Black Light: How Sensors Filter Spectral Variation of the Illuminant

DAVID H. BRAINARD, BRIAN A. WANDELL, AND WILLIAM B. COWAN

Abstract—Visual sensor responses may be used to classify objects on the basis of their surface reflectance functions. In a color image, the image data are represented as a vector of sensor responses at each point in the image. This vector depends both on the surface reflectance function and on the spectral power distribution of the ambient illumination. Algorithms designed to classify objects on the basis of their surface reflectance functions typically attempt to overcome the dependence of the sensor responses on the illuminant by integrating sensor data collected from multiple surfaces.

In machine vision applications, we show that it is often possible to design the sensor spectral responsivities so that the vector direction of the sensor responses does not depend upon the illuminant. We state the conditions under which this is possible and perform an illustrative calculation.

In biological systems, where the sensor responsivities are fixed, we show that some changes in the illumination cause no change in the sensor responses. We call such changes in illuminant *black illuminants*. It is possible to express any illuminant as the sum of two unique components. One component is a black illuminant. We call the second component the *visible component*. The visible component of an illuminant completely characterizes the effect of the illuminant on the vector of sensor responses.

Introduction

A USEFUL function of a color visual system is to classify objects on the basis of their surface reflectance functions. The spectral power distribution of the color signal arriving at the visual system depends on both the surface reflectance functions of the imaged objects and the spectral power distributions of the illuminants. This observation is often used to argue that it is impossible to estimate the surface reflectance function of an object from the vector of sensor responses at a single location in the image. Most algorithms designed to classify objects on the basis of their surface reflectance functions assume that the illuminant varies slowly across the image. They incorporate information from multiple locations in the image to separate the effects of surface reflectance functions and the illuminant spectral power distribution [1]–[9].

We reexamine the dependence of the color signal on the surface and illuminant. We show that when the surface reflectance functions and illuminant spectral power distributions encountered by the visual system are described by certain small-dimensional linear models, the color signal contains enough information so that the vector of sensor responses at each location in the image can be used directly to classify objects on the basis of their surface reflectance functions without any confound from the spectral power distribution of the illuminant. To do this, the spectral responsivities of the sensors must be matched to the surface and illuminant models. When the sensors, surfaces, and illuminants are matched, the vector direction of the sensor responses depends on the surface reflectance function but not on the spectral power distribution of the illuminant.

In the first part of this paper we describe conditions under which it is possible to match the sensors to the surface and illuminant models so that the vector direction of the sensor responses remains fixed as the spectral power distribution of the illuminant is varied. As an example, we show that these conditions hold for a realistic model of natural surfaces and daylight illuminants. We carry out the design of sensor responsivities for this model. These sensor responsivities remove, up to scalar multiplication, the dependence of the sensor responses on the spectral power distribution of the illuminant. In this sense, the sensors act to filter illuminant variation.

The analysis in the first part of the paper is useful for machine vision applications when the sensor spectral responsivities of the visual system may be chosen. In the study of biological visual systems, the sensor responsivities are given by the biology. In the second part of this paper, we analyze the role of any given sensor responsivities in filtering illuminant variation. Rather than asking what set of visual sensors filter a given class of illuminant variation, we ask what class of illuminant variations are filtered by a given set of visual sensors. When the surface reflectance functions encountered by the visual system are described by a small-dimensional linear model, there are a large number of illuminant variations that cause no change in the sensor responses. We call such variations *black illuminants*. Given a set of sensor responsivities and a linear model for surfaces, we show how to calculate the black illuminants.

It is possible to express any illuminant as the sum of two unique components. One component is a black illuminant. We call the second component the *visible com-*

Manuscript received February 1, 1988; revised May 27, 1988. This work was supported by Grant R01 EY03164 from the National Eye Institute and Grant NCC-2-307 from the NASA-Ames Research Center. Some of the material is based upon work supported under a National Science Foundation Graduate Fellowship.

D. H. Brainard and B. A. Wandell are with the Department of Psychology, Stanford University, Stanford, CA 94305.

W. B. Cowan is with the National Research Council, Optics Section, Division of Physics, Ottawa, Ont., Canada.

IEEE Log Number 8824425.

0018-9294/89/0100-0140$01.00 © 1989 IEEE

ponent. The visible component of an illuminant completely characterizes the effect of the illuminant on the vector of sensor responses. As an example, we compute the black illuminants for the human cones and a realistic model for natural surfaces. We find the black illuminant and visible components of several illuminants.

Definitions and Notation

We consider the situation where the light incident at the camera or eye is due to an illuminant reflected from a surface. An illuminant is described by its *spectral power distribution* $E(\lambda_n)$, which specifies the illuminant power at each wavelength. A surface is described by its *surface reflectance function* $S(\lambda_n)$, which describes the fraction of light reflected at each wavelength. The light incident at the camera or eye is called the *color signal* and is described by its spectral power distribution $C(\lambda_n)$. We measure all functions of wavelength at N_λ evenly spaced sample wavelengths λ_n with wavelength spacing $\Delta\lambda$. The relation between the illuminant spectral power distribution, surface reflectance function, and color signal spectral power distribution is given by

