume that there exists at least one point on each surface where the illumination resulting from mutual reflection from the opposite surface is negligible.

These assumptions mean in effect that we have adopted a two-zone, one-bounce model of mutual reflection. Although this reflects how our initial thinking about the problem developed, a two-zone model is rather restrictive. However, our equations—(1a) in particular—can be reinterpreted in a much more general way. They can be viewed as stating that for any point on side 1 where mutual reflection is present, the factor α_{21} represents the sum of all the contributions of light reflected from side 2 impinging upon it. Thus the factor α can be allowed to vary as a function of position; we can abandon assumption (e) and move to a one-bounce, infinite-zone model of mutual reflection.

The suitability of this interpretation is borne out in section 5, where we test the model on several simulated concave edges including ones derived from a full model of diffuse interreflection.

We are forced to keep assumption (f); however, a simple method of finding pixels where mutual reflection is negligible would be to examine the derivatives of ratios of RGB values. From figure 10 in section 5 we see that the derivatives of ratios can be expected to vanish at pixels where no mutual reflection is present.

Our simulation tests show that the model can be expected to work if one deals with surface reflectance functions for natural materials, but we do not expect the method to work when a one-bounce model is inappropriate. This would be the case for highly reflective surfaces or for geometric configurations in which mutual reflection plays a dominant role, such as inside highly concave enclosures.

2.4 Bound on Dimensionalities

In order that a solution be possible, there must be at least as many equations as unknowns. Counting components in the set $\{\epsilon_i, \sigma_j^{(1)}, \sigma_j^{(2)}\}$ we have $m + 2n$ unknowns, and α adds one more. Equations (3), (4), and (5) each provide p equations, for a total of $3p$ equations. Hence a solution is possible only if the condition $3p \geq m + 2n + 1$ holds. Therefore we cannot use $p = m = n = 3$ unless we impose a further constraint. Since we must only expect solutions to yield illumination and surface reflectance spectra up to a multiplicative constant, because the ϵ_i and σ_j occur in products in the color signal, we must make one further requirement on one of the ϵ_i or σ_j; we choose to set $\epsilon_1 \equiv 1$. That is, for surfaces we determine reflectance but not brightness.

Now our condition reads $p \geq (m + 2n)/3$ and we have a nonlinear set of equations for ϵ_i, $\sigma_j^{(1)}$, $\sigma_j^{(2)}$, and α that is at least sufficient for a solution. Since the set of equations is also clearly independent, in that one cannot write any equation in terms of the others, we would be confident of having fulfilled necessary and sufficient conditions for a solution if the set of equations were linear; however, because they are nonlinear we must proceed with caution.

3 Implementation and Results

We reduce the solution of the nonlienar set of equations (3), (4), and (5) to the solution of *linear* ones by breaking up the solution into a multi-stage algorithm. Starting with some reasonable initialization for ϵ_i, we iteratively solve in turn for $\sigma_j^{(1)}$, then $\sigma_j^{(2)}$, and finally ϵ_i. The condition $\epsilon_1 \equiv 1$ at each step of the iteration also forces α to converge.

In detail, with respect to equations (3), (4), and (5), we use the following algorithm:

Initialize $\epsilon_1 = \epsilon_2 = \epsilon_3 = 1$

Step 1: Use (3) to solve for $\sigma_j^{(1)}$ in terms of ϵ_i, since (3) is linear in these unknowns

Step 2: Use (5) to solve for $\alpha \cdot \sigma_j^{(2)}$ in terms of ϵ_i and $\sigma_j^{(1)}$.

Step 3: Use (4) to solve for $\epsilon_j \div \alpha$ in terms of $\alpha \cdot \sigma_j^{(2)}$.

Step 4: Now set $\epsilon_1 \equiv 1$: i.e., set $\epsilon_1 \div \alpha \leftarrow 1/\alpha$. Therefore, set $\alpha \leftarrow 1/\epsilon_1$ and $\epsilon_i \leftarrow \epsilon_i/\epsilon_1$.

Iterate until all values in the set $\{\epsilon_i, \sigma_j^{(1)}, \sigma_j^{(2)}, \alpha\}$ change less than a prescribed tolerance.

Since the equations are nonlinear, the order of the above steps is important. We found by trial and error that the steps in the order described lead in a stable fashion to the desired solution. The algorithm produced stable, unique results regardless of how the initial values for ϵ_i were chosen over a wide range. However, for the first initial value we imposed the constraint $\epsilon_1 \equiv 1$ from the outset. We also found that when the underlying finite-dimensional model did an extremely poor job of describing both surfaces, the algorithm was slow to converge and did not do as well in recovering surface reflectance. We explore this situation in section 5 by looking at very noisy, poorly modeled surfaces.

An important consideration is that the algorithm should obtain a unique solution. Uniqueness and speed of convergence can be established by converting the

problem to be solved into the form of a fixed-point problem (see, Burden et al. 1981). Choosing a particular algorithm corresponds to adopting a particular fixed-point version of a problem. There are often several possible options, and not all of them converge. We can investigate the convergence properties of our algorithm by explicitly converting it to the fixed-point form to which it is equivalent. Then the well-known contraction mapping theorem can be used to test uniqueness as well as give rates of convergence.

Following Steps 1, 2, and 3, we algebraically solve for functions

$$\vec{\sigma}^{(1)} = \vec{\sigma}^{(1)}(\vec{\epsilon}, \vec{\rho}^{(1)}) \tag{6}$$

$$\vec{\sigma}^{(2)} = \vec{\sigma}^{(2)}(\alpha, \vec{\epsilon}, \vec{\rho}^{(1)}, \vec{\mu}) \tag{7}$$

where

$$\vec{\mu} \equiv \vec{\rho}^{\,m(2)} - \vec{\rho}^{(2)}$$

and finally

$$\vec{\epsilon} = \vec{e}(\alpha, \vec{\epsilon}, \vec{\rho}^{(1)}, \vec{\mu}, \vec{\rho}^{(2)}) \tag{8}$$

where we denote by $\vec{e}$ the solution of equation (4) for $\vec{\epsilon}$. Now since we set $\epsilon_1 \equiv 1$, we can solve the ϵ_1 component of equation (8) for α, so that α is given in terms of ϵ_2, ϵ_3, and the camera RGB values $\vec{\rho}$, $\vec{\mu}$ only. Substituting this value of α into the remaining equations for ϵ_2 and ϵ_3 gives a set of two equations in fixed-point form:

$$\begin{aligned} \epsilon_2 &= e_2(\epsilon_2, \epsilon_3, RGB) \\ \epsilon_3 &= e_3(\epsilon_2, \epsilon_3, RGB) \end{aligned} \tag{9}$$

where we have denoted the set of observed camera RGB values simply as *RGB*.

To prove uniqueness we must show (a) that when the values of ϵ_2, ϵ_3 are allowed to range over a reasonable domain D, then the right-hand sides of equations (9) will also be confined to D; and (b) that the absolute values of all possible first partial derivatives of the two right-hand sides with respect to ϵ_2, ϵ_3 are bounded by a constant $K/2$ with $K < 1$. When conditions (a) and (b) are satisfied, the vector function $[e_2(\epsilon_2, \epsilon_3, RBG), e_3(\epsilon_2, \epsilon_3, RGB)]$ is guaranteed to intersect with the vector function $[e_2, e_3] \equiv [\epsilon_2, \epsilon_3]$ at a single point. These conditions also prove convergence of the sequence provided by the algorithm, with the rate of convergence being controlled by the bound of the derivatives.

