

Ambient illumination and the determination of material changes

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

Department of Computer Science, University of Toronto, Toronto, Ontario M5S 1A4, Canada

Received February 3, 1986; accepted May 20, 1986

The task of distinguishing material changes from shadow boundaries in chromatic images is discussed. Although there have been previous attempts at providing solutions to this problem, the assumptions that were adopted were too restrictive. Using a simple reflection model, we show that the ambient illumination cannot be assumed to have the same spectral characteristics as the incident illumination, since it may lead to the classification of shadow boundaries as material changes. In such cases, we show that it is necessary to take into account the spectral properties of the ambient illumination in order to develop a technique that is more robust and stable than previous techniques. This technique uses a biologically motivated model of color vision and, in particular, a set of chromatic-opponent and double-opponent center-surround operators. We apply this technique to simulated test patterns as well as to a chromatic image. It is shown that, given some knowledge about the strength of the ambient illumination, this method provides a better classification of shadow boundaries and material changes.

INTRODUCTION

The perception of color has been a subject of research in many disciplines, including image understanding. One approach to color-image understanding has mainly involved an application of algorithms originally developed for achromatic images to three-dimensional color space. These include edge detection,¹ clustering,²⁻⁴ region splitting,⁵ and region growing.^{6,7} These schemes regard color as a random variable that is analyzed statistically without regard to any model of specific physical processes of color generation. A second approach involves the explicit use of knowledge about the generation of color. One example is the work of Sloan,⁸ which shows that simple statements can be made about outdoor shadows and object colors. Similarly, Rubin and Richards^{9,10} demonstrated that by analyzing processes that cause changes in color, one can distinguish material changes from shadow boundaries in some situations. Finally, Shafer¹¹ presents a theoretical analysis of highlights and object-color reflection that provides a way to remove highlights from portions of images.

In this paper we consider the task, originally proposed by Rubin and Richards, of distinguishing material changes from shadow boundaries in chromatic image. A simple reflectance model is used to illustrate the effect on shadow boundaries of an ambient illumination having a different spectral power distribution from the direct illumination. It is clear from this model that in order to distinguish shadow boundaries reliably from material changes solely on the basis of local measurements of reflected light, some information about the ambient illumination is needed. In particular, it is shown that the technique proposed by Rubin and Richards¹⁰ can incorrectly label a shadow boundary as a material change when the spectral properties of the ambient illumination are sufficiently different from those of the direct illumination. Here we present a new method for labeling shadow boundaries in the presence of ambient illumination. The method is based on the output of biologically motivated

operators and on an *a priori* estimate for the strength of the ambient illumination.

A SIMPLE REFLECTANCE MODEL

The appearance of objects is a result of interaction between illumination and reflectance. Thus it is necessary to describe briefly the physics of reflection since it governs the appearance of these regions. In this section we present a simplified model of reflectance that is used in both computer graphics and image understanding⁹⁻¹² and that captures the important aspects of the reflection process.

In order to describe the spatial and chromatic properties of the reflected light, we use a reflection model composed of three components: ambient, diffuse, and specular. These reflectance components are properties of any illuminated object. In addition, we assume that there are two types of illumination: ambient (or indirect) and incident (or direct). The ambient illumination represents light incident from the environment, i.e., reflections of the direct light from all the materials in the scene, while the incident illumination is the light originating from specific light sources with no intermediate reflections. The product of illumination and reflectance yields the reflected intensity of the object. Therefore the ambient reflected intensity is the result of the interaction between the ambient illumination and the surface of the object; the diffuse reflected intensity reflects the light that is scattered generally equally in all directions and represents the color of the object itself, since it is the result of incident illumination interacting with the pigments within the object; the specular reflected intensity represents the highlights, the glossy reflection of objects, which is the result of the incident illumination's bouncing off the surface. Therefore, when all three components are taken into account, the image intensity of the reflected light $[I_{ref}(\lambda, \mathbf{x})]$ at a given point, given some illumination and viewing geometry, is formulated in a simple form as

$$I_{\text{ref}}(\lambda, \mathbf{x}) = I_{\text{ambient}}(\lambda, \mathbf{x})\rho_{\text{ambient}}(\lambda, \mathbf{x}) + I_{\text{incident}}(\lambda, \mathbf{x}) \times [\delta\rho_{\text{specular}}(\lambda, \mathbf{x}) + (1 - \delta)\rho_{\text{diffuse}}(\lambda, \mathbf{x})], \quad (1)$$

where $I_{\text{ambient}}(\lambda, \mathbf{x})$ and $I_{\text{incident}}(\lambda, \mathbf{x})$ are the intensities of the ambient and incident illuminations, respectively; $\rho_{\text{ambient}}(\lambda, \mathbf{x})$, $\rho_{\text{specular}}(\lambda, \mathbf{x})$, and $\rho_{\text{diffuse}}(\lambda, \mathbf{x})$ are the ambient, specular, and diffuse reflectances, respectively; and δ is a fraction of 1 expressing the magnitude of the surface specularity. Because the reflected ambient intensity is a result of interaction of some illumination (ambient) with the surface of the object, it has often been assumed that the ambient reflectance [$\rho_{\text{ambient}}(\lambda, \mathbf{x})$] can be represented as a linear combination of the diffuse and specular reflectances.^{11,12} (It is important to emphasize that this assumption is true for the ambient reflectance and not for the ambient illumination.) We can therefore replace the term that represents the ambient reflectance with a combination of the diffuse and specular ones. Furthermore, if we confine our analysis to regions that do not contain highlights, that is, ones in which the proportion of specular reflectance is almost constant (i.e., $\delta \approx \text{constant}$ and is usually close to zero), we can represent the reflectance in one term, which will be denoted as $\rho(\lambda, \mathbf{x})$. The formulation for the intensity of the reflected light is therefore