$$C(\lambda_n) = E(\lambda_n)\, S(\lambda_n). \tag{1}$$

A visual system contains several classes of color sensors. We use the subscript k to denote the kth sensor class, and we denote the number of sensor classes by P. For the human eye, $P = 3$ as there are three classes of cones in the retina. The information available to the visual system at a single point is contained in the vector of sensor responses whose entries are r_k, $k = 1, \cdots, P$. The response r_k of a sensor in the kth sensor class depends on the spectral power distribution of the color signal and the *spectral responsivity* of the sensor class. The spectral responsivity of the kth sensor class $R_k(\lambda_n)$ specifies the strength of the sensor response per unit energy as a function of wavelength. The sensor response is computed as

$$\begin{aligned} r_k &= \sum_{n=1}^{N_\lambda} R_k(\lambda_n)\, C(\lambda_n)\, \Delta\lambda \\ &= \sum_{n=1}^{N_\lambda} R_k(\lambda_n)\, S(\lambda_n)\, E(\lambda_n)\, \Delta\lambda. \end{aligned} \tag{2}$$

To simplify notation, we suppress the factor $\Delta\lambda$ by assuming that it is incorporated into the physical units.

If the surface reflectance functions and illuminant spectral power distributions encountered by a visual system are unrestricted, then any color signal can arise from any surface. The sensor responses cannot be used to classify objects based on their surface reflectance functions unless some additional constraints are used. Recent work [2]–[6], [10], [11] is based on the assumption that the surface reflectance functions and illuminant spectral power distributions encountered by the visual system are restricted. In particular, this work has assumed that the surfaces and illuminants are described by small-dimensional linear models. We say that the surfaces are described by a linear model if there exist N linearly independent *basis functions* $S_j(\lambda_n)$ and weights σ_j such that for any surface $S(\lambda_n)$ encountered by the visual system

$$S(\lambda_n) = \sum_{j=1}^{N} \sigma_j S_j(\lambda_n). \tag{3}$$

We call N the dimension of the linear model for surfaces. Similarly, we say that the illuminants are described by a linear model if there exist M linearly independent basis functions $E_i(\lambda_n)$ and weights ϵ_i such that for any illuminant $E(\lambda_n)$

$$E(\lambda_n) = \sum_{i=1}^{M} \epsilon_i E_i(\lambda_n). \tag{4}$$

We call M the dimension of the linear model for illuminants.

When the surfaces and illuminants are described by linear models, then the color signals formed when one of the illuminants reflects from one of the surfaces are also described by a linear model. For any such color signal, we have

$$C(\lambda_n) = \sum_{i=1}^{M} \sum_{j=1}^{N} \epsilon_i \sigma_j E_i(\lambda_n)\, S_j(\lambda_n) = \sum_{i,j} \kappa_{ij} C_{ij}(\lambda_n). \tag{5}$$

The basis functions for the color signal linear model are $C_{ij}(\lambda_n) = E_i(\lambda_n)\, S_j(\lambda_n)$. The weights κ_{ij} on each of the color signal basis functions are determined by the weights σ_j and ϵ_i of the underlying surface and illuminant. The dimension of the linear model for the color signals is at most NM. If the $C_{ij}(\lambda_n)$ are not linearly independent the dimension will be less than NM.

Filtering Illuminant Variation

Basic Ideas

When the surfaces and illuminants are described by small-dimensional linear models, the color signal is described by the sum in (5). It is useful to separate the terms of this sum into two groups: one group that contains terms that depend on the first basis function of the illuminant, and a second group that contains terms that depend on the other basis functions of the illuminant:

$$C(\lambda_n) = \left[\sum_{j=1}^{N} \kappa_{1j} C_{1j}(\lambda_n)\right] + \left[\sum_{i=2}^{M} \sum_{j=1}^{N} \kappa_{ij} C_{ij}(\lambda_n)\right]. \tag{6}$$

We call the first group the *principal illuminant term*, and the second group the *illuminant variation term*. As the spectral power distribution of the illuminant varies, the only variation in the principal illuminant term is multiplication by the scalar ϵ_1. If a visual system's sensors respond to the contribution to the color signal from the principal illuminant term but not to the contribution from the illuminant variation term then the vector of sensor responses to any surface will remain the same, up to the scale factor ϵ_1, as the illuminant is varied. For such sensors the vector direction of the sensor responses can be used to classify objects on the basis of their surface reflectance. Conversely, variation in the vector direction of the sensor responses can be relied upon to indicate variation in object surface reflectance. We say that such sensors filter the variation in the illuminant spectral power distribution.

Under what conditions can we design sensor spectral responsivities so that the sensors respond to the contribution to the color signal of the principal illuminant term and not to the contribution of the illuminant variation term? In the Appendix we show that if each of the basis functions $C_{ij}(\lambda_n), i = 2, \cdots M, j = 1, \cdots N$ in the illuminant variation term is linearly independent of the basis functions $C_{1j}(\lambda_n), j = 1, \cdots N$ in the principal illuminant term, then there exist such sensor spectral responsivities. This condition will be met for many choices of surface and illuminant basis functions. This condition will not be met, however, when the surface and illuminant linear models are identical. In this case, $C_{ij}(\lambda_n) = C_{ji}(\lambda_n)$ so that the principal illuminant term and the illuminant variation term contain a common basis function.