What this uniqueness check amounts to is carrying out the *first* iteration of our algorithm over a large domain D of *initial values* of ϵ_2, ϵ_3 and examining the resulting solution values e_2, e_3 after a first pass. If the values of e_2, e_3 are within the search space of initial values ϵ_2, ϵ_3 and if the partial derivatives of e_2, e_3 with respect to both ϵ_2 and ϵ_3 are everywhere sufficiently small, then the algorithm is guaranteed to converge to a unique solution.

Since the first-pass estimates e_2, e_3 of ϵ_2, ϵ_3 are generated using the particular observed sensor values $\vec{\rho}^{(1)}$, $\vec{\mu}$, $\vec{\rho}^{(2)}$ corresponding to the pixels under examination, we cannot make a blanket statement of uniqueness for our algorithm. There may be cases in which one cannot prove the theorem guaranteeing unique convergence, even when such cases do in fact converge correctly.

We implemented the uniqueness check as a preprocessing algorithm for the equation-solving algorithm. For a particular input set of sensor values, it is straightforward to search a wide domain of initial values of ϵ_2, ϵ_3 and take partial derivatives of the first-pass estimates e_2, e_3. We found in our simulations (section 5) that uniqueness was guaranteed except in those cases where the finite dimensional model itself represented the surfaces very poorly. Thus the preprocessing step is a useful filter for screening out any pixels for which the algorithm may not converge, such pixels generally corresponding to surface reflectances poorly captured by the finite dimensional model.

In the equation-solving algorithm itself, it is important to note the way in which α appears. Since it always appears in combination with another variable until the final normalizing step, the only place where the above algorithm's accuracy of solution for α and for all the variables is affected by the particular value of α is in forming $I_m^{(2)}(\lambda)$ in equation (2). As a result, the accuracy of solutions depends only very weakly on the particular value of α. We found (see simulations in section 5) that a change of α from 0.01 to 0.5 resulted in only a 6.7% change in the accuracy of the solution for α, with similar results for the other variables.

So long as the mutual reflection effect is indeed present, the solution proceeds to find α. No matter what the actual value of α the percent error in α will be nearly constant for a particular illumination and set of reflectances for the two surfaces. As well, the accuracy of the reflectances recovered depends only weakly on the particular value of α dictated by geometry. Of course, this near-independence with respect to α is true mathematically but noise will break this situation down somewhat. We investigate below the effect of noise on the solution. As well, as $\alpha \rightarrow 0$ we expect the method to fail; this would correspond to the case of negligible interreflection.

In section 5 we apply the algorithm to several cases involving synthesized color signals composed from naturally occurring lights and reflectances. In the next section we investigate the geometrical significance of the factor α.

4 Configuration Factors

One may ask whether in general there will exist points on a surface where we can effectively say that no mutual reflection effect exists. We can answer this question by examining the structure of the parameters α_{12} and α_{21}. The total irradiance onto surface B from surface A as a fraction of the total radiance emanating from surface B into the hemisphere above it is termed the *form-factor* in illumination engineering. Here we are interested instead in how much light reflected from one surface is intercepted at a *particular* spot on the second surface. A catalog of such *configuration factors* (total irradiance as a function of position on surface B incident from all points on surface A) has been assembled by Siegel and Howell (1981). In general, configuration factors can be calculated for any geometry and are independent of wavelength for ideal diffuse reflectors.

For illustrative purposes, consider for simplicity a relatively long (effectively semi-infinite) planar edge (figure 2). Let the fraction of light from the small area ΔA_1 intercepted by the upper surface A_2, divided by the total amount of light from ΔA_1 into the hemisphere above it, be F_{12}. Calculation of the quantity F_{12} will effectively yield the fraction of light F_{21} from A_2 intercepted by ΔA_1 as well, because of the reciprocity relation for configuration factors (Siegel & Howell 1981).

$$A_2F_{21} = \Delta A_1F_{12}$$

If $I(\lambda)$ is measured in W/m^2 of unprojected area then the power in a wavelength interval received over A_2 from light $I^{(1)}(\lambda)$ originating from ΔA_1 is given by

$$P^{(2)}(\lambda) = I^{(1)}(\lambda)\ \Delta A_1F_{12}$$

And the power received at ΔA_1 from $I^{(2)}(\lambda)$ radiated from all of A_2 is

$$P^{(1)}(\lambda) = I^{(2)}(\lambda)A_2F_{21}$$

assuming that $I^{(2)}(\lambda)$ is uniform over A_2.

In general, one finds that the factor F_{12} is given by the Fredholm area integral (Siegel & Howell 1981)

$$F_{12} = \int_{A_2} \frac{\cos\theta_1 \cos\theta_2}{\pi r^2}\, dA_2$$

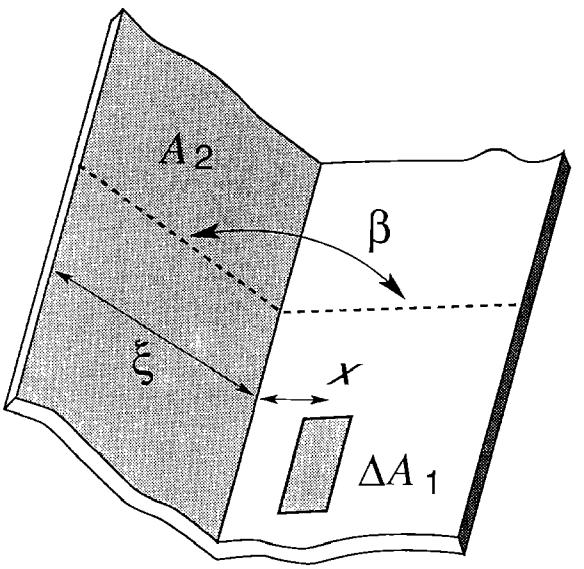

Fig. 2. A long edge formed from flat surfaces. The opening angle is β. ζ is the width of surface 2, with area A_2. The small surface area ΔA_1 is at distance x from the edge.

where r is the length of the vector from ΔA_1 to a point on A_2 and θ_1, θ_2 are the angles made between that vector and the normals to ΔA_1 and A_2.

Given this result, it is in fact straightforward to generalize to the case in which the second surface A_2 is curved, as in figure 3. This figure shows a general curved surface A_2 generated by a line moving parallel to itself and parallel to the plane of ΔA_1; calculating F_{12} for reflection onto A_2, the result is (Siegel & Howell 1981).

$$F_{12} = \frac{1}{2}(\sin\phi_1 - \sin\phi_2)$$

where ϕ_1, ϕ_2 are the angles from ΔA_1 to the top and bottom edges of A_2 (see figure 3). For figure 2, we have $\phi_2 = \pi/2$; from the geometry of figure 2 one finds that

$$\sin\phi_1 = \frac{x - \zeta\cos\beta}{(x^2 + \zeta^2 - 2x\zeta\cos\beta)^{1/2}}$$

Using the reciprocity relations for configuration factors given above and defining the dimensionless distance from the vertex as $X \equiv x/\zeta$, we find (Siegel & Howell 1981) that

$$F_{21} = \frac{\Delta A_1}{2A_2}\left[1 + \frac{\cos\beta - X}{(X^2 + 1 - 2X\cos\beta)^{1/2}}\right]$$

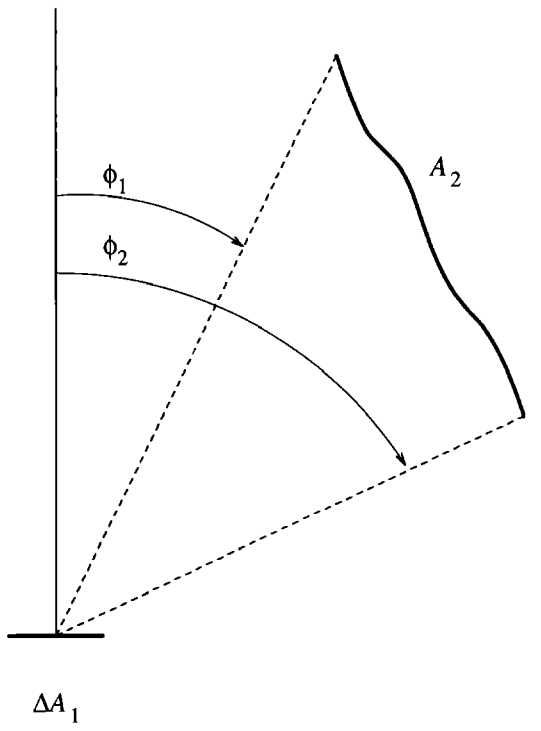

Fig. 3. Angles from ΔA_1 to top and bottom edges of a curved surface A_2.