$$I_{\text{ref}}(\lambda, \mathbf{x}) = I_{\text{ambient}}(\lambda, \mathbf{x})\rho(\lambda, \mathbf{x}) + I_{\text{incident}}(\lambda, \mathbf{x})\rho(\lambda, \mathbf{x}) = [I_{\text{ambient}}(\lambda, \mathbf{x}) + I_{\text{incident}}(\lambda, \mathbf{x})]\rho(\lambda, \mathbf{x}), \quad (2)$$

where $\rho(\lambda, \mathbf{x})$ is the reflectance of the object. From this analysis, it is easy to see that if there is no direct illumination, that is, $I_{\text{incident}}(\lambda, \mathbf{x}) = 0$, then the intensity of the reflected light at a given point will be

$$I_{\text{ref}}(\lambda, \mathbf{x}) = I_{\text{ambient}}(\lambda, \mathbf{x})\rho(\lambda, \mathbf{x}), \quad (3)$$

which represents the intensity of reflected light at a point that is not directly illuminated, hence a point in a shadow region. We therefore see that Eq. (2) represents the intensity at the lit region, while Eq. (3) represents the intensity at the shadow region.

IDEAL AND NONIDEAL SHADOWS

The common factor contributing to the appearance of both the shadow and lit regions of objects is the ambient illumination. This implies that understanding the effects of the ambient illumination on objects may give us important information regarding the disambiguation of these regions, enabling us to distinguish between them. In this section, we examine two conditions under which the relationships between shadow and lit regions change. In the first case, which we call the "ideal" case, we assume that the ambient illumination has the same spectral characteristics as the incident illumination. In the second case, which we call "nonideal," we assume that the ambient illumination differs in its spectral characteristics from the incident illumination. We use the term "chromatic components" to denote the responses of particular chromatic mechanisms. For example, the red component, which is denoted as R , is defined by

$$R(\mathbf{x}) = \int_0^{\infty} I_{\text{ref}}(\lambda, \mathbf{x})S_R(\lambda)d\lambda, \quad (4)$$

where $I_{\text{ref}}(\lambda, \mathbf{x})$ is the reflected intensity reaching the sensor (cone) and $S_R(\lambda)$ is the spectral sensitivity function of the red cone. For simplicity, we confine our discussion to the red and green components of the reflected intensities. The extension of the work to include a third chromatic component is straightforward.

In order to show how the ambient illumination affects objects, we first show that conditions hold for the ideal case and then extend the discussion to the nonideal case. We make use of Fig. 1 to illustrate the two cases. Suppose the original setup consists of an illuminant (I) that illuminates the observed object (O). Some of the light does not reach the object O because of a blocking object (B), thus creating a shadow region (S) and a lit region (L) on the observed object O. In the ideal case, the ambient illumination has the same spectral characteristics as the incident illumination. In other words, there is a difference only in the magnitude of the spectral distribution functions representing the incident and ambient illuminations, such that for any given wavelength the ambient illumination will be a fraction of 1 (denoted as α) of the incident illumination. Having assumed that the ambient reflectance is the same as the diffuse reflectance, we find that the ambient reflected intensity has the same spectral characteristics as the diffuse reflected intensity up to the same constant α . In such a case, if the total intensity of the reflected light in the lit regions of object O were the sum of the ambient and diffuse reflected intensities [as given by Eq. (2)], then we can represent the red and green components of this reflected intensity as $(\alpha R + R, \alpha G + G)$. Under ideal conditions, the red and green components of the total reflected intensity from the shadow parts of O (which are illuminated by ambient illumination only) are $(\alpha R, \alpha G)$.

One of the conditions under which an observed object is exposed to ambient illumination that does not have the same spectral characteristics as the incident illumination occurs when there is another object (or objects) casting its color on

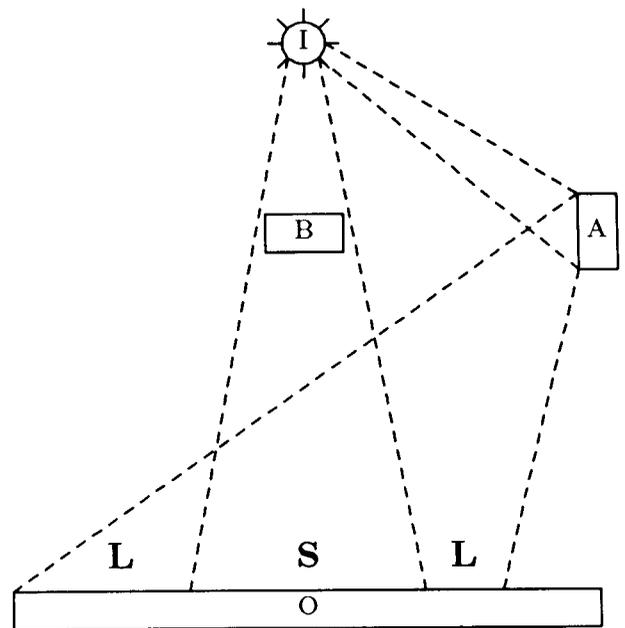


Fig. 1. A case in which the ambient illumination impinging upon an object does not have the same spectral characteristics as the incident illumination.