When the basis functions of the illuminant variation term are linearly independent of the basis functions of the principal illuminant term, we can design sensor spectral responsivities that filter the illuminant variation. If the P sensors of a visual system all filter the illuminant variation term, then the sensor responses r_k to any color signal encountered by the visual system are given by

$$r_k = \sum_{n=1}^{N_\lambda} R_k(\lambda_n)\, C(\lambda_n) = \sum_{n=1}^{N_\lambda} R_k(\lambda_n) \sum_{j=1}^{N} \kappa_{1j} C_{1j}(\lambda_n). \tag{7}$$

Equation (7) does not determine the sensor responsivities uniquely. There will generally be many different sets of P sensor responsivities for which (7) holds. Not all of these sets are equally useful for classifying objects on the basis of surface reflectance since some sets of sensor responsivities will fail to preserve information about the surfaces contained in the principal illuminant term. If we view the color signal weights κ_{1j} as a N dimensional vector, (7) defines a matrix mapping from this vector to the P dimensional vector of sensor responses. The kjth entry of the matrix is $\Sigma_{n=1}^{N_\lambda} R_k(\lambda_n)\, C_{1j}(\lambda_n)$. If this matrix does not have maximal rank, then the vector of sensor responses does not contain as much information as possible about the color signal weights in the principal illuminant term. If $P = N$ and this matrix has maximal rank, then the vector of sensor responses can be used to recover the weights κ_{1j}. Since $\kappa_{1j} = \epsilon_1 \sigma_j$, we can use the κ_{1j} to find the surface reflectance function weights σ_j up to an unknown scale factor ϵ_1. The σ_j can be used reconstruct the relative surface reflectance function using (3). Under what conditions can we find $P = N$ sensor responsivities such that the matrix mapping the κ_{1j} to the r_k has maximal rank? In the Appendix, we show that if the color signal basis functions $C_{1j}(\lambda_n)$ in the principal illuminant term are linearly independent, it is possible to find such sensor responsivities. We also show how, when the appropriate conditions on the $C_{ij}(\lambda_n)$ are met, to compute sensor responsivities such that the sensors do not respond to the illuminant variation term of the color signal and such that the vector of sensor responses can be used to recover the vector of weights κ_{1j}.

An Example Calculation

As an example, we have designed a set of sensor responsivities that filter the illuminant variation of natural daylight when the surfaces are described by a realistic linear model. Cohen [12] analyzed the surface reflectance functions of a large set of Munsell chips and concluded that they are well-described by a linear model with between three and six dimensions. Maloney [11], [13] extended Cohen's analysis to include a large set of natural surface reflectance functions measured by Krinov [14]. Maloney concluded that this larger set of surfaces is well-described by the same linear model. Following Maloney, we computed basis functions for a three-dimensional linear model for 462 Munsell chips measured by Nickerson [15]. This linear model also describes the surface reflectance functions measured by Krinov. Judd, MacAdam, and Wyszecki [16] analyzed a large set of daylight spectral power distributions and concluded they are well-described by a linear model with between three and five dimensions. They provide basis functions for a three-dimensional linear model for daylight illumination. For these linear models for surfaces and illuminants, we verified that the basis functions in the illuminant variation term are linearly independent of the basis functions in the principal illuminant term and that the basis functions in the principal illuminant term are linearly independent of each other. Thus, we were able to calculate a set of three sensor responsivities that filter the variation in daylight illumination when the surfaces are described by our three-dimensional linear model. These sensors also allow the recovery of the three weights κ_{1j} for the basis functions of the principal illuminant term.

In Fig. 1 we plot the three sensor responsivities that filter daylight variation. These sensor responsivities could be implemented in a machine vision system designed to classify objects on the basis of their surface reflectance functions. The negative sensor responsivities can be implemented by taking the difference between the responses of two sensors, each with nonnegative responsivity at every wavelength. Moreover, exciting new developments in spectroradiometer technology [17] may make it possible to measure the entire color signal spectral power distribution in real-time machine vision systems. In this case, the separation of the principal illuminant and illuminant variation terms could be handled with software.

Discussion

Comparison to Human Cone Responsivities

The three sensor responsivities shown in Fig. 1 are more sharply peaked than the human sensor responsivities. Fig.