To relate F_{21} to α_{21}, we must convert power received by ΔA_1 to power received per unit area via division by ΔA_1. Then

$$\alpha_{21}\, I^{(2)}(\lambda) = \frac{I^{(2)}(\lambda)\, A_2 F_{21}}{\Delta A_1}$$

Whereas the algorithm in section 3 extracts α_{21} from pixel RGB information, here we are constructing α_{21} from geometry—(the "forward" problem (Nayar et al. 1990)—and must decide how much of A_2 contributes to mutual reflection onto ΔA_1. For flat surfaces, it is simplest to take all of A_2 into account.

From the identification above, one has

$$\alpha_{21} = \frac{1}{2}\left[\, 1 + \frac{\cos\beta - X}{(X^2 + 1 - 2X\cos\beta)^{1/2}} \,\right]$$

This factor is plotted in figure 4 for representative opening angles β. One can see that as the position X from the vertex increases, the factor α_{21} decreases monotonically, with limit 0. Typical values, for $\beta = 120°$, *are 0.107 for* $X = 0.6$ and 0.028 for $X = 2.0$. Configuration factors between less simple surfaces may be arrived at directly from the definition or by using Stokes theorem to convert the double integral into a line integral around a contour (Sparrow & Cess 1978).

It is useful to note that in the limit $\zeta \to \infty$, $X \to 0$, and $\alpha_{21} \to (1 + \cos\beta)/2$. Therefore, for semi-infinite planes I_m is independent of position, in agreement with Horn (1977). In that case the α_{21} in (10), generated by integrating over an infinity of one-bounce contributions, goes over into our initial one-bounce model (1a). In the next section we test the model derived by reinterpreting the equations to mean that α sums up all the one-bounce contributions from one surface to another, where the surfaces are infinite and need not be planar.

For flat surfaces, figure 4 shows that there are areas where α is relatively small provided one surface is sufficiently wider than the other. For real surfaces, particularly convex ones, there will certainly be many sites classifiable as being far from edges or occlusions.

The form of α_{21} above shows a typical configuration factor that characterizes a mutual reflection effect from a surface A_2 to a position of interest on A_1. This value of α_{21} could be used in equations (2). However, the factor F_{12} above is not related to the number α_{12} discussed in section 2; F_{12} determines how much light from a small spot on A_1 would impinge onto A_2, and is therefore much smaller than F_{21}. In section 2, α_{12} characterizes how much light from A_1 arrives at a small region of A_2, and therefore might be of the same order as α_{21}.

5 Simulations

We test the method by synthesizing color signals from naturally occurring illumination spectra and surface spectral reflectance functions. By passing these synthetic color signals through bandpass filters representing film or other sensor sensitivities, we determine sensor values ρ_k, which then become the input values for our algorithm, along with the proportion α of the light from one surface reaching the second surface. The algorithm is then used to solve for an estimate of the surface spectral reflectance on each side of the edge, the illumination spectrum, and the mutual reflection factor α. In section 5.1 we construct the full mutual reflection input signal equation (2) and test the algorithm by solving the approximate model equations (3), (4), and (5). For simplicity we assume a geometry leading to $\alpha_{12} = \alpha_{21}$ to generate the input values of ρ^m from equation (2). This tests the model itself, but not the assumptions behind it. In section 5.2 we construct a complete infinite-bounce, infinite-zone diffuse mutual reflection edge composed of planar sides and apply our one-bounce model to it. This tests the model's assumptions.

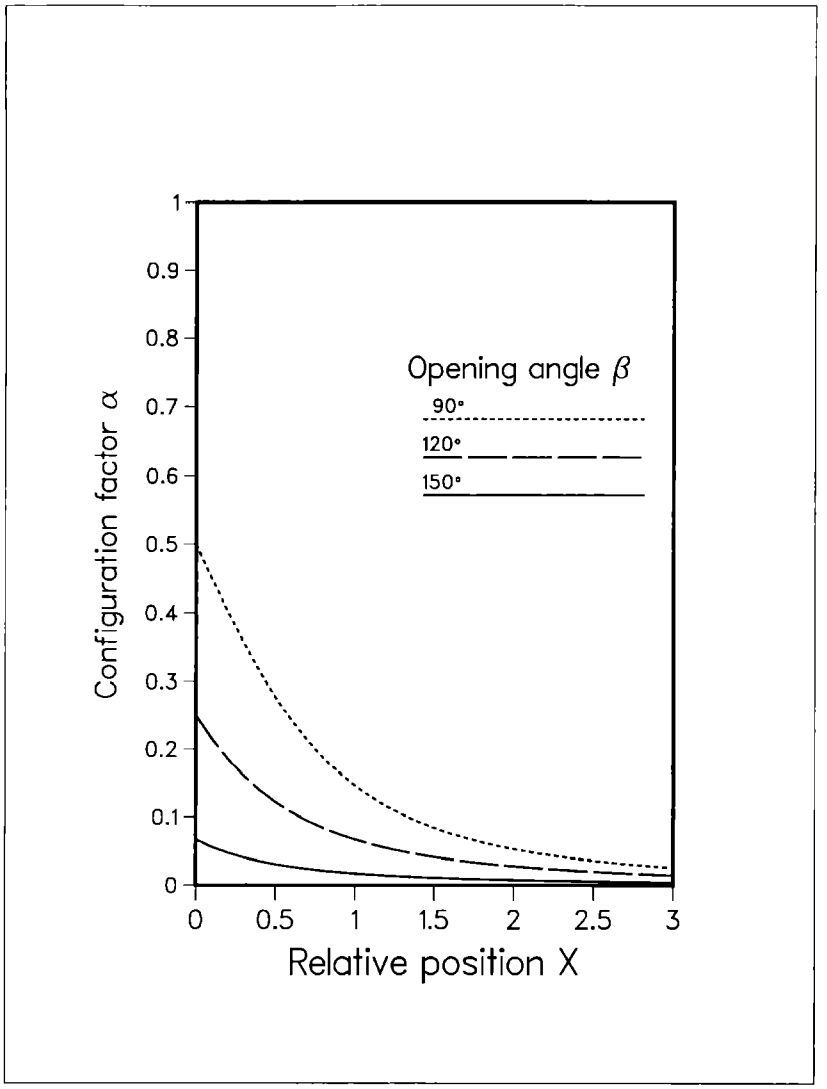

Fig. 4. Configuration factor α as a function of $X = x/\zeta$ for several opening angles β.

