

The Synthesis and Analysis of Color Images

BRIAN A. WANDELL

Abstract—I describe a method for performing the synthesis and analysis of digital color images. The method is based on two principles. First, image data are represented with respect to the separate physical factors, surface reflectance and the spectral power distribution of the ambient light, that give rise to the perceived color of an object. Second, the encoding is made efficient by using a basis expansion for the surface spectral reflectance and spectral power distribution of the ambient light that takes advantage of the high degree of correlation across the visible wavelengths normally found in such functions.

Within this framework, the same basic methods can be used to synthesize image data for color display monitors and printed materials, and to analyze image data into estimates of the spectral power distribution and surface spectral reflectances. The method can be applied to a variety of tasks. Examples of applications include the color balancing of color images and the identification of material surface spectral reflectance when the lighting cannot be completely controlled.

Index Terms—Color balancing, color constancy, color correction, efficient image encoding, image analysis, image synthesis, subspace methods.

LIST OF SYMBOLS

SYMBOL	SYMBOL DEFINITION
N	Number of wavelength samples.
λ_n	Sample wavelength value.
$E^x(\lambda_n)$	Ambient light spectral power distribution at x .
$E_i(\lambda_n)$	Basis function of ambient light representation.
e_i	Scalar weight of i th ambient light basis function.
ϵ	Vector of ambient light weights.
Ω_E	Ambient lighting matrix with respect to wavelength basis.
Λ_E	Ambient lighting matrix with respect to arbitrary basis functions.
$S^x(\lambda_n)$	Surface reflectance at x .
$S_j(\lambda_n)$	Basis function of surface reflectance.
σ_j^x	Scalar weight of i th surface basis function at x .
σ^x	Vector of surface weights at x .
$C^x(\lambda_n)$	Color signal at location x , $C^x(\lambda_n) = E^x(\lambda_n) S^x(\lambda_n)$.
$R_k(\lambda_n)$	Spectral sensitivity of k th sensor.

ρ_k^x	Scalar response of k th sensor at x .
ρ^x	Vector of sensor responses at x .
π	Best perpendicular to the sensor response vectors.
$D(E)$	Dimensionality of the ambient light representation.
$D(S)$	Dimensionality of the surface representation.

INTRODUCTION

THE two primary physical factors that influence the color appearance of objects in a still image are the spectral power distribution of the ambient light and the surface spectral reflectance of the objects in the image. The information one wishes to display in a synthetic image, and the information available to analyze in during image acquisition, is called the *color signal*. I use the term color signal to refer to the spectral power distribution of the light arriving at the recording device or the human eye. In a natural scene, the color signal is generally computed as the product of the spectral power distribution of the light incident on an object and the surface reflectance of that object. In this paper, I will report on a common representational framework for the synthesis and analysis of color images based upon the synthesis and analysis of the color signal.

In the *synthesis* of color images, we begin with specifications of the surface spectral reflectance function and ambient light spectral power distribution. From these, we calculate the expected color signal. We then seek the intensities of the color display guns such that the display will have the same effect on the human visual system as the actual color signal would have. The application of the synthesis methods is most significant for the problem of accurate rendering of color images.

In the *analysis* of color images, the sensor responses in the acquired image are determined by the color signal. We then seek a method of inverting the process, that is, recovering the underlying surface spectral reflectance and ambient light spectral power distribution functions that gave rise to the color signal. The methods for analysis of color images are useful both for the correct rendering of color images and for the more general vision problem of identifying surface properties of objects from image data. I believe that this method can play an important role in the pattern analysis of color images.

The central issue that unifies the synthesis and analysis of color images is a specification of how the surface and lighting *information* is encoded in the color signal. I characterize the information in the surface and light functions

Manuscript received August 19, 1985; revised February 21, 1986. Recommended for acceptance by S. W. Zucker. This work was supported by NEI Grant 2 R01 EY03164 and NASA Grants NCC 2-307 and 2-332. This work was performed while the author held an NRC Post-Doctoral Fellowship at the NASA Ames Research Center.

The author is with the Department of Psychology, Stanford University, Stanford, CA 94305.

IEEE Log Number 8609223.

by the number of degrees of freedom in a linear model that is used to approximate these functions. In image synthesis, significant efficiencies in image representation and computation can be obtained by making the number of independent degrees of freedom explicit. In image analysis, the number of independent degrees of freedom in the surface and light functions is the key factor that proscribes how many different classes of color sensors are necessary to decompose the color signal and estimate the surface and light functions. The principal result—due to Maloney and Wandell [12]—is that if the surface reflectance and light functions can be described using N degrees of freedom, then one requires $N + 1$ sensors to unlock the color signal and estimate the surface and light functions. This result has practical significance both for the accurate pattern analysis of color images and for the accurate rendering of color images.

OVERVIEW

In the next section, I will develop the basic notation and matrix operations common to the synthesis and analysis of color images. I will then describe the procedures for color synthesis, showing how we may relate the surface spectral reflectance functions, ambient light spectral power distribution, human photoreceptor responses, and display intensities on a monitor. In the third section, I will treat the analysis of color images. I will describe a version of Maloney and Wandell's [12], [20], [11] method for recovering the spectral power distribution of the ambient light and surface spectral reflectance functions of the objects from the sensor responses to the color signal.

The theoretical material in the next two sections is presented using general notation. The basic results apply to a wide class of representations of the ambient light and surface reflectance functions. In the final section, I will turn to the problem of selecting specific representations of these functions that are efficient with respect to a wide class of naturally occurring objects and lights.

PRELIMINARY CONSIDERATIONS

The Asymmetry Between Light and Surface

Suppose that we acquire information at sample wavelengths λ_n for $n = 1, N$. The color signal is obtained by multiplying the spectral power distribution of the ambient light at x , $E^x(\lambda_n)$ times the spectral reflectance function of the surface spectral reflectance at x , $S^x(\lambda_n)$. We denote the color signal at point x as $C^x(\lambda_n) = E^x(\lambda_n) S^x(\lambda_n)$.

In developing the mathematical formulation for the synthesis of color images, we can retain symmetry between the role played by the surface and light functions without penalty. This symmetry is evident in the definition of the color signal $C^x(\lambda_n) = E^x(\lambda_n) S^x(\lambda_n)$, and can be maintained throughout (cf. Sallstrom [14] and Buchsbaum [2]). In most images, however, there is a physical asymmetry between the surface spectral reflectance and the ambient light spectral power distribution (cf. Horn [8] and Oppenheim and Schaffer [13]). The asymmetry arises since the

rate of spatial variation of the ambient light spectral power distribution $E^x(\lambda_n)$ is generally lower than the rate of spatial variation in the surface spectral reflectance function $S^x(\lambda_n)$. The slow spatial variation of the spectral power distribution makes it possible to represent the ambient light at lower spatial resolution than the surface functions, permitting some additional computational savings.