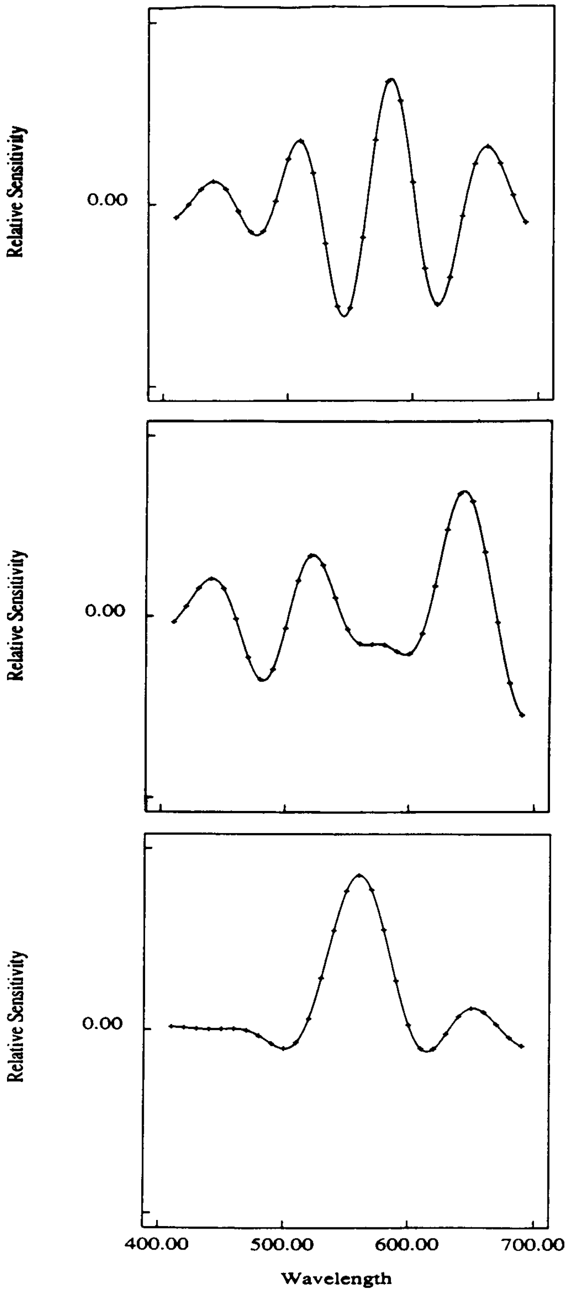

Fig. 1. Three sensor spectral responsivities that filter the illuminant variation term when the surfaces are described by a three-dimensional linear model for natural surfaces and the illuminants are described by the three-dimensional linear model for daylight provided by Judd, MacAdam, and Wyszecki. These sensors also allow the recovery of the color signal basis function weights κ_{1j} in the principal illuminant term.

2 shows the Smith and Pokorny estimate of the human middle-wavelength cone responsivity tabulated by Boynton [18]. The figure also shows the best linear fit to this responsivity using the computed sensor responsivities. The fit is poor. This suggests that the human cones do not completely filter the variation in daylight illumination. We show in the next section that the vector of human cone responses is in fact sensitive to variations in daylight illumination.

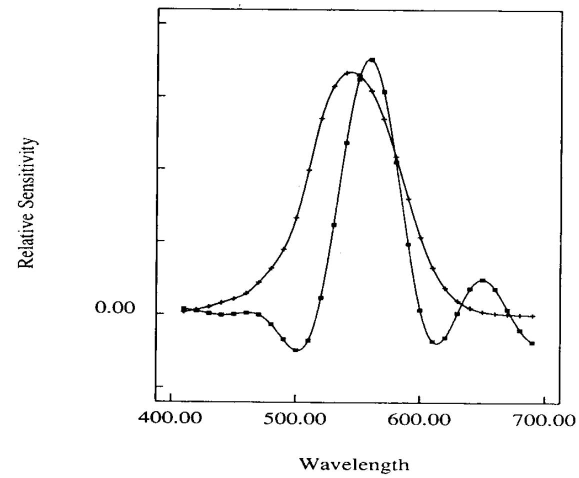

Fig. 2. Smith and Pokorny estimate of the human middle cone sensor responsivity along with the best linear fit using the spectral sensor responsivities computed in Fig. 1.

There may be several reasons why the human cone responsivities are unlike the spectral responsivities shown in Fig. 1. First, all biological photopigments studied to date have spectral responsivity functions that are slowly varying functions of wavelength. Dartnall and others have speculated that the biochemistry of photopigments requires that the human cones have this property [19]–[22]. Thus, one reason why biological cones do not completely filter variation in daylight illumination may be that it is difficult to construct photopigment molecules with the appropriate responsivities.

A second reason may be that the only design goal of biological visual systems is not to classify objects on the basis of their surface reflectance functions. We have purposefully designed sensor spectral responsivities that eliminate all information about illuminant variation from the vector of sensor responses. But the spectral power distribution of the illuminant may provide useful information to an organism: perhaps it is indicative of the time of day or of the weather conditions. A biological visual system may want sensors that permit it to notice variation in the illuminant. In that case, subsequent processing of the sensor responses would be required to classify objects on the basis of their surface reflectance.

Other Applications of Filtering

From a formal point of view there is nothing special about the particular division of the color signal into the principal illuminant term and the iluminant variation term. This grouping is useful when the design goal is to filter the effects of illuminant variation from the sensor responses. If we group the color signal basis functions differently, we can perform the same analysis and try to find sensor spectral responsivities that filter other parts of the color signal. For example, the algorithm developed by Maloney and Wandell [4], [5], [13] requires that the surfaces and illuminants be described by linear models with not more than $P - 1$ and P dimensions, respectively, where P is the number of sensor classes. If the linear models that describe surfaces and illuminants have $N > P - 1$ and $M > P$ dimensions, respectively, then we can group the terms in the color signal as follows:

$$C(\lambda_n) = \left[\sum_{i=1}^{P} \sum_{j=1}^{P-1} \kappa_{ij} C_{ij}(\lambda_n) \right] + \left[\sum_{i=P+1}^{M} \sum_{j=P}^{N} \kappa_{ij} C_{ij}(\lambda_n) \right]. \quad (8)$$

If the each of the basis functions in the second group are linearly independent of the basis functions in the first group, then we can find sensor responsivities that filter the extra dimensions in the surface and illuminant linear models. Such sensor responsivities would optimize the performance of the Maloney and Wandell algorithm.