5.1 Two-zone Edge

In this section, we use equations (2) to construct an edge. For surfaces we choose any pair of reflectances from the collection of 370 surface spectral reflectance measurements catalogued by Krinov (1947). From Ho et al. (1990) we know how the finite dimensional linear model of reflectance performs for these functions and we start by choosing a pair of reflectances that are typical in terms of the accuracy of the underlying model. We choose Krinov #53 ("heather, dense growth before flowering") and Krinov #54 ("river valley with meadows") as typical curves. For these reflectances, the underlying model with 3 basis functions taken from Cohen gives errors of 11.04% and 8.11%, respectively. Here, we define an error statistic that measures departures from the correct curve over the entire function of wavelength. It is essentially a rescaled Standard Error divided by the sample root mean square in order to make it a fraction rather than an absolute figure. Error is defined as

$$\left\{ \frac{\Sigma(y - y')^2}{\Sigma y^2} \right\}^{1/2}$$

where y is an actual value and y' is the estimate.

For illumination, we can use one of the standard daylights tabulated by Judd et al. (1964) or any other tabulated daylight, for example, those measured by Dixon (1978). Dixon included ultraviolet wavelengths

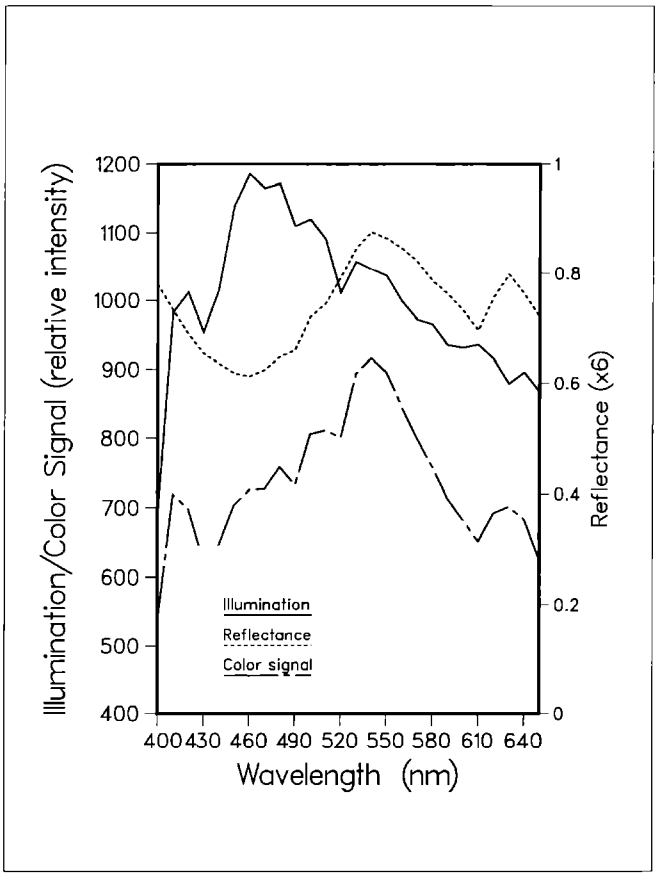

Fig. 5. Typical case: color signal formed from product of surface spectral reflectance Krinov #54 and Dixon spectrum for Australian daylight. The color signal has been scaled up by a factor of 6.

in developing a principal-component analysis for Australian daylight, but here we truncate to the visible, 400 nm–650 nm. The spectrum for the mean vector for daylight in Bendigo, Australia, is shown in figure 5 along with Krinov reflectance curve #54 and the resulting color signal composed of the product of the two spectra. For the domain D of initial value ϵ_2, ϵ_3's to use in the uniqueness check preprocessing step we took the largest absolute values of ϵ_2 and ϵ_3 for all of Judd's curves and then doubled that region. We suppose for simplicity that the shape and geometry of the two surfaces leads to geometrical factors of $\alpha_{21} = \alpha_{12} \equiv \alpha = 0.05$ for mutual reflection.

The algorithm of section 3 applied to the two color signals generated from this illumination spectrum and Krinov reflectances #53 and #54 can be tested by forming sensor values $\rho_k^{(1)}$, $\rho_k^{(2)}$, and ρ_k^{m}. To develop these signals, we used sensitivity functions (see figure 6) corresponding to Kodak filters #25 (red), #58 (green), and #47B (blue). The algorithm results are quite good: the value of α found by the solution method is 0.0495, so that the ratio of the estimated α to the actual one is 0.990. As pointed out above, using the approximate model algorithm the accuracy for this ratio remains nearly constant no matter what the actual value of α is.

To develop an error figure for the spectral reflectance function, it is important to recall that the algorithm gives a solution only up to an overall multiplicative constant. Therefore, the shape of the reflectance curve is derived, but not its absolute scale. To make a fair comparison, we scale by an appropriate factor and then calculate the error statistic. A best match of the derived curve and the actual one is found by solving for the multiplicative factor by minimizing the squared residuals after multiplying by the unknown factor. The normal equation for this minimization problem yields the factor.

As shown in figure 7, the spectral curve for reflectance 2 is quite close to that for Krinov #54; the overall error for the curve found by our algorithm is 9.71%.

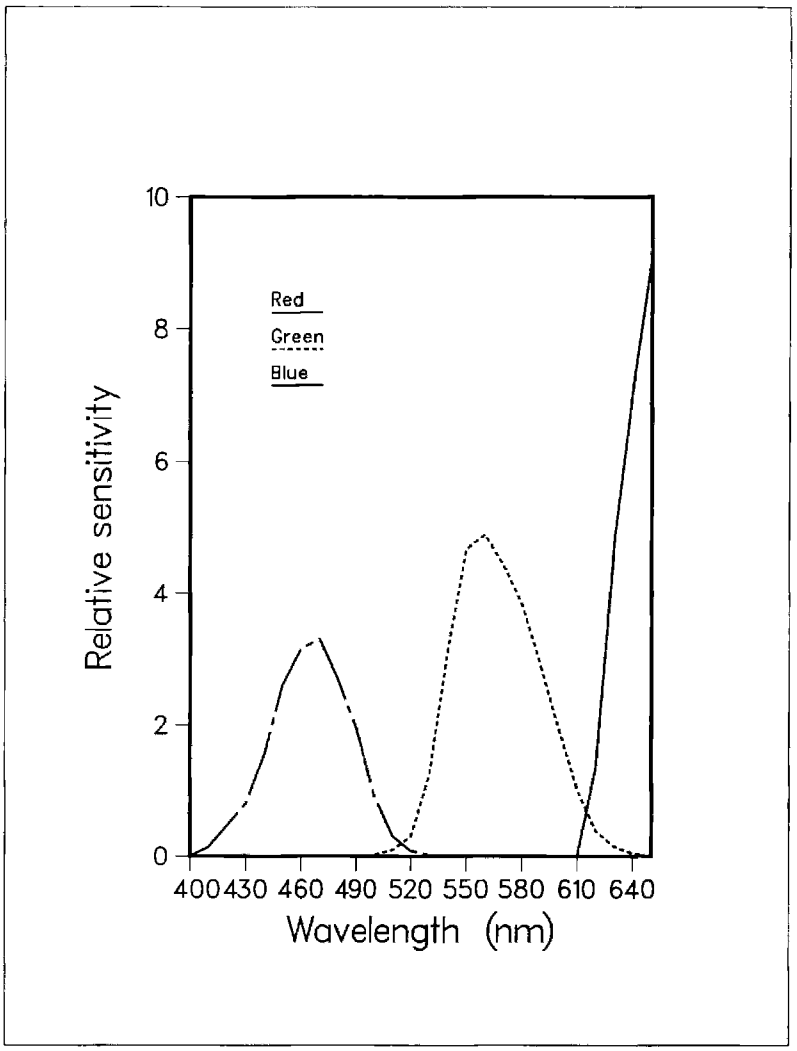

Fig. 6. Kodak filters #25 (red), #58 (green), and #47B (blue).