The most significant consequence of the difference in the rate of spatial variation, however, occurs for the analysis of color images. The uniqueness results for the recovery of the surface and light spectral power distributions from the color signal indicate that we cannot uniquely recover the light and surface functions if both vary freely over space. Essentially unique recovery is possible only if one of these functions varies more slowly over space than the other. We adopt the assumption that spatial variation of the ambient light is slower than the spatial variation of the surface reflectance function (cf. Horn [8] and Oppenheim and Schaffer [13]).

To evaluate the plausibility of this idea, the reader might notice that in natural scenes in which the ambient light varies more rapidly spatially than the underlying surface—for example, when viewing a slide projected on a screen—the human visual system interprets the spatial variation as if it were due to a surface variation, and not from its true cause, the ambient light variation. The analysis we will use here also assumes that there is an asymmetry between light and surface functions. In any local region of the image, we assume that there is only spatial variation in the surface reflectance function, and no spatial variation in the ambient light spectral power distribution. We will indicate this assumption in our notation by suppressing the superscript x when we refer to the parameters of the spectral power distribution of the ambient light. Thus, we will write $E(\lambda_n)$ instead of $E^x(\lambda_n)$, so that in a small region of the image, $C^x(\lambda_n) = E(\lambda_n) S^x(\lambda_n)$.

I emphasize that this assumption does not imply that the ambient light remains constant over the entire image, but only that the spectral power distribution remains constant over a local region of the image in which there is significant spatial variation in the surface spectral reflectance function.

SYNTHESIS OF COLOR IMAGES

Introduction

In generating color images, we wish to determine the appropriate display gun intensities so that the display will have the same *visual effect* as the actual color signal $C^x(\lambda_n)$. In their important work on color standards, the members of the CIE have spent more than 50 years describing various methods and representations for achieving this result. The various methods and representations (using XYZ values, YIQ values, etc.) are logically equivalent, but have sometimes been adopted because of odd technical limitations in the computation. For purposes of the analysis here, I prefer to use a calculation advocated

by Cornsweet [6] based on the photoreceptor absorptions of the human eye. This is the method most closely based on the physical factors governing color perception. Although the method is logically equivalent to the several methods suggested by the CIE and has an excellent physical interpretation, no standard photoreceptor representation has been adopted by the CIE. The estimates by Smith and Pokorny [15] and described by Boynton [1] are widely used at present.

Basic Equations: Wavelength Representation

To determine the display gun intensities that will appear the same as the color signal to the human eye, we first compute the effect that the color signal $C^x(\lambda_n)$ would have on the three types of photoreceptors in the human eye. Suppose that the sensitivity of the receptors in the eye is described by the three sensitivity functions $R_k(\lambda_n)$ for $k = 1, 3$. If we have measurements of the surface reflectance function and spectral power distribution function at N wavelength values for λ_n , $n = 1, N$, then the strength of the k th receptor response at location x is

$$\rho_k^x = \sum_{n=1}^N C^x(\lambda_n) R_k(\lambda_n). \quad (1)$$

If one begins with the specification of the color signal in terms of the surface and light functions, and if there are N sample points along the wavelength dimension, then the calculation requires $2N$ multiplications and $N - 1$ additions per pixel. To compute the three receptor responses requires $6N$ multiplications and $3(N - 1)$ additions.

The effect of setting the intensity values of the three color guns on the receptors can be calculated as follows. Suppose that the red, green, and blue display guns have spectral emission functions $G_j(\lambda_n)$ for $j = 1, 3$. Denote the intensity of the j th gun as I_j . When the display gun intensities at pixel x are set to the vector of value $I^x = (I_1^x, I_2^x, I_3^x)$, the response of the k th class of photoreceptors will be

$$\rho_k^x = \sum_{j=1}^3 I_j^x \sum_{n=1}^N G_j(\lambda_n) R_k(\lambda_n). \quad (2)$$

Equation (2) can be written as a matrix equation of the form

$$\rho^x = \Gamma I^x \quad (3)$$

where the k, j th entry of the matrix Γ is

$$\sum_{n=1}^N G_j(\lambda_n) R_k(\lambda_n). \quad (4)$$

The gun intensities that are visually equivalent to the color signal may be determined by first solving (1) for the values of ρ^x and then solving (3) for I^x . Two issues arise immediately in computing the display gun intensities.

First, the solution may have negative values. This will occur when the required color image is outside the gamut of colors that can be displayed by the particular device. I will discuss a method of treating this problem later in the paper when the theory has been more fully developed.

Second, there is the issue of computational efficiency. The entries of the matrix Γ are fixed by the display gun spectral emission functions and the human photoreceptor sensitivity curves. This matrix (or its inverse which generally plays a bigger role) need only be calculated once, and therefore does not represent a significant computational burden. The calculation of the receptor responses ρ^x to a color signal $C^x(\lambda_n)$ must be made for every point in the image. This is the most expensive part of the computation in the synthesis of color images. The cost of this computation is governed by the number of sample points along the wavelength dimension for the functions $E^x(\lambda_n)$ and $S^x(\lambda_n)$. Since the amount of computation is proportional to the sampling rate N , one method of limiting the computation is to reduce the sampling rate N .

A second option is to transform the representation of the surface spectral reflectance and ambient light spectral power distribution functions from a wavelength representation to a more efficient representation. Great savings in efficiency can be obtained because the wavelength functions describing naturally occurring lights and surfaces are highly correlated at nearby wavelengths. To improve computational efficiency, we must seek representations of the spectral power distributions of the lights and spectral reflectance functions of the surfaces that are decorrelated. We will consider specific representations later in this paper. The general effect of applying such transformations is analyzed in the following section.

Basic Equations: Transformed Representation

In order to efficiently encode ambient light and surface vectors, we will seek alternative representations, related to the wavelength representation by a linear transformation. These take the form

$$E(\lambda_n) = \sum_{i=1}^{D(E)} \epsilon_i E_i(\lambda_n) \quad (5)$$

and

$$S^x(\lambda_n) = \sum_{j=1}^{D(S)} \sigma_j^x S_j(\lambda_n) \quad (6)$$

where $D(E)$ is the chosen dimensionality of the light representation and $D(S)$ is the dimensionality of the surface representation. In order to span any particular subspace, it is sensible to select functions $E_i(\lambda_n)$ or $S_j(\lambda_n)$ that are orthonormal in the sense that

$$\sum_{n=1}^N E_i(\lambda_n) E_j(\lambda_n) = 0 \quad (7)$$

when $i \neq j$ and the length of each function is one. This is not necessary—the entire development remains consistent for any choice of linearly independent functions. However, for any subspace one wishes to span, it is always possible to find such a set of functions so that choosing orthonormal functions does not restrict the analysis, and by choosing an orthonormal basis, the recovery of the weights ϵ_i and σ_j^x from the wavelength function is very easy. We refer to these sets of functions $E_i(\lambda_n)$ and $S_j(\lambda_n)$ as *basis functions*.