Analogy to Prefiltering

The role of the sensor responsivities in filtering the color signal is analogous to the role of prefiltering in digital sampling applications. When a signal is to be sampled at discrete points, it is common to filter the signal before the sampling operation is performed. This prefiltering limits the signal's bandwidth and prevents aliasing distortion. The effect of prefiltering is to remove some of the degrees of freedom in the input signal. Suppose the original signal is band-limited and periodic. Then we can express this signal as the weighted sum of a finite number N of Fourier basis functions. If there are only $N/2$ sample measurements, we may decide to use a low-pass filter to remove $N/2$ of the Fourier basis terms. The $N/2$ sample measurements are adequate to reconstruct the reduced signal. If one does not eliminate the additional terms, then the input signal has more degrees of freedom than the number of samples. In this case, the mapping from signals to sample responses is many to one and the sample values cannot be uniquely inverted to estimate the signal. Prefiltering serves to match the number of degrees of freedom in the signal to the number of sample measurements.

There is a precise analogy between prefiltering a signal to avoid aliasing and the method we describe here. Prefiltering the signal renders the sample values insensitive to the higher frequency Fourier basis terms. Our choice of sensor spectral responsivities renders the sensor responses insensitive to the illuminant variation term. The proper choice of low-pass filter permits the sample values to be used to reconstruct the low-frequency component of the input signal. The proper choice of sensor responsivities permits the sensor responses to be used to reconstruct the component of the color signal described by the principal illuminant term. The classic low-pass prefiltering analysis is a special case of the general analysis presented here. The low-pass analysis is applicable to convolution systems. Since the sensor responses are not the output of a convolution system, we have used methods applicable to general linear systems.

Black Illuminants

Basic Ideas

We have shown how to select sensor responsivities that respond to some of the color signal basis functions and filter the others. The analysis can be used when it is possible to design the sensor responsivities. When we analyze human performance, the sensor responsivities are determined by biology and are not under our control. In this section, we extend our analysis to clarify the role played by any given sensor responsivities in filtering illuminant variation.

If a visual system has only a few classes of sensors and the surface reflectance functions encountered by the visual system are restricted, then there will be many pairs of physically distinct illuminants such that the sensor responses to any surface are unchanged as the illuminant is changed from one to the other. We call the difference between such a pair of illuminants a *black illuminant*. A black illuminant is always defined with respect to an ensemble of surfaces and a set of sensor responsivities. The concept of a black illuminant is closely related to the concept of a *metameric black* surface studied by Wyszecki and Stiles [23]–[25] and the forbidden subspace analysis in West and Brill [26]. All black illuminants are positive at some wavelengths and negative at others and thus cannot be realized physically. Nevertheless, they can be used to construct pairs of physically realizable illuminants that leave the sensor responses to any surface unchanged. The illuminant pair is constructed by adding the spectral power distribution of a black illuminant to a physically realizable illuminant.

If the surfaces encountered by a visual system are described by a small-dimensional linear model, it is possible to compute all the black illuminants for the visual system. A black illuminant $E(\lambda_n)$ must satisfy the equation

$$0 = \sum_{n=1}^{N_\lambda} R_k(\lambda_n)\, S_j(\lambda_n)\, E(\lambda_n) \quad (9)$$

for all surface reflectance basis functions $S_j(\lambda_n)$ and all sensor responsivities $R_k(\lambda_n)$. Conversely, any illuminant that satisfies (9) is a black illuminant. Equation (9) defines a set of linear constraints on the black illuminants. Thus, in this case the black illuminants are a linear subspace of all possible illuminants. In the Appendix, we describe how to compute a set of basis functions for this subspace.

We call the subspace of illuminants orthogonal to the black illuminants the subspace of *visible components*. Any illuminant can be written uniquely as the sum of two components, one in each of these orthogonal subspaces. This decomposition of an illuminant reveals the effect of the sensor responsivities in filtering the illuminant. The black illuminant component is completely filtered by the sensors. The sensor responses depend only on the visible component.

An Example Calculation

An an example, we have computed basis functions for the black illuminant and visible component subspaces using our three-dimensional linear model for natural surfaces and the human cone responsivities. We used the Smith and Pokorny estimates of the cone responsivities tabulated by Boynton [18]. In our calculations, we represent all functions of wavelength by their values at 10 nm intervals between 400 and 700 nm. Thus, the space of all illuminants has 31 dimensions. Our calculations reveal that for this example the subspace of black illuminants has 22 dimensions, while the subspace of visible illuminants has nine dimensions.

The computed basis functions allow us to find spectral power distributions that are completely contained in the subspaces of black illuminants and visible components. Fig. 3 shows an illuminant from each subspace. The top panel shows a black illuminant. The bottom panel shows an illuminant from the visible component subspace. For the human cones and our linear model for surfaces, the black illuminants typically have spectral power distributions that are rapidly varying functions of wavelength. The visible component spectral power distributions are typically more smoothly varying.