Table 1. The Response of the Red and Green Mechanisms to the Reflected Intensities of the Shadow and Lit Regions Under Ideal and Nonideal Lighting Conditions

Condition	Reflected Intensity	
	Shadow Region	Lit Region
Ideal	$(\alpha R, \alpha G)$	$(\alpha R + R, \alpha G + G)$
Nonideal	$(r + \alpha R, \alpha G)$	$(r + \alpha R + R, \alpha G + G)$

the observed object. Suppose that there is another object (A) illuminated by I, which casts its color onto O. The object A need not be physically close to O nor necessarily exist in the image. It is used merely to illustrate that the ambient illumination may be biased toward some color other than the incident illumination. Furthermore, let us assume that the color object that A adds is red and that this red color impinges upon both the shadow and the lit regions and is thus reflected off them, since it acts as an illuminant. We therefore have to account for a new factor, which we call the *additional ambient illumination*. Because the regions belong to the same object (O), this additional ambient illumination affects both the shadow and lit regions the same, thus resulting in an additional ambient reflected intensity, of which its red component (the only component in our example) is denoted as r . Therefore the red and green components of the reflected intensity from O are $(r + \alpha R + R, \alpha G + G)$ for the lit region and $(r + \alpha R, \alpha G)$ for the shadow region. A summary of the two cases and their effects on the reflected intensities of the objects is given in Table 1. Note that it is assumed in this paper that *both* the shadow and the lit regions are influenced by the additional ambient illumination. If that were not true, and only one of the regions would have been illuminated by the additional ambient illumination, then the color of the two regions would have been very different, making the task of identifying them as part of the same object complicated.

DETERMINING MATERIAL CHANGES

To this point we have formalized the relationship between shadow and lit regions under different illumination conditions given a simple model of reflection. In this section we will examine the techniques suggested by Rubin and Richards^{9,10} for determining material changes.

Initially, Rubin and Richards⁹ suggested the *spectral crosspoint condition*, which stated that the sign of the change in amplitude of a spectral component across a shadow boundary must be the same as the sign for any other spectral component. More formally, this condition states that given two regions X and Y across a discontinuity and intensity samples I taken at two wavelengths λ_1 and λ_2 the following is a test for a material change:

$$(I_{X\lambda_1} - I_{Y\lambda_1})(I_{X\lambda_2} - I_{Y\lambda_2}) < 0. \quad (5)$$

This condition was found to be inadequate as a means of finding material changes since "in a crosspoint, spatial and spectral information are hopelessly intertwined."¹⁰ Therefore a second and independent condition was introduced, the *opposite slope sign condition*.¹⁰ The opposite slope sign condition requires that, given two spectral components mea-

sured on both sides of a shadow boundary, on each side the same component must be the maximum. Formally, the opposite slope sign condition is

$$(I_{X\lambda_2} - I_{X\lambda_1})(I_{Y\lambda_2} - I_{Y\lambda_1}) < 0, \quad (6)$$

which again is meant to signal a material change. Indeed, assuming ideal ambient lighting conditions for the image, either expression (5) or (6) can be satisfied at an ideal shadow boundary.

The question arises whether the opposite slope sign condition holds in the nonideal case. As can be observed from Table 1, the treatment of the nonideal case is complicated by the term representing the ambient illumination (in our example, r) that is added to both shadow and lit regions, which are considered to be the two sides of the discontinuity. The complications arise in cases in which the relationship between the shadow and lit regions is not proportional, thus suggesting a material change, although they may in fact be two regions of the same material. To be exact, it may be the case that $(r + \alpha R + R) < (\alpha G + G)$ in the lit region and $(r + \alpha R) > \alpha G$ in the shadow region, or vice versa, therefore causing a change in the sign of the red versus green components, which indicates a material change according to the opposite slope sign condition. Table 2 illustrates this point by using a numerical example taken from an image that will be presented in the section headed Results. Although the opposite slope sign condition indicates that there is a material change, the two samples in Table 2 were obtained from the same object in two different regions: one that was lit directly and the other in a shadow. In summary, we would like to avoid the labeling of such changes as material changes and would rather identify them as changes caused by shadows. This is the nature of the problem that we attempt to solve.

MEASUREMENTS OF THE EFFECTS OF AMBIENT ILLUMINATION

In this section we quantify how the ambient illumination affects the reflected intensity of the object. Without loss of generality, we will limit our discussion to two spectral components of the reflected intensity of the object, namely, red and green. The discussion can be similarly applied to the blue and yellow components without any further assumptions.

Let us define the vector describing the red and green components of the reflected intensity of the observed object as (R, G) . This vector describes the reflected intensity that is a result of the direct illumination. If the spectral proper-

Table 2. An Example Taken from an Image Taken Over a Shadow Boundary That Is Classified as a Material Change by the Opposite Slope Sign Condition

Reflected Intensity of Lit Region (X)	Reflected Intensity of Shadow Region (Y)
$I_{X\lambda_1} = 151$	$I_{Y\lambda_1} = 41$
$I_{X\lambda_2} = 176$	$I_{Y\lambda_2} = 16$
Opposite Slope Sign Condition: $(I_{X\lambda_2} - I_{X\lambda_1})(I_{Y\lambda_2} - I_{Y\lambda_1}) = (25)(-25) < 0$ (Material Change)	