The usual wavelength representation is equivalent to choosing a vector of weights with respect to the set of functions $E_i(\lambda_n)$ such that

$$E_i(\lambda_n) = 1 \quad \text{if } i = n \text{ and } 0 \text{ otherwise} \quad (8)$$

and $S_j(\lambda_n)$ such that

$$S_j(\lambda_n) = 1 \quad \text{if } j = n \text{ and } 0 \text{ otherwise.} \quad (9)$$

We can express the relationship between the vector of wavelength sample measurements that characterize the ambient light $E(\lambda_n)$ and the column vector of weights that characterize the transformed representation of the ambient light $\epsilon = (\epsilon_1, \epsilon_2, \dots)$ as

$$E(\lambda_n) = E\epsilon. \quad (10)$$

The columns of the matrix E are equal to the basis vectors $E_i(\lambda_n)$. Similarly, we can express the relationship between the vector of sample measurements of the surface reflectance that characterizes the surface with respect to wavelength and the vector of weights that characterize the surface with respect to the transformed representation $\sigma^x = (\sigma_1^x, \sigma_2^x, \dots)$ as

$$S^x(\lambda_n) = S\sigma^x. \quad (11)$$

The columns of the matrix S are the basis vectors $S_j(\lambda_n)$. To compute the matrix required to transform from one representation to another, one may use the linear equations (10) and (11) to compute the vector in either wavelength or the transformed domain. When the vectors in the two domains are linearly independent and of the same dimension, a transformation is always possible without loss of information. To improve computational efficiency, we would like to choose the dimensionality in the transform domain to be smaller than the dimensionality in the original representation ($D(E) < N$ and $D(S) < N$), but with little loss of precision. When we choose a smaller dimensionality in the transformed domain than in the wavelength domain, the equations may have no exact solution. In that case, we will always choose *the least squares solution* to (10) and (11) as our representation. When the bases are orthonormal, the least squares fit is found by multiplying by the matrix transpose S' or E' .

Relationship Among Physical Quantities

In order to compute the response of a receptor ρ_k^x from the surface and light functions, we use the equation¹

¹In general, the effective surface reflectance at a point in the image may depend upon the geometry of the image. For example, specular surface reflections depend strongly upon the angles among the lighting source, the surface, and the viewer. Although we do not specifically indicate such dependencies, we assume that the surface reflectance function $S^x(\lambda_n)$ will generally depend upon the viewing geometry. Since our analysis here is for a single fixed image, we simply assume that the surface spectral reflectance has been properly specified for the particular viewing geometry. An analysis of the effects of specularity at the physical level may be found in Torrance and Sparrow [18] and Torrance *et al.* [19]. An analysis of the effects of incorporating more sophisticated specular effects in computer-generated images may be found in Cook and Torrance [5]. I will consider the issue of specularity again in the section on the selection of basis functions.

$$\rho_k^x = \sum_{n=1}^N E(\lambda_n) S^x(\lambda_n) R_k(\lambda_n). \quad (12)$$

Using matrix notation, we can express the relationship among the sensors, the ambient light, and the surface reflectance function in the form

$$\rho^x = \Omega_E \omega^x \quad (13)$$

where ω^x is the column vector $(S^x(\lambda_1), S^x(\lambda_2) \dots)'$. When the surface is represented with respect to the wavelength basis terms $S_j(\lambda_n) = 1$ if $i = n$ and 0 otherwise, then the kj th entry of the matrix Ω_E is

$$E(\lambda_j) R_k(\lambda_j). \quad (14)$$

I reserve the use of the Greek symbols for omega (Ω and ω) to refer to the special case in which the physical entities are represented with respect to the wavelength basis functions.

The relationship among the surface, the light, and the sensor responses may be characterized as a simple matrix product. The surface vector is mapped by a linear transformation into the vector of sensor responses. The linear transformation is determined by two factors: the ambient light $E(\lambda_n)$ and the receptor sensitivities $R_k(\lambda)$. Across viewing conditions for a particular camera or for the human eye, the receptor sensitivities are unchanging and known. Thus, the only unknown factor that governs the linear transformation Ω_E is the ambient light spectral power distribution. For that reason, we make E explicit when we write down the matrix representing the linear transformation. We will call matrices such as Ω_E *lighting matrices*, although the reader should note that the spectral power distribution of the ambient light and the spectral sensitivities of the sensors both play a role in forming the effective lighting matrix for any particular device.

The relationship among the surface, lights, and sensors takes on a parallel form when we transform the surface representation to a new set of basis vectors. If we use

$$S^x(\lambda_n) = \sum_{j=1}^N \sigma_j^x S_j(\lambda_n), \quad (15)$$

then the matrix equation becomes

$$\rho^x = \Omega_E S \sigma^x = \Lambda_E \sigma^x \quad (16)$$

where the columns of the matrix S are simply the N vectors $S_j(\lambda_n)$. The matrix S is fixed and we have incorporated the change of basis of the surface representation into the lighting matrix to define a new lighting matrix $\Lambda_E = \Omega_E S$. The kj th entry of Λ_E is now

$$\sum_{n=1}^N E(\lambda_n) S_j(\lambda_n) R_k(\lambda_n). \quad (17)$$

This formula generalizes the formula in (14) for which the surface representation is $S_j(\lambda_n) = 1$ if $j = n$ and 0 otherwise. We will use expression (17) to define the entries of the lighting matrix in general.

The matrix Λ_E can be decomposed into a set of fixed primitive matrices that depend only upon the choice of

basis functions. If we express the spectral power distribution of the ambient light with respect to the basis vectors $E_i(\lambda_n)$ so that the lighting vector is ϵ , then decomposition of Λ_E is

$$\Lambda_E = \sum_{i=1}^{D(E)} \epsilon_i \Lambda_i \quad (18)$$

where the kj th entry of Λ_i is

$$\sum_{n=1}^N E_i(\lambda_n) S_j(\lambda_n) R_k(\lambda_n). \quad (19)$$

This emphasizes the fact that the effect of the light may be described as the sum of the action of several different linear transformations, one linear transformation per degree of freedom in the lighting model.

Summary of Results

The set of matrix operations that capture the relationship between these various quantities may be summarized as follows. Using the rule that the effect of the display on the receptor responses should equal the effect of the surfaces and light on the receptors, we may write

$$\Gamma I^x = \Lambda_E \sigma^x. \quad (20)$$

Inverting Γ ,

$$I^x = \Gamma^{-1} \Lambda_E \sigma^x. \quad (21)$$

If we wish to make the role of the ambient light parameters explicit, we may also write this equation using the expansion of the lighting matrix from (18):

$$I^x = \left[\sum_{n=1}^{D(E)} \epsilon_n \Gamma^{-1} \Lambda_n \right] \sigma^x. \quad (22)$$

Equation (21) expresses the complete relationship among the surface spectral reflectances, light spectral power distribution, and display gun intensities. Once the basis functions are made explicit, and the vector characterizing the lighting ϵ and vectors characterizing the surfaces σ^x are provided with respect to that basis, then all the terms on the right-hand side of the equation are known so that the gun intensities on the left-hand side may be computed.