The black illuminant analysis is useful because any illuminant can be written uniquely as the sum of a black illuminant and an illuminant from the visible component subspace. Only the visible component effects the vector of sensor responses. We can find the black illuminant and visible components of any illuminant by expressing the illuminant spectral power distribution as the weighted sum of the basis functions for the two subspaces and separating the two components in the sum. Fig. 4 shows several examples. In the first two rows, two physically nonrealizable illuminants with sinusoidal spectral power distributions are decomposed. These decompositions also suggest that the space of black illuminants contains mostly functions that vary rapidly as a function of wavelength, while the space of visible illuminants contains more smoothly varying functions. In the third row, a spectral power distribution which is typical of a change in daylight illumination is decomposed [16]. The variation in daylight contains a significant visible component. This shows that the human cone responsivities do not filter all variation in daylight illumination. In the last row, a fluorescent illuminant spectral power distribution is decomposed. The fluorescent illuminant spectral power distribution contains several sharp spikes. The visible component is much more smoothly varying: the cone responsivities do filter the sharp spikes of the fluorescent illuminant. Note that the color rendering properties of an illuminant only depend on its visible component. To decide whether an artificial illuminant will have color rendering properties similar to those of daylight illuminants, for example, the visible components of each should be compared. Our analysis provides a way to compare illuminants that takes the surface reflectances and sensor responsivities into account.

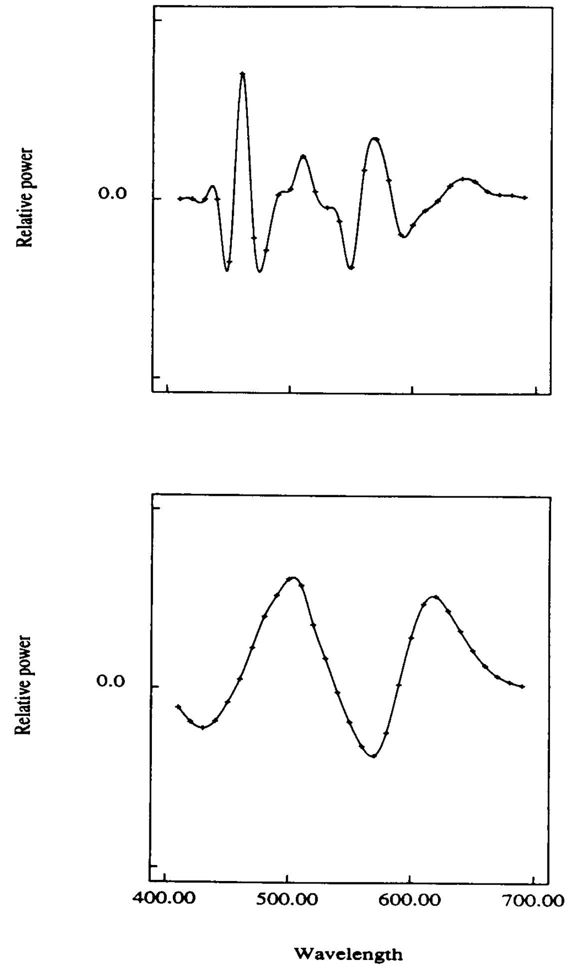

Fig. 3. Examples of illuminants from the black illuminant subspace (top) and visible component subspace (bottom). The basis functions for the subspaces were computed with respect to a three-dimensional linear model for natural surfaces and the human cone spectral responsivities.

Discussion: Relation to Other Work

The analysis of illuminants into two components presented in this section is closely related to the work of Wyszecki and Stiles [23]–[25] who studied *metameric black* surfaces. A metameric black surface is analogous to a black illuminant. A metameric black surface is defined with respect to a particular illuminant. It is a physically nonrealizable surface that when added to any other surface causes no change in the vector of sensor responses. By reversing the role of surfaces and illuminants in our analysis, we can find the metameric black surfaces with respect to any linear model for illuminants.

Another related analysis is that of Cohen and Kappauf [27], [28]. They studied the effect of the human cone responsivities in filtering the color signal without considering the underlying surfaces and illuminants. They point out that any color signal can be written as the sum of two orthogonal components, which they call a *metameric*