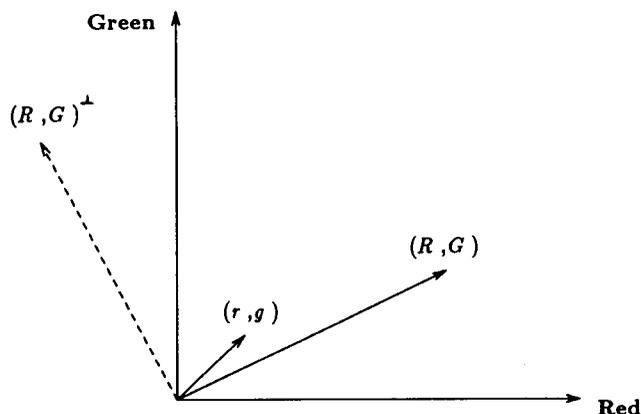


Fig. 2. The red and green components of the vectors describing the diffuse reflected intensity of an object, its normal, and the additional ambient reflected intensity.

ties of the ambient illumination are the same as the direct illumination, then the red and green components of the shadow region are proportional to the vector (R, G) . In addition, we define (r, g) , the reflected components due to the ambient illumination, as $r = \int_0^\infty I_{\text{ambient}}(\lambda, \mathbf{x}) \rho_{\text{ambient}}(\lambda, \mathbf{x}) S_R(\lambda) d\lambda$ and $g = \int_0^\infty I_{\text{ambient}}(\lambda, \mathbf{x}) \rho_{\text{ambient}}(\lambda, \mathbf{x}) S_G(\lambda) d\lambda$. Note that we assume that the additional ambient reflected intensity is not solely red, as was illustrated in Fig. 1 by r , but rather has two spectral components, red (r) and green (g). The vectors are illustrated in Fig. 2.

In the ideal case the spectral components on either side of a shadow boundary are related by a multiplicative constant (see Table 1). For a general value of (r, g) this will be the case only if (r, g) is in the same direction as (R, G) . A measure of the difference from the ideal case is therefore

$$\text{Amount of pull} = (r, g)(R, G)^\perp, \quad (7)$$

where $(R, G)^\perp$ is orthogonal to (R, G) and therefore $(R, G)^\perp \equiv (-G, R)/|(R, G)|$. In order to measure the *relative* amount of pull, we divide by the length of the original vector. Since the sign of the result is not important, but its magnitude is, we take the absolute value of the result. Thus, when put formally, the pull factor is computed as

$$\text{Pull factor} = \frac{(r, g)(-G, R)}{|(R, G)|^2} = \frac{|gR - rG|}{R^2 + G^2}. \quad (8)$$

This quantity is important because the stronger the effect of the additional ambient illumination, the more the shadow region differs from the lit region, thus making the association of both with the same material more unlikely.

PRACTICAL ESTIMATION OF THE CHANGE IN REFLECTED INTENSITY

To this point, we have examined the factors that contribute to the different behavior of ideal and nonideal shadows. We have also suggested a measurement that reflects the relationship between the additional ambient reflected intensity and the original diffuse reflected intensity. We now turn to apply this information to the problem of identifying material changes through the use of a set of operators that respond differently under the two illumination conditions. In this section we introduce the operators suggested for this task

and examine their responses under the different conditions, along with their correspondence to the theoretical analysis of the reflection vectors described in the previous section.

Let us first consider a set of operators that have spatial center-surround antagonistic organization and share the same spectral intensity functions in both the center and the surround. Examples of these units, which we will refer to as monochromatic opponent units, are $R + \text{center}/R - \text{surround}$, $G + \text{center}/G - \text{surround}$, where R and G represent the responses of the mechanisms bearing the R - and G -cone spectral sensitivity functions. A formulation of the response of one of these units, say, $R + \text{center}/R - \text{surround}$, is given by

$$\begin{aligned} \text{RESP}(\mathbf{x}; \sigma) &= \mathbf{G}(\mathbf{x}; \sigma_c) * L_R(\mathbf{x}) - \mathbf{G}(\mathbf{x}; \sigma_s) * L_R(\mathbf{x}) \\ &= \int_{-\infty}^{\infty} \mathbf{G}(\mathbf{x} - \mathbf{r}; \sigma_c) L_R(\mathbf{r}) d\mathbf{r} \\ &\quad - \int_{-\infty}^{\infty} \mathbf{G}(\mathbf{x} - \mathbf{r}; \sigma_s) L_R(\mathbf{r}) d\mathbf{r}, \end{aligned} \quad (9)$$

where $*$ is the convolution operator and $L_R(\mathbf{x})$ is the logarithm of the response of the red-sensitive mechanism to the input image, that is, $L_R(\mathbf{x}) = \log R(\mathbf{x})$, and $R(\mathbf{x})$ is defined in Eq. (4). We chose the logarithm since it provides operators with a simple behavior to multiplicative scaling of stimuli. It is also interesting to note that, after adaptation effects have been taken into account, the logarithm provides a rough approximation to photoreceptor responses for low-contrast stimuli.^{13,14} $\mathbf{G}(\mathbf{x}; \sigma_i)$ is a Gaussian of the form

$$\mathbf{G}(\mathbf{x}; \sigma_i) = \frac{1}{2\pi\sigma_i^2} \exp\left(\frac{-|\mathbf{x}|^2}{2\sigma_i^2}\right). \quad (10)$$

It is evident that these operators are difference-of-Gaussians (DOG's) in space and therefore are bandpass filters with chromatic sensitivity, each one tuned to a different band of wavelengths. If these operators were input the corresponding spectral component of images (e.g., and $R + \text{center}/R - \text{surround}$ operator with the red component of an image), then they are bound to detect any discontinuity in that component. These discontinuities may be material changes, ideal shadow boundaries, or nonideal shadow boundaries. In terms of the analysis presented in the previous section, the monochromatic opponent units provide information about the *total* change in a given chromatic component. For example, a unit of the type $R + \text{center}/R - \text{surround}$ will respond to the total change in both r and g along a discontinuity, signaling how much the red component of the reflected intensity varied from one side of the discontinuity to the other. Two examples of responses of such units are given in Figs. 3(a)–3(d).