We may assess the computational burden as follows. In the typical color display, we can control the intensities of three color guns; thus, the vector I^x is three dimensional. The matrix Γ^{-1} is 3 by 3 and its values are fixed by the display phosphors and human photopigment sensitivities. The selection of the basis functions to represent the surfaces and lights completely fixes the matrices Λ_n (19). These matrices are 3 by $D(S)$, so that the product of $\Gamma^{-1} \Lambda_n$ is 3 by $D(S)$ and may be precomputed. They do not represent a significant computational burden.

Thus, we see that the amount of computation depends primarily on the dimensionality of the surface representation $D(S)$. This number plays the same role as the number of sample points on the wavelength dimension N when the equations are written with respect to the wavelength

representation. Significant computational savings can be achieved when the representation with respect to an alternative set of basis functions permits us to use a value of $D(S)$ that is significantly smaller than the value of N .

Display Gamut Limitations: When the display monitor cannot generate the desired signal, one of the required gun intensities I^x will be negative. In general, such an outcome cannot be entirely avoided since the span of colors obtainable by additive light mixtures is certainly less than the span of colors obtainable choosing arbitrary surface and light functions. In general, the range of display intensities that may be required is determined by the range spanned by the columns of the matrix

$$\Gamma^{-1} \Lambda_E. \quad (23)$$

Once the surface and lighting basis functions are chosen, the only freedom left in determining the span of the columns of this matrix is the weights of the ambient light spectral power distribution ϵ_i .

By suitable choice of the light and surface basis functions, however, the number of instances in which these gamut errors occur can be greatly limited. For example, Buchsbaum and Gottschalk [3] have considered the spectra of lights that are smoothly varying in the sense that the Fourier transform of these spectra—when treated as functions of wavelength—only have power in the low-frequency region.² If these functions are restricted to be bandlimited (i.e., have no power above some frequency of modulation), then the corresponding color signal will be bandlimited and generally fall within a restricted color gamut.

In representing objects outside the color gamut of the display, one may choose two methods. First, one may simply adjust the individual point outside the gamut, for example, by setting the negative value to zero. This method distorts the color representation of that point alone. This solution may be adequate, depending upon the significance of the color of that particular portion of the image.

An alternative method of solution is to adjust all of the display values together, attempting to move all of the colors to within the display gamut. This causes some error at each point in the display, but preserves the relationship among the different objects. A reasonable method of making this type of adjustment is to adjust the value of the lighting vector ϵ_i in order to avoid negative display intensities.

The advantage of this type of correction derives from the psychophysical observation that human observers are reasonably good at discounting the appearance of the ambient light when judging the surface appearance of objects (cf. Helson [7]). This method of color correction in the synthesis of images will be evaluated more fully elsewhere.

²Note: Frequency in this case refers to the rate of variation of the light and surface functions with respect to the wavelength of the ambient light, a rather unfortunate (multiple) use of the frequency concept.

The Display of Color Camera Data

When the synthesis of the color image on the screen depends upon data acquired from a color camera rather than data generated by an internal model, the relationship between the camera responses and the desired display values may be rather complex. The difficulty arises when the camera sensors are not within a linear transformation of the human receptor sensitivities. In this case (cf. Horn [9]), no simple calculation relating the camera responses to the desired display values is possible.

The methods described in the next section—on image analysis—may be used to decide upon the correct color display values. The idea is the following. Even if the camera sensors are not within a linear transformation of the human photoreceptor responses, their response can still be used to estimate the surface spectral reflectances and light spectral power distributions in the image. These estimated values can then be used to determine the display output. I describe the procedure for using camera data to estimate the surface and the light parameters in the next section. I will return to the problem of correcting camera output at the end of the paper.

THE ANALYSIS OF COLOR IMAGES

In this section, I will treat the analysis of image data. Beginning with the sensor responses acquired in a camera ρ^x , we ask what we can learn about the surface reflectance functions and the ambient light spectral power distribution. The development in this section is based on work done by Maloney and me [11], [20], [12].

The key equation that relates the response of the sensors in a camera system to the ambient light spectral power distribution and surface reflectance functions is (16) which we now write as

$$\rho^x = \Lambda_E \sigma^x. \quad (24)$$

The known values in the acquired image are on the left-hand side of the equation, and the unknown values are on the right-hand side of the equation. An extensive mathematical analysis of the recovery procedure is contained in Maloney's [11] dissertation. The main results are summarized in Maloney and Wandell [12]. I now describe the computational method used in the Stanford color analysis package, written by Brainard and me.³

Conditions for Recovery: Spatial Variation

We begin by considering what can be recovered from the color signal when the surface and light functions are both free to vary over space. Recall that the response of the sensors ρ^x is completely determined by the color signal, which is given by $C^x(\lambda_n) = E^x(\lambda_n) S^x(\lambda_n)$. To evaluate the uniqueness of recovery, note that we can introduce an arbitrary function of both space and wavelength $G^x(\lambda_n)$ into the equation for the color signal so that

$$C^x(\lambda_n) = E^x(\lambda_n) S^x(\lambda_n) \quad (25)$$

and

$$C^x(\lambda_n) = (E^x(\lambda_n) G^x(\lambda_n)) \left(\frac{S^x(\lambda_n)}{G^x(\lambda_n)} \right). \quad (26)$$

The uniqueness of the recovery can be characterized, then, by describing the limitations on the function $G^x(\lambda_n)$. Clearly, without some limitation on the structure of the surface and light functions, unique recovery is impossible. In order to obtain unique recovery, we must introduce restrictions on the variation of the light and surface functions. The consequences of the restrictions I describe here are thoroughly analyzed by Maloney [11].

A Condition on Spatial Variation: Suppose that the scalar intensity of both the ambient light and surface reflectances are free to vary at each point. If we scale the ambient light at x by a factor $g(x)$ and we scale the surface reflectivity by a factor $1/g(x)$, then the color signal remains unchanged. This emphasizes that we cannot discriminate among the alternative decompositions of the signal into $E^x(\lambda_n)$ and $S^x(\lambda_n)$ versus $E^x(\lambda_n) g(x)$ and $S^x(\lambda_n)/g(x)$.

Equations (25) and (26) demonstrate that the identical color signal can arise from many different interpretations of the functions describing the ambient light and the surface reflectances. To permit a unique recovery of the light and surface functions, we are forced to adopt some assumption about the spatial variation of the light and surface functions. The general view we will adopt—to be made precise below—is that the rate of variation over space of the surface reflectance is greater than the rate of variation over space of the ambient light.

The basic premise of the analysis is illustrated in Fig. 1. We analyze the image sequentially by dividing it into overlapping spatial areas. Within each of the areas, we analyze the image under the assumption that the ambient light in that small region is constant. We choose the image areas with two considerations in mind. First, within each area, there must be significant variation in the sensor responses. As we shall see, this is necessary in order to obtain estimates of the light and surfaces in that area. Second, we choose the areas with a high degree of overlap. This is not necessary, but it is useful so that we can detect slow variations of the spectral power distribution of the ambient light. The complete estimate of the ambient light function is determined by combining the separate estimates of the ambient light into a smoother, interpolated function. The necessary variation in the spectral reflectance function of the surface reflectance functions required within each of the areas is specified in the next section in which I describe the recovery method.