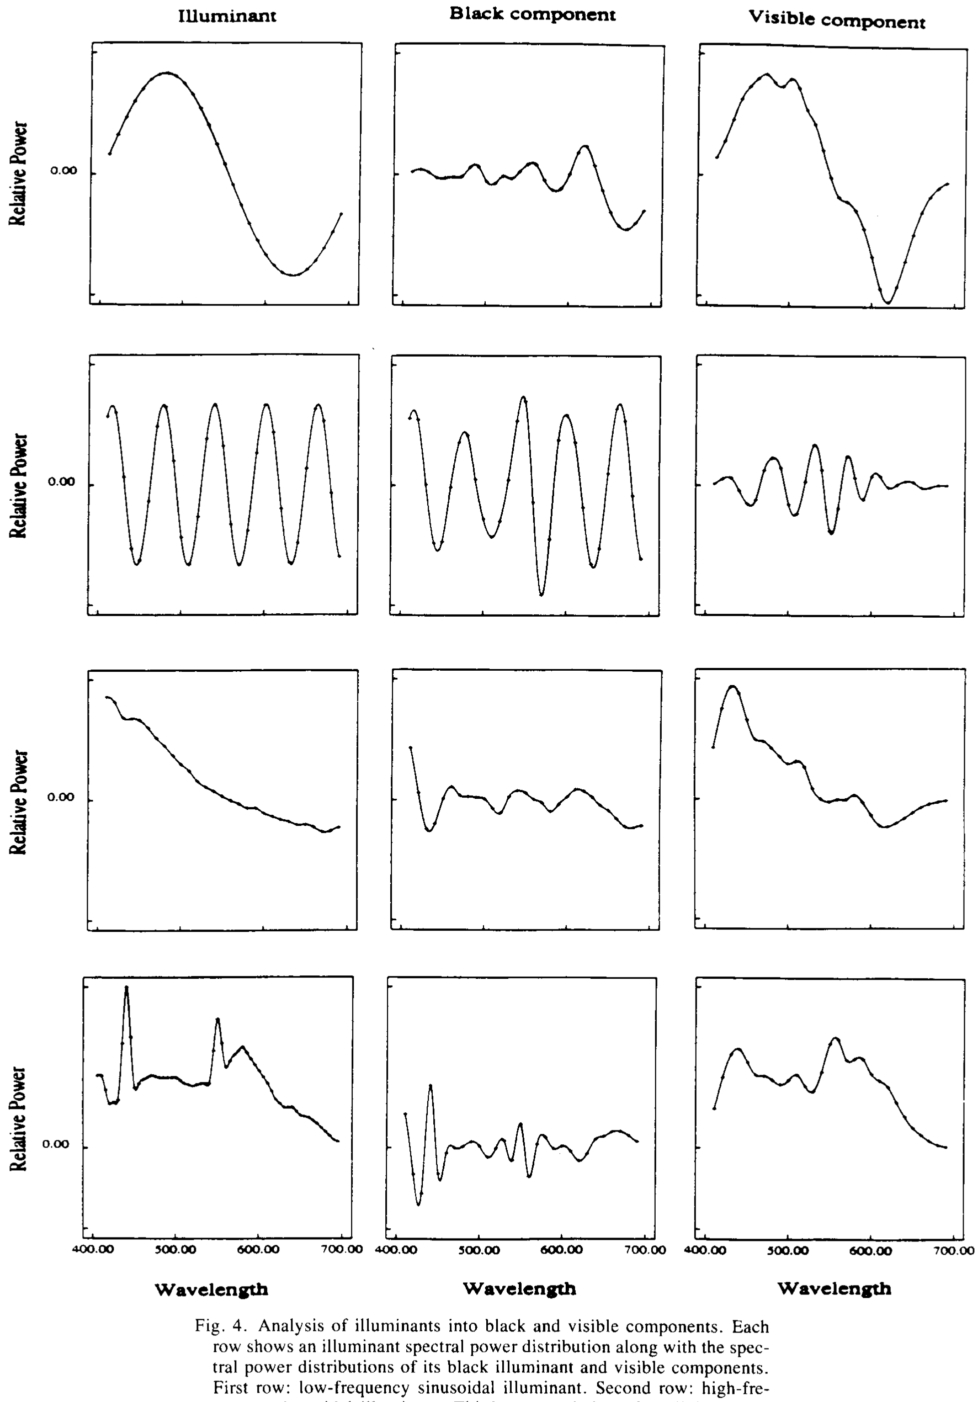

Fig. 4. Analysis of illuminants into black and visible components. Each row shows an illuminant spectral power distribution along with the spectral power distributions of its black illuminant and visible components. First row: low-frequency sinusoidal illuminant. Second row: high-frequency sinusoidal illuminant. Third row: variation of daylight illumination. Fourth row: fluorescent illuminant.

black and a *fundamental metamer*. The metameric black component is filtered by the sensors. The sensor responses depend only on the fundamental metamer.

Summary

This paper has considered the role of the sensors in filtering illuminant variation. In the first part of the paper, we showed that the color signal can contain enough information to allow the proper sensors to separate the surface reflectance function from the illuminant spectral power distribution. In the second part of the paper, we showed that the sensors always act to filter some component of the illuminant. These two analyses help clarify both what information about the environment is contained in the color signal and the important role of the sensor spectral responsivities in visual processing.

Appendix: Computational Methods

Introduction and Notation

We use vectors and matrix notation to describe the computational procedures. Let $\boldsymbol{e}$ be a N_λ dimensional column vector whose nth entry is $E(\lambda_n)$. Let $\boldsymbol{s}$ be a N_λ dimensional column vector whose nth entry is $S(\lambda_n)$. Let $\boldsymbol{c}$ be a N_λ dimensional column vector whose nth entry is $C(\lambda_n)$. In addition, let $\boldsymbol{S}$ be an N_λ by N_λ diagonal matrix whose nth diagonal entry is $S(\lambda_n)$. Then we can write (1) as

$$\boldsymbol{c} = \boldsymbol{S}\boldsymbol{e}. \tag{10}$$

Let $\boldsymbol{r}$ be a P dimensional column whose kth entry is r_k. Let $\boldsymbol{R}$ by a P by N_λ matrix whose knth entry is $R_k(\lambda_n)$. The rows of $\boldsymbol{R}$ are the P sensor spectral responsivity functions. Using this matrix notation, (2) becomes

$$\boldsymbol{r} = \boldsymbol{R}\boldsymbol{c} = \boldsymbol{R}\boldsymbol{S}\boldsymbol{e} = \boldsymbol{L}_s\boldsymbol{e} \tag{11}$$

where $\boldsymbol{L}_s = \boldsymbol{RS}$ is a P by N_λ matrix that describes the relation between the illuminant and the sensor responses for a particular surface $\boldsymbol{s}$.