We now consider double-opponent operators of the same organization and behavior described by physiologists as existing in the cortex of primates.^{15,16} These operators, like the monochromatic opponent units, are spectrally sensitive DOG's, although their spectral sensitivity is a combination of two different spectral sensitivity functions. Examples of such operators are $R + G - \text{center}/R - G + \text{surround}$, $B + Y - \text{center}/B - Y + \text{surround}$. The response of such a unit, say,

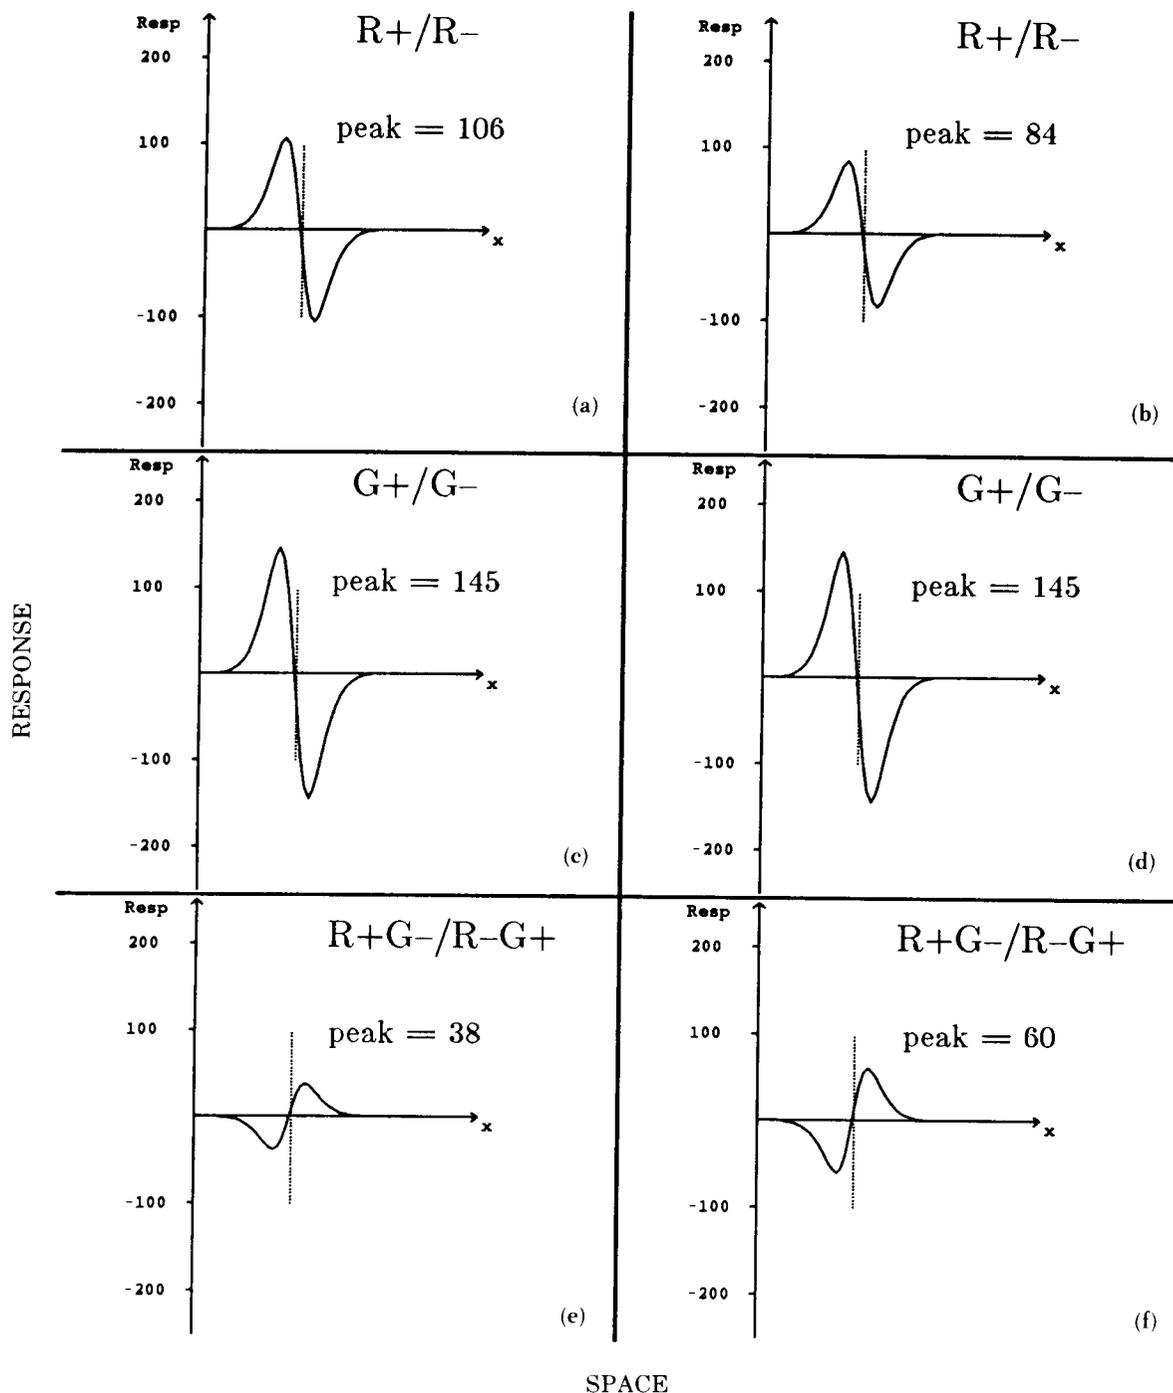


Fig. 3. Responses of monochromatic opponent units and double-opponent operators to two sets of test patterns that include a step-edge in space in which the spectral density changes. The red and green components of the test patterns are (150, 200) followed by (90, 100) for plots (a), (c), and (e) and (90, 100) followed by (60, 50) for plots (b), (d), and (f). The additional ambient reflected intensity was assumed to be 30 units of red, and $\alpha = 0.5$. Plots (a) and (b) are responses of an $R+$ center/ $R-$ surround operator, plots (c) and (d) are of a $G+$ center/ $G-$ surround operator, and plots (e) and (f) are of an $R+G-$ center/ $R-G+$ surround operator. The radius of the center of all operators used is 7 pixels, and σ_s/σ_c is 2.5.

$R + G-$ center/ $R - G+$ surround, is computed as

$$\begin{aligned} \text{RESP}(\mathbf{x}; \sigma) &= \mathbf{G}(\mathbf{x}; \sigma_c) * [L_R(\mathbf{x}) - L_G(\mathbf{x})] \\ &\quad - \mathbf{G}(\mathbf{x}; \sigma_s) * [L_R(\mathbf{x}) - L_G(\mathbf{x})] \\ &= [L_R(\mathbf{x}) - L_G(\mathbf{x})] \text{DOG}(\mathbf{x}). \end{aligned} \tag{11}$$

The difference is that double-opponent operators respond to nonproportional changes in their two spectral components. It has already been shown¹⁷ that if the discontinuity repre-

sents a proportional change, then other operators, such as achromatic DOG's, respond to the discontinuity while the double-opponent operator does not, thus indicating that although there may have been a discontinuity, it is due to an ideal shadow. By contrast, if there are changes due to shadows that are not proportional, that is, nonideal, and as a result, there exists some additional ambient reflected intensity, then the double-opponent operators do respond; this is illustrated in Figs. 3(e) and 3(f).

In summary, the two sets of operators, the monochromatic opponent units and the double-opponent units, signal different aspects of the information needed. The monochromatic opponent units provide information about the total change in each chromatic component, while the double-opponent units provide information about changes in the relative amounts. We anticipate that the stronger the additional ambient reflected intensity becomes (relative to the diffuse reflected intensity), the stronger the response of the double-opponent operator will be. This point is illustrated in Fig. 3, where responses of the units to two test patterns are plotted. Each test pattern is composed of two chromatic stimuli juxtaposed in space (we use in this example only one spatial dimension). In the left-hand test pattern [Figs. 3(a), 3(c), and 3(e)], r is weak relative to R and G , while on the right-hand test pattern [Figs. 3(b), 3(d), and 