In each local region, then, the spatial variation in the sensor response is assumed to be due entirely to the surface reflectance, and not at all to spatial variation in the ambient light. With this assumption, the estimate of the spectral power distribution of the ambient light is unique

³The manual for the package can be obtained by writing the author. The package is written in C on a UNIX BSD 4.2 system.




Fig. 1. The analysis of the image data proceeds by selecting regions of the image with enough complexity in order to form independent estimates of the surface and the light. Within each region indicated in the figure, there is enough data to permit an estimate of the surfaces and light functions within that region. As we move from region to region, the light and surface functions estimates are forced to be consistent at the points of overlap.

up to a single scalar $\alpha E(\lambda_n)$ and $(1/\alpha) S^x(\lambda_n)$ with weak conditions. The value of α is fixed for all points within each region of the analysis. (The final estimate of the ambient light will be a smoothed version of this piecewise constant function.) Further, since the regions may be chosen to overlap, the requirement of consistency at points of overlap leaves us with only a single unknown scalar for the entire image. By convention, we choose α to normalize the total power in the ambient light to unity.

A Condition on the Number of Color Sensors: Let us now consider the second important condition that is necessary for recovery. This is a condition on the number of color sensors necessary for recovery.

Suppose that the dimensionality of the surface representation $D(S)$ is equal to the number of sensor classes, that is, the dimensionality of ρ^x . In that case, again, no unique solution of the equations across space is possible. To see this, let the reader select any vector whatsoever to represent the weights of the vector of the ambient light E . This choice specifies the lighting matrix Λ_E . With this lighting vector chosen and the data ρ^x known, we may invert the lighting matrix and solve for the surface vectors σ^x . This solution provides a completely acceptable account of the image data. For each arbitrary choice of the lighting parameters, we can generate a set of surface vectors that are consistent with the observed sensor responses. It follows that when the number of dimensions in the surface representation equals (or exceeds) the number of sensor classes, we do not have enough data to obtain a unique solution.

This is the first important difference between the synthesis of images and the analysis of images. When we synthesize images, we can use a surface representation with dimensionality equal to or greater than the number of independent display guns. In order to analyze an acquired color image, we must use a surface representation whose dimensionality is strictly less than the number of sensor classes. The result emphasizes what is possible—

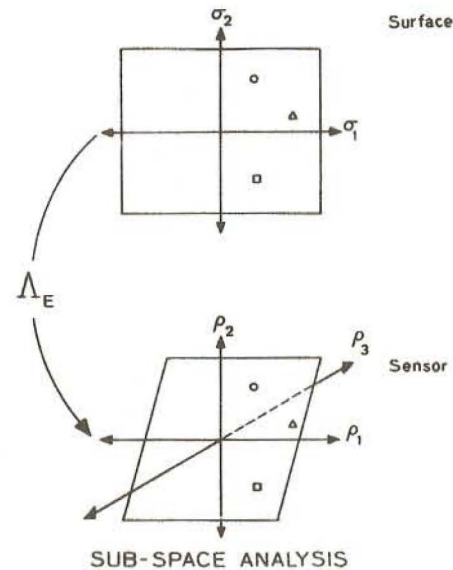


Fig. 2. The surface vectors are represented in a two-dimensional space, shown at the top of the figure. The effect of the light is to act as a linear transformation that maps these two-dimensional vectors into a subspace of the three-dimensional space of receptor responses.

and what is not possible—for systems with different numbers of color sensors. A system with three color sensors, for example, can only recover surfaces and lights that fall within a two-dimensional linear model. Errors will occur should the image data be composed of surfaces and lights requiring a higher dimensional representation. The significance of these errors for a system with three sensors will depend on the nature of the subsequent processing and the severity of the failure of the two-dimensional model.

The Recovery Method

From (16),

$$\rho^x = \Omega_E S \sigma^x = \Lambda_E \sigma^x, \quad (27)$$

we see that when the dimensionality of the surface vector σ^x is less than the dimensionality of the receptor vector ρ^x , the observed data will be restricted to a subspace of the receptor responses. For example, when the set of surface vectors is a two-dimensional space, the image of the surface vectors under a fixed light will span a plane within the three-dimensional space of sensor catches. This is illustrated in Fig. 2. The plane spanned in receptor space is determined by the columns of the matrix Λ_E . The columns of Λ_E are determined entirely by the lighting parameters ϵ . It follows, therefore, that the identity of the subspace is determined by the values of the light parameters, and the position within the subspace is determined by the surface parameters σ^x .

We can use this fact to achieve recovery in a three-step process. First, we identify the subspace which contains the set of receptor responses. Second, we use the identity of the subspace to estimate the lighting vector ϵ . Finally, we use the lighting parameters to determine the pseudoinverse of the lighting matrix Λ_E^{-1} . This is the final step in

the problem, since with this matrix, we can determine the estimated surface parameters σ^x from the receptor responses ρ^x .

I will now describe the implementation of each of the surface and light recovery steps used in the Stanford color analysis package. This will permit me to describe the points at which failures can occur.

Identifying the Subspace: The first step in the recovery procedure is to attempt to identify the subspace that contains the set of sensor vectors ρ^x . For image analysis, we will generally wish to use a surface and light representation whose dimensionality is as large as possible in order to minimize the error in the analysis. The best fit can be obtained using the linear models with the largest degrees of freedom. Thus, we will assume that $D(E)$ and $D(S) + 1$ equal the number of sensors since any other assumption will produce an inferior fit. With this assumption, we expect the data to fall in a subspace of the sensor response that can be uniquely characterized by the unit normal vector, perpendicular to all of the vectors in the subspace. Call the unit normal vector, perpendicular to the subspace, π . The dot product of this vector with each of the sensor vectors ρ^x will be zero.

In fact, of course, the data are not likely to fall precisely within a subspace, but rather to span the sensor response space. Our strategy, however, is to find the vector that is "most perpendicular" to the data from the sensor vectors, under the assumption that perturbations outside of the subspace are due to noise. To find the most perpendicular vector, we form the matrix Δ whose rows consists of the sensor values in a region of the image. We then solve the homogeneous linear equation

$$\Delta\pi = 0 \quad (28)$$

where 0 is the zero vector.

In any particular image, the best perpendicular vector π may not be very perpendicular to the data set. That is, the sensor values may not be well confined to a proper subspace. This represents the first potential failure of the recovery method.

The principal difficulty in this case is that the dimensionality of the surfaces and lights that generate the image are actually equal to or greater than the number of sensors in the camera. If this is the case, nothing can be done and accurate recovery is impossible using this particular sensing apparatus. The recovery procedure then has the virtue that it warns the user that accurate recovery cannot take place since the data are not well confined to a subspace. I discuss the selection of good basis functions in the next section.

Notice that in order for the subspace to be properly defined, there must be enough independent surfaces to define the subspace. In general, the required number is $D(E) - 1$ (Maloney [11]). In practice, we require the spatial partitions of the image to be large enough so that each of the overlapping areas in Fig. 1 has enough surface variation to permit a solution.