Since the error of the best fit using the underlying finite dimensional model is 8.11%, the ratio of our model error to the finite dimensional model error is only 1.197. The fit for surface 1, which has reflectance Krinov #53, is similar. We show in figure 8 the results of the algorithm in recovering the illumination compared to the actual spectrum. Again, the increase in error is quite small—we obtain an error of 4.10% as compared to the finite dimensional model error of 3.97%.

In table 1, we show the results of running the algorithm on sensor values derived from the same reflectances as above for all five of Judd's standard daylights, corresponding to five different correlated color temperatures, as well as the Dixon spectrum. As can be seen, the results are similar for all illuminations. As well, uniqueness was validated by the preprocessing step in these cases.

A more stringent test is carried out by utilizing atypical reflectance curves that are not well modeled by the basis functions. For this test we use Krinov #162 ("grass, young, green") along with Krinov #54. A best fit using the finite dimensional model itself has an error of 26.65% for Krinov #162. The results for the algorithm are again reasonably good, as seen in table 2.

To take into account noise, we add a random perturbation to each ρ-vector separately; each component independently receives additive noise equal to a percentage of the vector magnitude. For this test we use the Dixon spectrum coupled to reflectances Krinov #53 and Krinov #54. Running the algorithm repeatedly, we

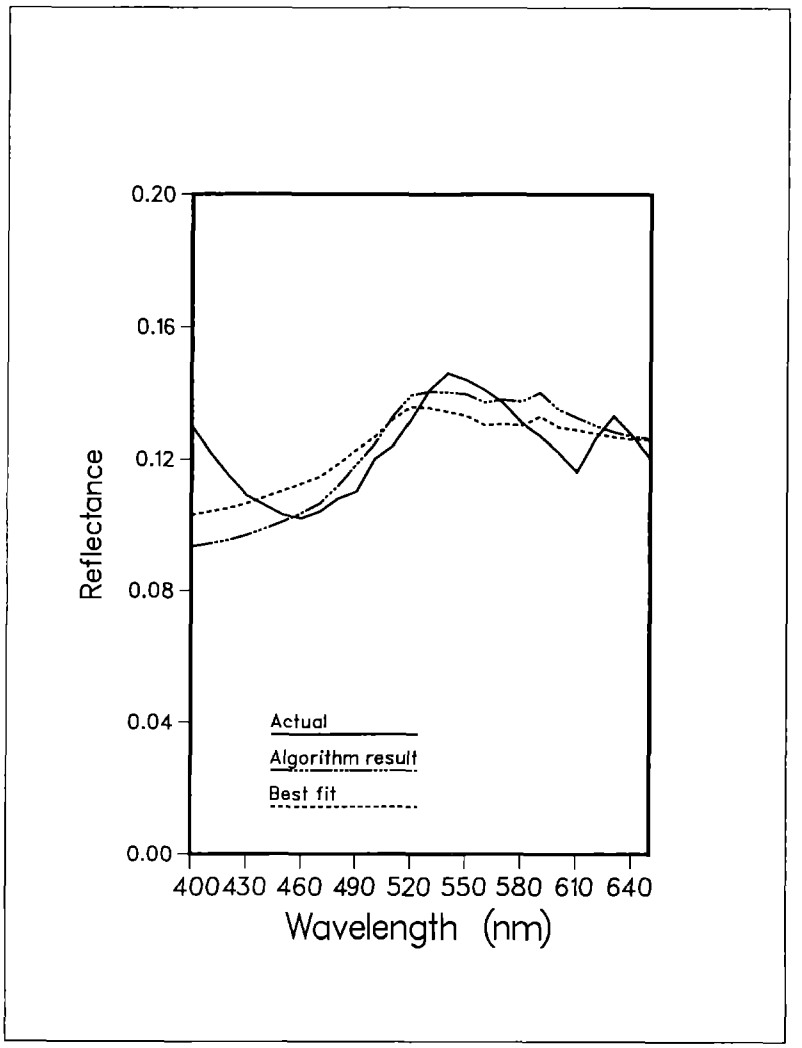

Fig. 7. Comparison of the best fit given by dimension-3 finite dimensional model with surface spectral curve resulting from the algorithm. The best-fit error (here, the reflectance curve is Krinov #54) has error = 8.11%, whereas the algorithm result has error = 9.71%, a factor of 1.197 worse than the best fit. When the geometrical mutual reflection factor α is 0.05, the algorithm gives the value 0.04949, so that the ratio of the estimate to the actual α is 0.9898.

arrive at average results, shown in table 3. Here, the 0% noise case is shown for comparison and all other figures are averages. For a few very poorly fit noisy signals at the higher noise level, the algorithm is less stable and a greater tolerance must be used for halting iteration. For each of these cases the uniqueness check also showed that the results were unreliable. So for a large noise component the results are not good. Nevertheless, the results are reasonably good for low noise, and in that regime the stability of the algorithm is very good. The number of iterations through the algorithm was typically 6 to 8.

If use is made in our model of broad-band sensor functions such as those in the human visual system, instead of bandpass filters common in video and film, then the results of the algorithm are found to be somewhat more accurate since a broader sampling of more frequencies results in a better representation of the entire color signal spectrum in terms of the finite dimensional model weights.

5.2 Infinite-zone Edge

In Drew and Funt (1990), we simulated the full physics (infinite-zone, infinite-bounce) of the interreflection at

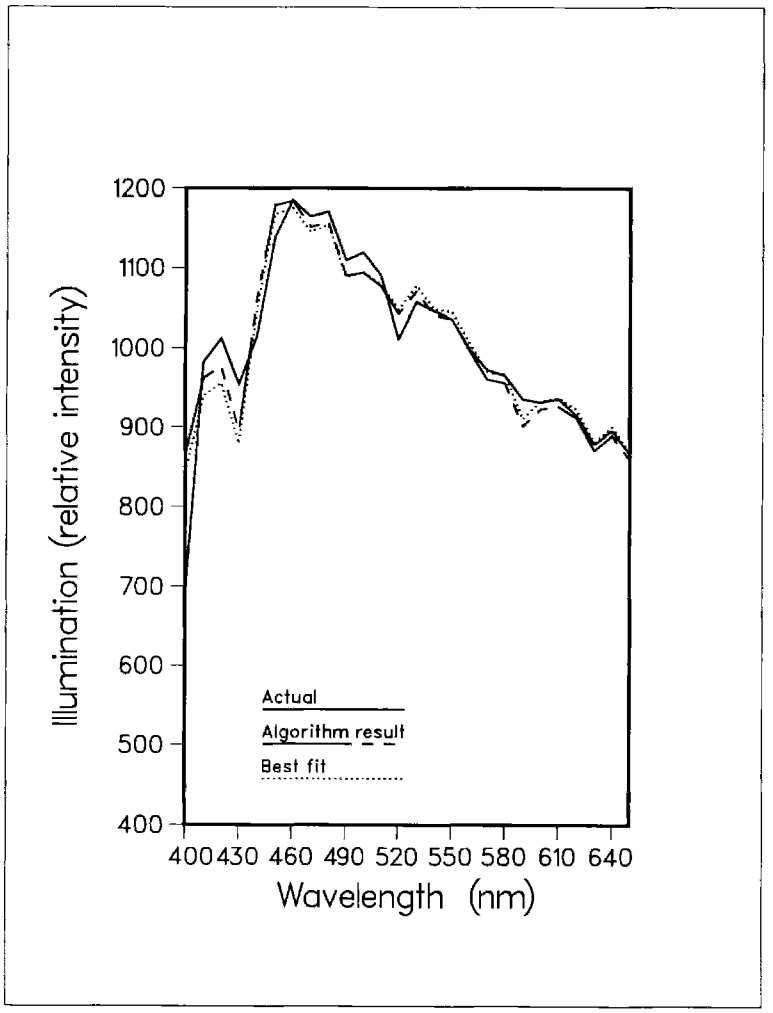

Fig. 8. The original Dixon illumination spectrum for Australian daylight and results of the mutual reflection algorithm applied to the color signal of figure 5—error 4.10%.

an edge under diffuse illumination by employing the radiosity method (Goral et al. 1984). However, unlike the usual approach in computer graphics, we carried out the radiosity calculation separately on each wavelength instead of on the usual three color bands and therefore could not use standard computer graphics radiosity tools to synthesize test images. In generating a correct color distribution at the edge it is crucial that one apply the filter functions to generate RGB values only after calculating the entire color signal spectrum for each pixel.