3(f)], r is strong relative to R and G . Therefore the red-opponent unit responds more strongly to the left-hand test pattern [Figs. 3(a) and 3(b)], while the double-opponent unit responds more strongly on the right [Figs. 3(e) and 3(f)]. Since the double-opponent units are sensitive to two chromatic bands (either red and green or blue and yellow), when comparing the response of a double-opponent unit (say, red and green) with its corresponding monochromatic opponent units, we will have to consider the responses of *two* units, the red ones and the green ones. Thus it is reasonable to compute the ratio between the peak amplitude (absolute value) of the response of a double-opponent unit with the square root of the sum of peak amplitudes of the responses of the corresponding monochromatic opponent units squared. So, for example, in the red-green case, we compute the ratio between the peak response of the $R+G-$ center/ $R-G+$ surround unit and the square root of the sum of peak responses of the $R+$ center/ $R-$ surround and $G+$ center/ $G-$ surround units squared. That is,

$$\text{Relative Amplitude Response} = \frac{|\text{Peak Response of } R+G- / R-G+|}{[(\text{Peak Response of } R+/R-)^2 + (\text{Peak Response of } G+/G-)^2]^{1/2}}. \quad (12)$$

If the ratio is greater than the anticipated pull factor, a threshold measurement computed by a higher-level process, then the discontinuity in the area is not considered a material change but rather a change caused by shadow effects.

What remains to be defined is the extent of the area in which we compute this relative amplitude response. This area has to be small enough so as not to include the responses resulting from other discontinuities and large enough to cover both the peaks of the monochromatic opponent units and the double-opponent ones, which may not occur in the same spatial location. The area in which we compare the amplitude responses is the one covered by the center field of both operators, which are of the same size, since this area essentially excludes (with high probability) peaks that are a result of two adjacent discontinuities,¹⁸ yet it is large enough to guarantee the detection of the peak of the responses.¹⁹ The appropriate size of operators relevant to the image (or parts of it) remains an open problem.

RESULTS

First we verify that the relative amplitude response is monotonically related to the pull factor. We have simulated different cases of object colors and different levels of ambient

illumination and have computed both the relative amplitude responses of the operators and the corresponding pull factors. The relative amplitude responses were computed as the ratio between the highest peaks of the double-opponent units and the highest peaks of the corresponding monochromatic opponent units, all within an area bounded by the radius of the center of the operators. For example, for the red-green case, the ratio was computed as the one between the highest peak of the $R+G$ center/ $R-G+$ surround unit and the sum of the highest peaks of the $R+$ center/ $R-$ surround and $G+$ center/ $G-$ surround units within a round area of radius 7 pixels. The pull factor was computed by using Eq. (8). The results of the comparisons are summarized in Fig. 4.