Finally, notice that if the relative spectral power distri-

bution is constant, but the absolute level of the ambient light varies, the data will continue to fall within a subspace defined by the relative spectral power distribution. Variation in the power of the ambient light may occur due to the geometry of point source illuminations, shadowing, or other factors. This type of spatial variation does not pose a problem for the recovery procedure, but it emphasizes the fact that absolute levels of the surface spectral reflectance and ambient light distributions, in particular for widely separated spatial locations, are not uniquely recovered.

As an alternative to solving the homogeneous linear equation (28) in order to find the perpendicular vector, one may perform a singular value decomposition of the data matrix Δ . The direction corresponding to the smallest singular value would then be selected as the best normal direction to the data.

Recovering the Light Vector: We now wish to recover the light vector ϵ from the vector perpendicular to the sensor data π . We can accomplish this by expanding (28) in the following way.

First, notice that the requirement that π be perpendicular to each observed sensor vector may be put as the set of equations

$$\pi' \rho^x = \pi' (\Lambda_E \sigma^x) = 0 \quad (29)$$

where π' is the transpose of π , and (29) holds for every surface vector σ^x within the model. We may take advantage of the symmetric role played by the light and the surface at each point in the image to rewrite the matrix product $\Lambda_E \sigma^x$ as

$$\Lambda_E \sigma^x = \Lambda_{S^x} \epsilon \quad (30)$$

where the ik th entry of Λ_{S^x} is

$$\sum_{n=0}^N E_i(\lambda_n) S^x(\lambda_n) R_k(\lambda_n). \quad (31)$$

It follows that for any surface within the linear model,

$$\pi' \Lambda_{S^x} \epsilon = 0. \quad (32)$$

For each independent surface in the model, we have a linear equation in ϵ . We may form a matrix Π whose rows are all of the form $\pi' \Lambda_{S^x}$, where the choice of surface function $S^x(\lambda_n)$ varies from row to row. (Note: The surfaces $S^x(\lambda_n)$ need not be in the actual image since the equation must hold for any surface whatsoever.)

Since it is impossible to measure the absolute intensity of the ambient lighting—there is always a free scalar in the estimation of ϵ —it is convenient to choose ϵ so that its first entry is one. By building Π from enough independent surface vectors $S^x(\lambda_n)$, we can solve for epsilon using

$$\Pi\epsilon = 0. \quad (33)$$

Since the first entry of ϵ is normalized to one, we may convert this homogeneous linear equation into a nonhomogeneous equation by moving the first column of Π to the right-hand side. The equation then becomes a non-

homogeneous equation with one less dimension, and we may solve for the remaining entries of ϵ .

Solving for Surfaces Once the Light Is Known: Once the lighting vector is known, the remainder of the problem is straightforward. The dependency of the lighting matrix Λ_E on the lighting vector is made explicit in (17). This matrix may be explicitly computed and its pseudoinverse computed using conventional techniques. The surface vectors can then be computed using

$$\Lambda_E^{-1} \rho^x = \sigma^x. \quad (34)$$

Two difficulties may arise at this stage in the analysis. First, the estimated Λ_E may not be one to one. In this case, there will be no unique solution for the surfaces unless the dimensionality of the surface representation is reduced.

Second, the recovered surface weights σ^x , while consistent with the model, may in fact give rise to estimates of the surface reflectance functions that are not physically realizable, that is, that contain negative values of surface reflectivity. When this happens broadly throughout an image, it is an indication that the linear model for the lights and surfaces is not providing a very good fit to the actual lights and surfaces. This too requires reevaluating the choice and number of basis functions used in the estimation procedure.

Noise Analysis

Two categories of error can enter into the calculation. The first source of error is measurement noise introduced at the level of the sensor responses. The second source of error is failures of the linear models of lights and surfaces to properly characterize the surfaces and lights. The problem of sensor noise is relatively straightforward. If we model sensor noise by a vector η^x added to true sensor vector ρ_0^x

$$\rho^x = \rho_0^x + \eta^x, \quad (35)$$

then the error analysis will depend upon the distribution of the noise vector η^x . Generally, for this source of noise, the typical assumptions of independence of the noise vectors across the image are adequate. The analysis presented in the previous sections is the maximum likelihood calculation under the assumption that the entries of the noise vector are Gaussian and independent.

Model Errors: If the dimensionality of the linear model representations of lights and surfaces $D(E)$ or $D(S)$ are allowed to become arbitrarily large, then model error can be made arbitrarily small. As the dimensionality of the linear representations is restricted to provide increasing efficiency, then an additional error will be introduced by the model approximations to the actual surface and light functions.

Suppose that the dimensionalities of the surface and light functions required to produce an essentially error-free representation of the data are M and N . To reduce the size of the computation, we will seek linear models in

$$\rho^x = \Lambda_I \sigma^x + \dots + \Lambda_{D(E)} \sigma^x + \dots + \Lambda_M \sigma^x$$

Fig. 3. Figure showing the full model matrix equations with the error terms inside of the reduced model.

which $D(E) < M$ and $D(S) < N$. The sensor data will be correctly modeled by an equation of the form

$$\rho^x = \sum_{n=1}^M \epsilon_n \Lambda_n \sigma^x \quad (36)$$

in which σ^x is an N -dimensional vector. The analysis, however, will be based upon the model equation

$$\rho^x = \sum_{n=1}^{D(E)} \epsilon_n \Lambda_n \sigma^x \quad (37)$$

in which σ^x is a $D(S)$ -dimensional vector. The difference between the measured sensor data (36) and the best model estimates (37) primarily depends on the ambient light and the surfaces in the image. The error terms that express the difference between these two equations may be described by an additional vector with contributions from various sources, as illustrated in Fig. 3. The distribution of this error vector, then, depends upon the surfaces in the image as well as the ambient light falling on the surfaces. Since we cannot specify the sampling distribution of the actual surfaces without some more knowledge of the specific task of the visual sensor (a robotic inspection system will sample a different class of surfaces from a color correction circuit in a home video recorder), the distribution of the error signal cannot be specified without further information concerning the class of materials and ambient lights that the system will encounter.

THE EFFICIENT REPRESENTATION OF SURFACES AND LIGHTS

Introduction

The choice of the linear models for the ambient light and surface spectral reflectance terms is crucial in building an efficient computation in the synthesis of images, and it is crucial in permitting one to use a camera with only a few color sensors in the analysis of image data. In this section, I will review the selection of basis functions that permit a more efficient representation and computation of the lighting and surface parameters.

Fourier Basis

The principal means available to us for reducing the computational burden is to reduce the dimensionality of the surface spectral reflectance N . Were the surface spectral reflectance functions and the ambient light spectral

power distributions essentially bandlimited functions, then from the sampling theorem, we know that specifying these functions at, say, four sampling points along the wavelength dimension would uniquely determine the identity of the function.

Stiles and Wyszecki [17] and Stiles *et al.* [16] performed an early and important analysis of surfaces reflectance functions in which they suggested that these functions are low-pass functions of wavelength in the visible range. Maloney [11] has put forth the intriguing hypothesis—based on the physical chemistry of surface reflectances—that these functions are generally forced to be low pass in the visible range.