We used this method on the edge of figure 2, using Krinov spectra #54 and #248 and incident illumination D65. Figure 9 shows the variation in RGB on surface 2 that occurs with increasing distance from the edge. For comparison, the (constant) RGB values resulting from $S^{(2)}(\lambda)E(\lambda)$ in the absence of interreflection are also plotted. The ratios R/B and G/B graphed in figure 10 show that not only the intensities, but the colors too change with distance from the edge. We can see from figure 10 that derivatives of ratios approach zero as one gets farther from the edge.

Running our mutual reflection analysis algorithm on the calculated RGB image data, we find that reflectances on each side are recovered quite accurately. Except at pixels very close to the edge, the algorithm obtains a reflectance $S^{(2)}(\lambda)$ that is virtually independent of position. The error with respect to the actual reflectance spectrum is 4.37%, which translates into a color difference ΔE of 3.54 units in CIELUV uniform color space (see Wyszecki and Stiles 1982). For surface 2, the best least-squares approximation using the finite

Table 1. Errors using typical reflectances: Krinov #53 and #54, $\alpha = 0.5$, and Judd's North American daylights of five different correlated color temperatures as well as Dixon's Australian daylight. The finite dimensional model best curves have errors of 11.04% for Krinov #53 and 8.11% for Krinov #54. Columns show: (1) α recovered by algorithm; (2) ratio of the recovered α to the actual one; (3) error for illumination spectrum found by algorithm; (4) absolute error of reflectance curve found by algorithm for surface 1; (5) error of reflectance curve for surface 1 compared to error in underlying finite dimensional model; (6, 7) similarly for reflectance curve for surface 2. For the Dixon case, the finite dimensional model illumination modeling error is 3.97%, so that the ratio $E(\lambda)$ Error / $E(\lambda)$ $\text{Error}_{\text{FDM}}$ is 1.032.

Daylight	α	$\frac{\alpha_{EST}}{\alpha_{ACT}}$	$E(\lambda)$ Error (%)	$R^{(1)}(\lambda)$ Error (%)	$\frac{R^{(1)}(\lambda)\text{ Error}}{R^{(1)}(\lambda)\text{ Error}_{\text{FDM}}}$	$R^{(2)}(\lambda)$ Error (%)	$\frac{R^{(2)}(\lambda)\text{ Error}}{R^{(2)}(\lambda)\text{ Error}_{\text{FDM}}}$
4800K	0.04971	0.9943	0.30	11.35	1.028	9.81	1.209
5500K	0.04971	0.9942	0.29	11.34	1.027	9.79	1.206
6500K	0.04973	0.9945	0.31	11.34	1.027	9.75	1.202
7500K	0.04972	0.9944	0.31	11.32	1.026	9.73	1.200
10000K	0.04972	0.9944	0.31	11.33	1.026	9.70	1.196
Dixon	0.04949	0.9898	4.10	11.34	1.027	9.71	1.197

Table 2. Poorly fit case: reflectances Krinov #162 and #54, $\alpha = 0.05$—errors for five Judd daylights and for Dixon daylight (best model fit has error 26.65% for Krinov #162 and 8.11% for Krinov #54); α recovered by algorithm; compared to correct α; illumination spectrum error found by algorithm; reflectance curve 1 error; compared to error in underlying finite dimensional model; similarly for reflectance curve 2. For the Dixon case, the finite dimensional model illumination modeling error is 3.97%, so that the ratio $E(\lambda)$ Error / $E(\lambda)$ $\text{Error}_{\text{FDM}}$ is 1.428.

Daylight	α	$\frac{\alpha_{EST}}{\alpha_{ACT}}$	$E(\lambda)$ Error (%)	$R^{(1)}(\lambda)$ Error (%)	$\frac{R^{(1)}(\lambda)\text{ Error}}{R^{(1)}(\lambda)\text{ Error}_{\text{FDM}}}$	$R^{(2)}(\lambda)$ Error (%)	$\frac{R^{(2)}(\lambda)\text{ Error}}{R^{(2)}(\lambda)\text{ Error}_{\text{FDM}}}$
4800K	0.04988	0.9975	3.47	32.29	1.212	11.13	1.371
5500K	0.04987	0.9974	3.41	32.11	1.205	11.04	1.361
6500K	0.04988	0.9976	3.37	31.93	1.198	10.95	1.349
7500K	0.04987	0.9974	3.19	31.83	1.194	10.89	1.342
10000K	0.04987	0.9974	2.86	31.68	1.189	10.80	1.331
Dixon	0.04964	0.9929	5.66	31.95	1.199	10.90	1.343

Table 3. Noise: reflectances Krinov #53 and #54, Dixon spectrum, $\alpha = 0.05$—finite dimensional model best fits have $E(\lambda)$ $\text{Error}_{\text{FDM}}$ = 3.97%, $R^{(1)}(\lambda)$ $\text{Error}_{\text{FDM}}$ = 11.04%, and $R^{(2)}(\lambda)$ $\text{Error}_{\text{FDM}}$ = 8.11%. Columns show: α recovered by algorithm; compared to correct α; illumination spectrum error found by algorithm; reflectance curve 1 error; reflectance curve 2 error.

Noise (%)	α	$\frac{\alpha_{\text{EST}}}{\alpha_{\text{ACT}}}$	$E(\lambda)$ Error (%)	$R^{(1)}(\lambda)$ Error (%)	$R^{(2)}(\lambda)$ Error (%)
0	0.04983	0.9967	4.13	11.33	9.75
5	0.04385	0.8771	17.13	16.80	15.05
10	0.03912	0.7823	38.50	29.030	28.34