In the tests that we ran, four different diffuse reflected intensities were picked, and for each one the amount of ambient reflected intensity was varied. For each test, the relative amplitude response and the pull factor were computed. As can be observed from the results, there is a proportional relationship between the relative amplitude responses and the pull factor. This leads us to the conclusion that the pull factor can be estimated by the practical method introduced, which makes use of the relative responses of the monochromatic opponent units and the double-opponent operators. Therefore, by using a threshold on the acceptable pull factor, it is possible to hypothesize whether the discontinuity is a shadow boundary or a material change. For larger values of this threshold, more discontinuities will be accepted as shadow boundaries.

We applied this technique to the problem of identifying material changes, as is shown in Fig. 5. The size of the original image is 256×256 pixels, the radius of the center for both the monochromatic opponent units and the double-opponent units is 7 pixels, and the ratio between the stan-

dard deviations of the surround to the center (σ_s/σ_c) is 2.5. The original image [Fig. 5(a)] has three peppers, one of which (1) is red, which casts its reflection on a green pepper (2). The image was convolved with an $R+G-$ center/ $R-G+$ surround operator to yield possible material changes [Fig. 5(b)—the white boundaries]. It is evident that some of the shadows in the green pepper (2) are ones that are affected by the red reflection of the red pepper, and therefore we tested the image for possible changes of classification in order to verify it. Figure 5(c) is a result of the reclassification with a predicted pull factor of 0.2, showing that some of the possible material changes are probably just shadows (black boundaries); Fig. 5(d) shows that increasing the pull factor to 0.3 changes the classification of more of the material changes to possible shadows. What we see, then, is that if some global process is able to predict the pull factor, then there is a way of correcting the classification even under complex illumination conditions.

SUMMARY

The relationship between shadow and lit regions in images can be complicated by the existence of spectrally biased ambient illumination that is reflected off objects in the environment. In contrast to previous approaches for inferring

Comparison Results				
Test Number	Diffuse Intensity (R, G)	Ambient Intensity (r, g)	Pull Factor	Relative Amplitude Response
1	(110,160)	(20,0)	0.08	0.13
	(110,160)	(30,0)	0.13	0.17
	(110,160)	(40,0)	0.17	0.20
	(110,160)	(50,0)	0.21	0.23
	(110,160)	(60,0)	0.25	0.26
	(110,160)	(100,0)	0.42	0.33
2	(50,60)	(10,0)	0.10	0.12
	(50,60)	(20,0)	0.20	0.20
	(50,60)	(30,0)	0.29	0.26
	(50,60)	(40,0)	0.39	0.30
	(50,60)	(50,0)	0.49	0.32
	(50,60)	(60,0)	0.59	0.34
3	(100,50)	(15,05)	0.02	0.05
	(100,50)	(30,10)	0.04	0.07
	(100,50)	(45,15)	0.06	0.08
	(100,50)	(60,20)	0.08	0.09
4	(150,100)	(10,20)	0.06	0.08
	(150,100)	(15,30)	0.09	0.11
	(150,100)	(20,40)	0.12	0.13
	(150,100)	(25,50)	0.15	0.15
	(150,100)	(30,60)	0.18	0.19

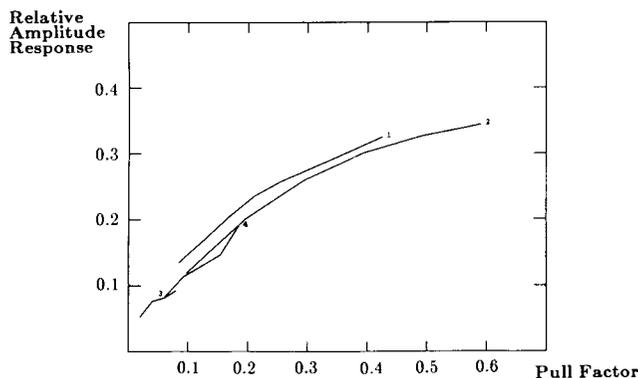


Fig. 4. Comparisons between the pull factor and the relative amplitude responses of the units (red and green opponent units and $R+G$ -center/ $R-G$ +surround operator).

material changes, which assumed that the ambient illumination has the same spectral characteristics as the incident illumination, we have presented a technique that takes into account ambient illumination that is different in its distribution. A measurement that is based on a biologically motivated set of operators has been introduced and found to be related directly to the relative strength of the ambient illumination. It is important to remember that this is an early visual process and as such will not determine perfectly, nor decisively, whether the discontinuities are indeed material changes. Some other process, a higher-level one that involves knowledge about the scene and its structure, may perform the final decision making regarding this classification. Furthermore, it is clear that the threshold that we introduced has to be provided to our technique, since it may vary from one image to another or even from one region to another in the same image. Therefore an additional computation must be carried out by higher-level spatiochroma-