There is wide informal agreement that natural surfaces and lights are generally smoothly varying. There are certain exceptions, such as fluorescent light and the surface reflectance functions of the rare earth metals, but the fact that the overwhelming majority of functions are smoothly varying suggests that in designing a system with only a few color sensors, it is far better to select a representation of the surfaces and lights that is slowly varying across wavelength. Unless there is some specific knowledge of the working environment to override this choice, it seems generally preferable to represent the spectra as the weighted sum of several low-frequency terms across the visual spectrum.

Specular Reflectance: When using a low-frequency representation, the functions $S_j(\lambda_n)$ are illustrated in panel (b) of Fig. 4. The function S_1 is the dc component of the spectral reflectance function, and the functions S_2 and S_3 represent the sinusoidal and cosinusoidal components at one cycle across the visible region. Notice that when this representation is chosen, the weight of the first dimension S_1 is a mixture of the specular reflectance at that point in the image, added together with the dc component of the matte portion of the surface reflectance function. Using such a representation, then, it is relatively straightforward to systematically vary the specularity. In practice, this can be achieved by simply manipulating the value of the first dimension of the surface reflectance function vector.

In Fig. 5, I have analyzed how well the Munsell chips, a widely used set of color surfaces, can be described as bandlimited functions. The figure is derived as follows. First, I used the spectral reflectance functions of the Munsell chips measured by Nickerson (1943)⁴ to calculate the least squares fit of Fourier basis terms to the surface spectral reflectance functions of each of the Munsell chips, measured from 400 to 700 nm. I then normalized the root mean-squared error (rmse) of the best fit by the vector length of the surface reflectance function. This is the square root of the proportion of variance accounted for. On the vertical axis, I have plotted the normalized rmse of the median of the 462 Munsell chips available in our data as a function of the number of Fourier basis terms. In Fig. 4(a), I have plotted the worst of the 462 fits to the sample data, along with the three basis functions.

⁴I thank the Munsell Corporation for making these measurements available.




Fig. 4. (a) The surface reflectance function (open symbols) and best fitting estimate (filled symbols) of the Munsell chip with the largest residual error (worst fit) of all 462 chips. (b) The three Fourier basis functions used to fit the surface reflectance functions of the Munsell chip samples.

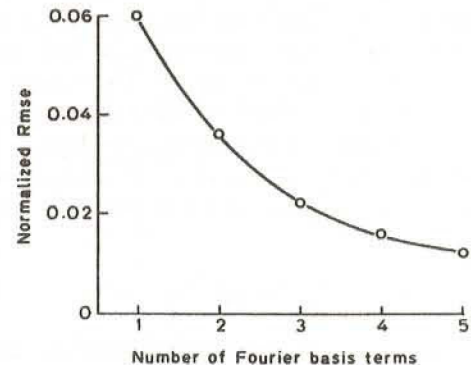


Fig. 5. Each of 462 Munsell chip surface reflectance functions was fit by the sum of sinusoidal and cosinusoidal functions. This figure plots the normalized rmse of the median of 462 Munsell chip fits as a function of the number of Fourier basis terms used in the fit. The percent of variance accounted for may be computed as $\sqrt{1 - y}$ where y is the normalized rmse.

Clearly, the fit of the Munsell spectra is rather good. This should not be taken as evidence that *all* surface reflectance functions fall well within a three-dimensional model. Rather, it should be seen as an interesting case of a complex set of objects that can be well described by a linear model with only a few degrees of freedom. The point of the analysis is to demonstrate that the Fourier representation is efficient for representing spectral reflectance functions and ambient light because of the high degree of correlation across the wavelength spectrum. It has the further advantage of identifying specular reflectance with the dc component of the surface reflectance function.

An alternative method for selecting the representation, which is particularly appropriate when the set of surfaces and lights are known in advance, is to perform a principal component analysis of the distribution of surface reflectance functions and spectral power distributions. The principal component analysis—sometimes called the Karhunen–Loeve transformation—explicitly seeks that set of basis functions that simultaneously minimizes the correlation among dimensions and identifies those dimensions that are most descriptive of the data set.

General analyses of the degrees of freedom in spectral data have been applied to the spectral power distribution of ambient daylights by Judd *et al.* [10] and Dixon (1978) and to the surface reflectances of the Munsell chips by Cohen [4]. In addition, Maloney [11] has analyzed the data from a large set of natural objects. All authors have concluded that their data sets are well described using linear models with a small number of degrees of freedom.⁵

Relationship to Physical Processes: It is interesting to note that the linear model for the variations in the ambient daylight spectral power distribution is not a reasonable model for the physical processes that give rise to the observed spectral power distributions (compare the argument in Witkin and Pentland [21]). For example, blackbody radiators form a one-parameter nonlinear family of spectral power distribution curves. Yet Maloney has shown (personal communication) that the set of principal components of the natural daylights also provides a rather good approximation to the spectral power distribution seen in blackbody radiators. The physical effects that cause the fluctuation in the spectral power distribution of daylight include such factors as Rayleigh scattering and the absorption and scattering by atmospheric moisture. Were one to provide a serious physical model of these effects, the linear model would never be selected. Normally such processes are modeled by a multiplicative filtering action and nonlinear molecular actions. The linear model used here is not closely tied to the physics of the imaging environment, but rather is simply a useful mathematical method. Its usefulness depends on our ability to accurately describe the wavelength functions of the desired quantities. The fact that the underlying physical processes are not well modeled by the linear sum of different physical processes is not relevant to the performance of the algorithm.

The Color Camera Revisited: In Fig. 6, I indicate how the analysis and synthesis calculations developed in this paper can be put together in order to obtain a rendering of an image sensed by a color camera. As Horn [9] has pointed out, if the camera sensor spectral sensitivities are within a linear transformation of the spectral sensitivities of the human eye, the problem may be easily solved. The

⁵Judd *et al.* [10] estimated the principal components of daylight at an early stage in the technical development of these methods. The reported functions are not precisely orthogonal. Unfortunately, the original data have been misplaced so that we cannot reanalyze the original data. Maloney (personal communication) is preparing a new analysis of a large set of such data.

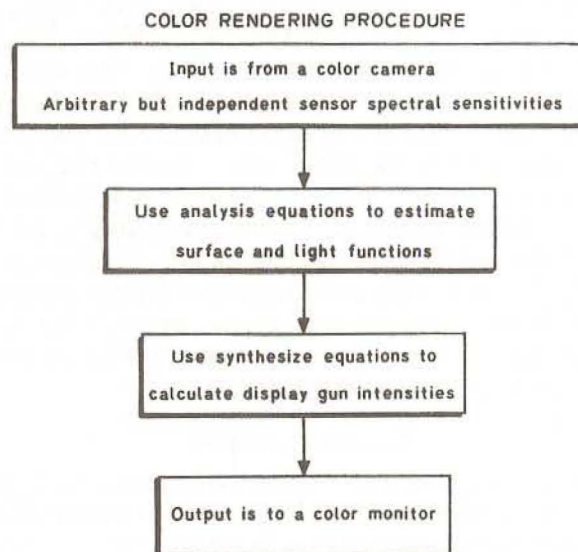


Fig. 6. Flow diagram of how the analyses and synthesis procedures may be used together to convert the outputs of a camera with arbitrary sensor spectral to display an image so that it will appear to an observer as if he were at the camera position.

more difficult—and common—case is when the camera sensors are not related to the visual pigment sensitivities by a linear transformation.