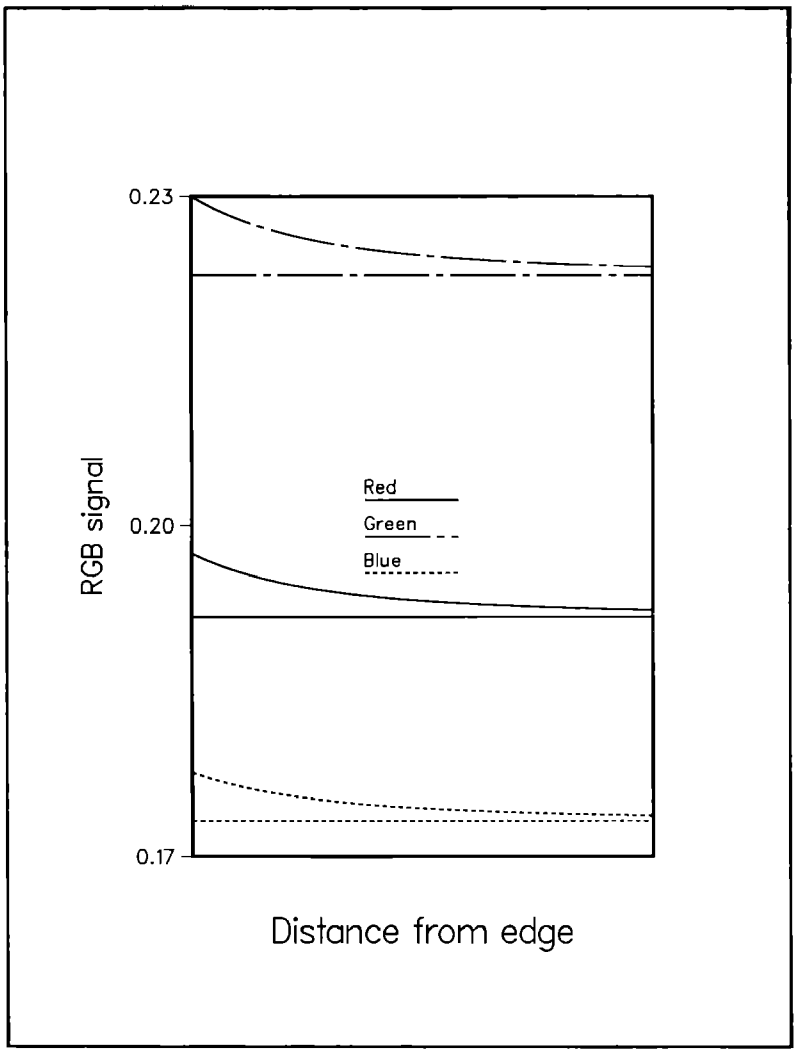

Fig. 9. RGB variation on one side of an edge due to mutual reflection.

dimensional model with dimensionality 3 gives an error of $\Delta E = 3.25$; clearly, the mutual reflection method contributes only a small additional error. For illumination, the error of the recovered spectrum is virtually zero.

The recovered values of α vary as shown in figure 11. Here we have scaled α to take into account the scaling of the illumination E that gives $\epsilon_1 = 1$. For comparison we also show the theoretical configuration factor α as determined by equation (10). We display that factor offset so that it goes to zero and not to the small value it takes on at the pixel farthest from the edge. As one would expect since the recovered α maps the contributions from all bounces into a one-bounce factor, it is marginally larger than the theoretical α generated using equation (10) that models an infinite-zone, one-bounce situation.

6 Conclusions

We have examined the effects of mutual reflection in color images. Since mutual reflection causes the spectrum of the light reflected from a surface to vary with location on the surface even though the surface reflectance is constant, it provides clues to the actual surface reflectance and to the geometry of the situation. Tests show that the algorithm is quite robust.

The method has the potential for eliminating mutual reflection effects and in most cases recovers a good approximation of the complete spectral information. Thus the method yields more than enough information to supply color descriptors for the object surfaces that are independent of ambient illumination; therefore, color constancy is obtained. Sensor values from only three

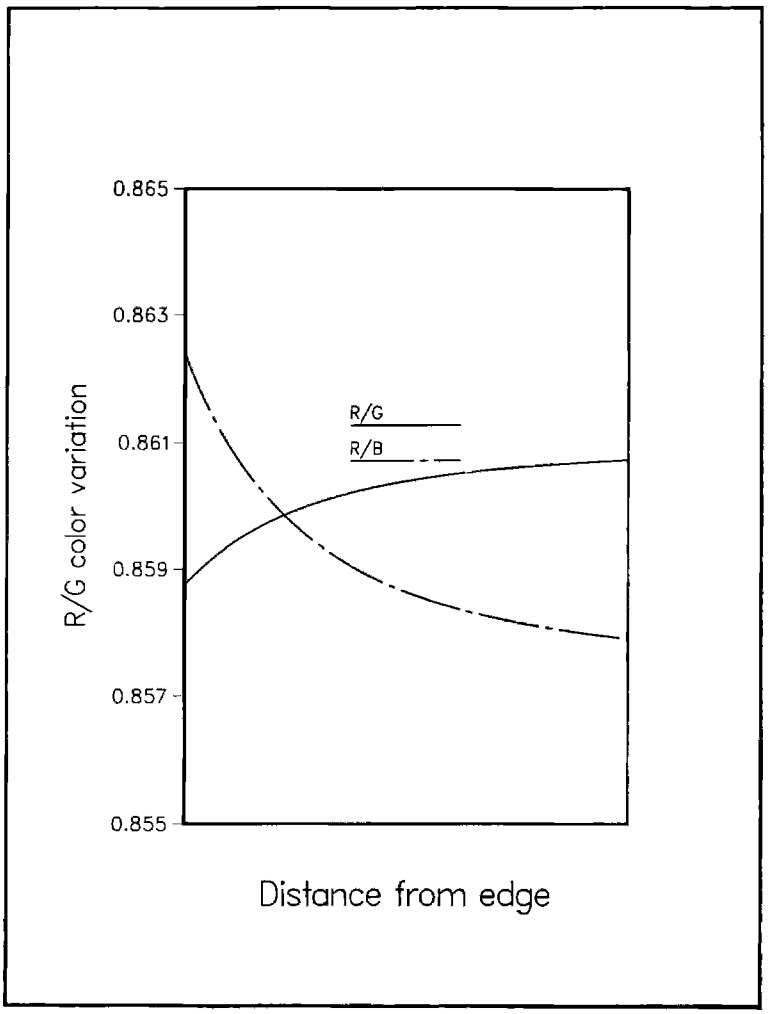

Fig. 10. Color variation: ratios of R/B and G/B for figure 9.

pixels need to be measured to obtain a result and the method will work on surfaces that do not meet the surface complexity conditions needed to make Maloney and Wandell's algorithm work. The elimination of mutual reflection is a desirable feature in a vision system since it is an effect that introduces extra complexities; that is, it is known to create problems in shape-from-shading schemes (Forsyth & Zisserman 1989, 1990).

Clearly, the interreflection model put forward here, that uses basis function decompositions for an accurate whole-spectrum model of reflection, can be extended to the modeling of multisurface enclosures for graphics applications (cf. Goral et al. 1984). Multiple interreflections could be incorporated at the expense of having a more difficult mathematical system to solve; for enclosures, such multiple bounces may add a significant contribution. Ignoring multiple bounces for 3, 4, etc., surfaces results in a simple set of linear equations similar to the linear equations for $I_m(\lambda)$ in section 2.

The algorithm still works even if both surfaces have the *same* color, with the same or differing brightness. This case might correspond to a folded piece of uniform material. The recovery of α here gives shape information near the fold. The method works in this one-color case because the reflected light is a different color from the illumination, and the extra measurements of sensor values supply enough information to disentangle the two.

The results show that, in situations in which mutual reflection is present, measurements of the sensor values taken separately from the two different interreflecting