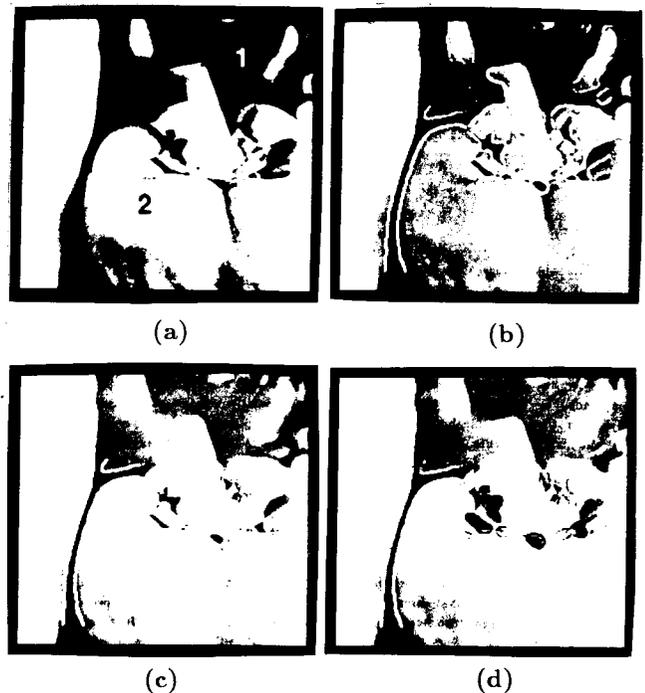


Fig. 5. An application of the algorithm to a chromatic image. (a) Is the original image, with a red pepper (1) casting its color on a green pepper (2). (b) Is the result of the algorithm run with predicted pull factor of zero; the bright lines are the discontinuities hypothesized as material changes. (c) Is the result of the algorithm with predicted pull factor of 0.2, and (d) with a pull factor of 0.3; the darker lines are discontinuities that were reclassified as shadow boundaries. The images in (b), (c), and (d) are shown with reduced contrast so that the discontinuities will be easier to see.

tic mechanisms, ones that should take into account information such as spatial organization and reflectance properties of objects. Nevertheless, the technique that we presented is definitely more flexible, accurate, and robust in its performance.

ACKNOWLEDGMENTS

This work was financially supported in part by the Natural Science and Engineering Research Council of Canada and by the Department of Computer Science at the University of Toronto.

John K. Tsotsos is also a Fellow of the Canadian Institute for Advanced Research.

REFERENCES

1. R. Nevatia, "A color edge detector and its use in scene segmentation," *IEEE Trans. Syst. Man Cybern.* **SMC-7**, 820-826 (1977).
2. G. B. Coleman and H. C. Andrews, "Imaging segmentation by clustering," *Proc. IEEE* **67**, 773-785 (1979).
3. P. A. Nagin, A. R. Hanson, and E. M. Riseman, "Studies in global and local histogram-guided relaxation algorithms," *IEEE Trans. Pattern Anal. Machine Intell.* **PAMI-4**, 263-276 (1982).
4. B. J. Schacter, L. S. Davis, and A. Rosenfeld, "Scene segmentation by cluster detection in color spaces," *SIGART Newsl.* **58**, 16-17 (1976).
5. R. Ohlander, K. E. Price, and D. R. Reddy, "Picture segmentation using a recursive region splitting method," *Comput. Graphics Image Process.* **8**, 313-333 (1978).
6. M. D. Levine and S. I. Shaheen, "A modular computer vision

- system for picture segmentation and interpretation, Part 1," in *Proceedings of IEEE Conference on Pattern Recognition and Image Processing* (Institute of Electrical and Electronics Engineers, New York, 1979), pp. 523-533.
7. M. Nagao, T. Matsuyama, and Y. Ikeda, "Region extraction and shape analysis in aerial photographs," *Comput. Graphics Image Process.* **10**, 195-223 (1979).
 8. K. E. Sloan, "World model driven recognition of natural scenes," Ph.D. dissertation (University of Pennsylvania, Philadelphia, Pa., 1977).
 9. J. M. Rubin and W. A. Richards, "Color vision and image intensities: when are changes material?" *Biol. Cybern.* **45**, 215-226 (1982).
 10. J. M. Rubin and W. A. Richards, "Color vision: representing material changes," AI Memo 764 (MIT Artificial Intelligence Laboratory, Cambridge, Mass., 1984).
 11. S. A. Shafer, "Using color to separate reflection components," Technical Rep. TR-136 (Department of Computer Science, University of Rochester, Rochester, N.Y., 1984).
 12. R. L. Cook and K. E. Torrance, "A reflectance model for computer graphics," *Comput. Graphics* **15**, 307-316 (1981).
 13. R. M. Boynton and D. N. Whitten, "Visual adaptation in monkey cones: recordings of late receptor potentials," *Science* **170**, 1423-1426 (1970).
 14. W. A. H. Rushton, "Peripheral coding in the nervous system," in *Sensory Communication*, W. A. Rosenblith, ed. (MIT Press, Cambridge, Mass., 1961).
 15. M. S. Livingstone and D. H. Hubel, "Anatomy and physiology of a color system in the primate visual cortex," *J. Neurosci.* **4**, 309-356 (1984).
 16. C. R. Michael, "Color vision mechanisms in monkey striate cortex: dual-opponent cells with concentric receptive fields," *J. Neurophys.* **41**, 557-576 (1978).
 17. R. Gershon, "Empirical results with a model of color vision," in *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition* (Institute of Electrical and Electronics Engineers, New York, 1985), pp. 302-305.
 18. M. S. Longuet-Higgins, "The distribution of intervals between zeros of a stationary random function," *Phil. Trans. R. Soc. London Ser. A* **254**, 557-599 (1962).
 19. D. J. Fleet, A. D. Jepson, and P. E. Hallett, "A spatio-temporal model for early visual processing," Technical Rep. RBCV-TR-84-1 (Department of Computer Science, University of Toronto, Toronto, 1984).