In that case, we may proceed as indicated in Fig. 6. We use the camera sensor responses and the procedure in the analysis section to form an estimate of the ambient light spectral power distribution E . We then calculate estimates of the surface reflectances:

$$\sigma^x = \Lambda_E^{-1} \rho^x. \quad (6)$$

Using these estimates of the surface and ambient light distributions, we set the display intensities using the synthesis calculations in (21), which I repeat here:

$$I^x = \left[\sum_{n=1}^{D(E)} \epsilon_n \Gamma^{-1} \Lambda_n \right] \sigma^x. \quad (7)$$

To the extent that the camera data permit an accurate recovery of the surface and light functions, the synthesized colors will have the same visual effect as if the observer had been at the position of the camera.

CONCLUSIONS

The methods described here for the synthesis and analysis of color images are based on the physical factors—lighting, reflectance, and absorption—that give rise to color information. The insistence on referring all calculations of sensor responses back to the underlying physical processes permits us to treat several different problems in the theory of color images from a common framework. In particular, the equations for selecting display gun intensities, and the equations for analyzing color images, are both based on a common formula, (27).

Maloney and Wandell's [12] method for analyzing the color signal into its surface and light components has applications to several different problems in color science.

The analysis of the image data into surface and light components permits us to correct the acquired sensor signal for unwanted effects of ambient lighting. The ability to perform this correction will improve the art of color balancing images. The ability to estimate the surface reflectance of objects in the image from cameras that are not strictly matched to the human eye should also improve the color rendering of images. Finally, the ability to use *object* surface reflectance—rather than gray-level *image* intensity in which confound surface and light information are confounded—should become an important aid in the development of visual inspection systems that are robust with respect to the ambient lighting.

ACKNOWLEDGMENT

The author would like to thank D. Brainard, L. T. Maloney, C. McConnell, and D. Varner for useful discussions and their comments on this paper.

REFERENCES

- [1] R. M. Boynton, *Human Color Vision*. New York: Holt, Rinehart and Winston, 1979.
- [2] G. Buchsbaum, "A spatial processor model for object color perception," *J. Franklin Inst.*, vol. 310, 1980.
- [3] G. Buchsbaum and A. Gottschalk, "Chromaticity coordinates of frequency-limited functions," *J. Opt. Soc. Amer.*, vol. 1, pp. 885-887, 1984.
- [4] J. Cohen, "Dependency of the spectral reflectance curves of the Munsell color chips," *Psychon. Sci.*, vol. 1, pp. 369-370, 1964.
- [5] R. L. Cook and K. E. Torrance, "A reflectance model for computer graphics," *Comput. Graphics*, vol. 15, pp. 307-316, Aug. 1981.
- [6] T. N. Cornsweet, *Visual Perception*. New York: Academic, 1970.
- [7] H. Helson, "Fundamental problems in color vision. I. The principle governing changes in hue, saturation and lightness of non-selective samples in chromatic illumination," *J. Exp. Psychol.*, vol. 23, pp. 439-476, 1938.
- [8] B. K. P. Horn, "Determining lightness from an image," *Comput. Graphics Image Processing*, vol. 3, pp. 277-299, 1974.
- [9] —, "Exact reproduction of colored images," *Comput. Vision, Graphics, Image Processing*, vol. 26, pp. 135-167, 1984.
- [10] D. B. Judd, D. L. MacAdam, and G. W. Wyszecki, "Spectral distribution of typical daylight as a function of correlated color temperature," *J. Opt. Soc. Amer.*, vol. 54, p. 1031, 1964.
- [11] L. Maloney, "Computational approaches to color vision," Ph.D. dissertation, Stanford Univ., Stanford, CA, 1984.
- [12] L. T. Maloney and B. Wandell, "A computational model of color constancy," *J. Opt. Soc. Amer.*, vol. 1, pp. 29-33, 1986.
- [13] A. V. Oppenheim and R. W. Schaffer, Ed., *Digital Signal Processing*. Englewood Cliffs, NJ: Prentice-Hall, 1975.
- [14] P. Sallstrom, "Colour and physics; Some remarks concerning the physical aspects of human colour vision," *Inst. Phys. Rep.* 73-09, 1973.
- [15] V. Smith and J. Pokorny, "Spectral sensitivity of the foveal cone photopigments between 400 and 500 nm," *Vision Res.*, vol. 15, pp. 161-171, 1975.
- [16] W. S. Stiles, G. Wyszecki, and N. Ohta, "Counting metameric object-color stimuli using frequency-limited spectral reflectance functions," *J. Opt. Soc. Amer.*, vol. 67, no. 6, pp. 779-784, 1977.
- [17] W. S. Stiles and G. W. Wyszecki, "Counting metameric object colors," *J. Opt. Soc. Amer.*, vol. 52, pp. 313-327, 1962.
- [18] K. E. Torrance and E. M. Sparrow, "Theory for off-specular reflection from roughened surfaces," *J. Opt. Soc. Amer.*, vol. 57, pp. 1105-1114, 1967.
- [19] K. E. Torrance, E. M. Sparrow, and R. C. Birkebak, "Polarization, direction distribution, and off-specular peak phenomena in light reflected from roughened surfaces," *J. Opt. Soc. Amer.*, vol. 56, pp. 916-925, 1966.
- [20] B. Wandell and L. Maloney, "Computational methods for color identification," in *Proc. OSA Annu. Meet.*, 1984, p. 21.
- [21] A. Witkin and A. Pentland, "Perception and generic modeling," *J. Opt. Soc. Amer.*, vol. 12, p. 1261A, 1984.



Brian A. Wandell was born in New York, NY, on October 6, 1951. He received the B.S. degree in mathematics and psychology from the University of Michigan, Ann Arbor, in 1973, and the Ph.D. degree from the University of California, Irvine, in 1977.

In 1975 he spent a year as a Visiting Scholar at Columbia University, New York, NY, in 1976 a year as a Visiting Scholar at Harvard University, Cambridge, MA, and in 1977 he spent a year and a half as a Post-Doctoral Fellow at the University of Pennsylvania, Philadelphia. In January 1979 he was appointed to the Department of Psychology, Stanford University, Stanford, CA, where he is currently an Associate Professor. Apart from his work on computational vision, he is generally engaged in measuring and analyzing human visual performance.

Dr. Wandell is currently a member of the Editorial Board of *Vision Research* and a member of the Association for Research in Vision and Ophthalmology, the Optical Society of America, and Director of the NASA-Ames-Stanford Center of Excellence in model-based human performance.