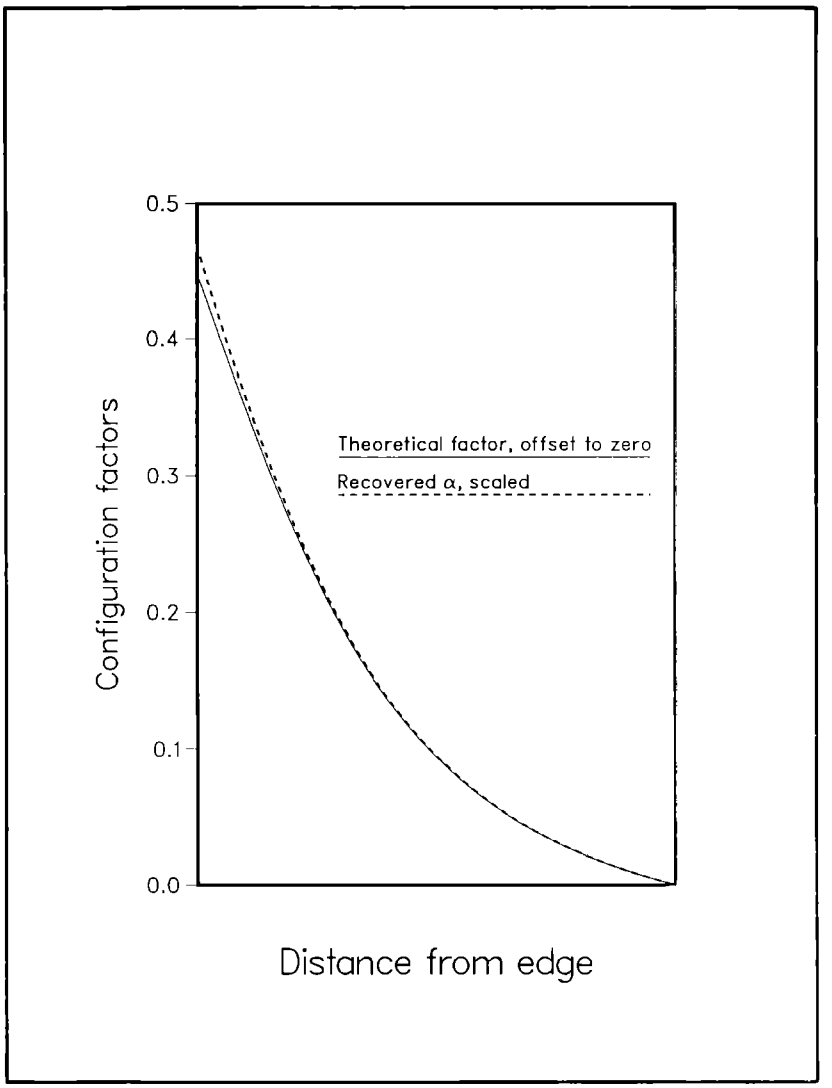

Fig. 11. Variation of one-bounce, two-zone α factor, where the sensor RGB inputs are derived from a full multi-zone model. Theoretical one-bounce, infinite-zone factor is shown for comparison.

surfaces, as well as from within the mutual reflection zone, can convey sufficient information to allow complete recovery of the two surface spectral reflectance functions as well as of ambient illumination. We also recover the factor α, which summarizes information about the relative geometry of the two surfaces.

7 Acknowledgments

M.S. Drew is indebted to the Centre for Systems Science at Simon Fraser University for partial support; B.V. Funt thanks both the CSS and the Natural Sciences and Engineering Research Council of Canada for their support.

References

Beck, J. 1972. *Surface Color Perception*, Cornell Univ. Press: Ithaca, N.Y.

Blackwell, K.T., and Buchsbaum, G. 1988. Quantitative studies of color constancy. *J. Opt. Soc. Amer. A.*, 5:1772–1780.

Brainard, D.A., Wandell, B.A., and Cowan, W.B. 1989. Black light: How sensors filter spectral variation of the illuminant. *IEEE Trans. Biomed. Eng.* 36:140–149, 572.

Brill, M.H., and West, G. 1986. Chromatic adaptation and color constancy: A possible dichotomy. *Col. Res. Appl.* 11:196–204.

Buchsbaum, G. 1980. A spatial processor model for object colour perception. *J. Franklin Inst.* 310:1–26.

Burden, R.L., Faires, J.D. and Reynolds, A.C. 1981. *Numerical Analysis.* Prindle, Weber & Schmidt: Boston.

Cohen, J. 1964. Dependency of the spectral reflectance curves of the Munsell color chips. *Psychon. Sci. Amer. A*, 6:318–322.

Dixon, E.R. 1978. Spectral distribution of Australian daylight. *J. Opt. Soc. Amer.* 68:437–450.

Drew, M.S., and Funt, B.V. 1990. Calculating surface reflectance using a single-bounce model of mutual reflection. *Proc. Intern. Conf. Comput. Vision*, Osaka, December 4–7.

D'Zmura, M. and Lennie, P. Mechanisms of color constancy. *J. Opt. Soc. Amer.* (10):1662–1672.

Forsyth, D., and Zisserman, A. 1989. Mutual illumination. *Proc. conf. Comput. Vision Patt. Recog.*, San Diego, pp. 466–473.

Forsyth, D., and Zisserman, A. 1990. Shape from shading in the light of mutual illumination. *Image Vision Comput.* 8:42–49.

Funt, B.V., and Ho, J. 1988. Color from black and white. *Proc. 2nd Intern. Conf. Comput. Vision*, pp. 2–8, Dec. 5–8, Tarpon Springs, FL.

Funt, B.V., and Ho, J. 1989. Color from black and white. *Intern. J. Comput. Vision* 3(2).

Gershon, R., Jepson, A.D, and Tsotsos, J.K. 1986. Ambient illumination and the determination of material changes. *J. Opt. Soc. Amer.* 3:1700–1707.

Gershon, R., Jepson, A.D. and Tsotsos, J.K. 1987. From [R,G,B] to surface reflectance: Computing color constant descriptors in images. *Proc. 10th IJCAI*, Milan, pp. 755–758.

Goral, C.M., Torrance, K.E., Greenberg, D.P. and Battaile, B. 1984. Modeling the interaction of light between diffuse surfaces. *Computer Graphics* 18:213–222.

Ho, J., Funt, B.V., and Drew, M.S. 1990. Separating a color signal into illumination and surface reflectance components: Theory and applications. *IEEE Trans. Patt. Anal. Mach. Intell.* 12:966–977.

Horn, B.K.P. 1977. Understanding image intensities. *Artificial Intelligence* 8:201–231.

Horn, B.K.P.. 1986. *Robot Vision*. MIT Press: Cambridge, MA.

Judd, D.B., MacAdam, D.L., and Wyszecki, G. 1964. Spectral distribution of typical daylight as a function of correlated color temperature. *J. Opt. Soc. Amer.* 54:1031–1040, August.

Krinov, E.L. 1947. Spectral reflectance properties of natural formations. *Technical Translation* TT–439, National Research Council of Canada.

Koenderink, J.J., and van Doorn, A.J. 1983. Geometrical modes as a general method to treat diffuse interreflections in radiometry. *J. Opt. Soc. Amer.* 73:843–850.

Maloney, L.T. 1985. Computational approaches to color constancy, Ph.D. dissertation, Stanford University, Applied Psychology Laboratory.

Maloney, L.T. 1986. Evaluation of linear models of surface spectral reflectance with small numbers of parameters. *J. Opt. Soc. Amer.* A, 3:1673–1683.

Maloney, L.T., and Wandell, B.A. 1986. Color constancy: a method for recovering surface spectral reflectance. *J. Opt. Soc. Amer.* A, 3:29–33.

Nayar, S.K., Ikeuchi, K., and Kanade, T. 1990. Shape from interreflections. *Proc. Intern. Conf. Comput. Vision*, Osaka, December 4–7.

Parkkinen, J.P.S., Hallikainen, J., and Jaaskelainen, T. 1989. Characteristic spectra of Munsell colors. *J. Opt. Soc. Amer.* A, 6:318–322.

Siegel, R., and Howell, J.R. 1981. *Thermal Radiation Heat Transfer.* Hemisphere Publishing Corp: New York.

Sparrow, E.M., and Cess, R.D. 1978. *Thermal Radiation Heat Transfer.* Hemisphere Publishing Corp.: New York.

Wandell, B.A. 1987. The synthesis and analysis of color images. *IEEE Trans. Patt. Anal. Mach. Intell.* 9:2–13.

Wyszecki, G., and Stiles, W.S. 1982. *Color Science: Concepts and Methods, Quantitative Data and Formulas*. Wiley: New York.

Yuille, A. 1987. A method for computing spectral reflectance. *Biological Cybernetics* 56:195–201.