We can also use vector and matrix notation to express the color signal linear model of (5). We write

$$\boldsymbol{c} = \boldsymbol{C}_m\kappa. \tag{12}$$

κ is an NM dimensional column vector whose entries are the weights κ_{ij}. $\boldsymbol{C}_m$ is an N_λ by NM matrix whose columns, denoted by $\boldsymbol{c}_{ij}$, are the vector representations of the color signal basis functions $C_{ij}(\lambda_n) = E_i(\lambda_n)\, S_j(\lambda_n)$ ordered lexiographically by the values of i and j. We call the $\boldsymbol{c}_{ij}$ the *basis vectors* for the color signal linear model.

Sensor Responsivities

How can we compute P sensor responsivities that filter the illuminant variation? Here we describe the calculation for the case when the number of sensors P is equal to the number of dimensions in surface linear model N. The extension to the case $P \neq N$ is straightforward.

Recall that the rows of the matrix $\boldsymbol{R}$ are the $P = N$ sensor responsivities and that the columns of $\boldsymbol{C}_m$ are the NM color signal basis vectors. To calculate the effects of the color signal basis vectors on the sensor responses, we compute

$$\boldsymbol{\Phi} = \boldsymbol{R}\boldsymbol{C}_m. \tag{13}$$

$\boldsymbol{\Phi}$ is an N by NM matrix. Each column of $\boldsymbol{\Phi}$ is the vector of sensor responses to the one of the color signal basis vectors. The first N columns of $\boldsymbol{C}_m$ span the color signals in the principal illuminant term. The corresponding columns in the matrix $\boldsymbol{\Phi}$ are the vectors of sensor responses to these basis vectors. To ensure that all of the color signals in the principal illuminant term are both detectable and distinguishable, we require that the first N columns of $\boldsymbol{\Phi}$ be independent. The remaining $N(M-1)$ columns of $\boldsymbol{C}_m$ span the color signals in the illuminant variation term. The corresponding columns in the matrix $\boldsymbol{\Phi}$ are the vectors of sensor responses to these basis vectors. To ensure that the sensors filter the illuminant variation term, we require that the remaining $N(M-1)$ columns of $\boldsymbol{\Phi}$ be zero.

If the NM columns of $\boldsymbol{C}_m$ are independent and $N_\lambda \geq N$, then it follows from elementary theorems of linear algebra that we can find a matrix $\boldsymbol{R}$ of N sensor responsivities that satisfy (13) for any matrix $\boldsymbol{\Phi}$. The $\boldsymbol{\Phi}$ described above has the special form that its last $N(M-1)$ columns are all zero. In this case we can relax the condition of independence of the columns of $\boldsymbol{C}_m$. The last $N(M-1)$ columns of $\boldsymbol{C}_m$ may be linearly dependent on each other. We only require that the first N columns of $\boldsymbol{C}_m$ be linearly independent of each other and that each of the last $N(M-1)$ columns of $\boldsymbol{C}_m$ be linearly independent of the first N columns.

Given that $\boldsymbol{C}_m$ satisfies these conditions, we can choose a $\boldsymbol{\Phi}$ as described above and solve for the sensor responsivities $\boldsymbol{R}$ using (13). When $N_\lambda > N$ there are many possible solutions. To restrict the solution, we add additional linear constraints on $\boldsymbol{R}$. These constraints restricted the sensor responsivities in each row of $\boldsymbol{R}$ to be a weighted sum of NM basis functions. In our computations, $NM = 9$ and we used basis functions that were all Gaussian functions of wavelength, each centered on a different wavelength. When we compare our computed sensor responsivities to the human cone responsivities, it is important to take the effect of our extra linear constraints into account. A very good fit to the human cone responsivities is possible using a linear combination of the nine Gaussian basis functions that constrained our sensor responsivities. The poor fit shown in Fig. 2 is not an artifact of this choice of basis functions.

It is easy to verify that any $\boldsymbol{R}$ that satisfies (13) for a $\boldsymbol{\Phi}$ as described above allows the recovery of the weights κ_{1j} of the principal illuminant term from the vector of sensor responses. Recall from Section III that the kjth entry of the matrix that maps the κ_{1j} to the r_k is given by $\sum_{n=1}^{N_\lambda} R_k(\lambda_n) \cdot C_{1j}(\lambda_n)$. Call this matrix $\boldsymbol{\Phi}_N$. Comparison with (13) shows that $\boldsymbol{\Phi}_N$ is exactly the first N columns of the N by NM matrix $\boldsymbol{\Phi}$. Since the first N columns of $\boldsymbol{\Phi}$ are linearly independent by construction, $\boldsymbol{\Phi}_N$ is invertible.

Matrix multiplication of the vector of sensor responses by the inverse $\mathbf{\Phi}_N$ gives the vector whose entries are the κ_{1j}.

Black Illuminants

When a black illuminant is added to the current illumination, there is no change in the vector of sensor responses for any surface. Because of the linearity of the relation [(11)] between the response vector $\boldsymbol{r}$ and the illuminant vector $\boldsymbol{e}$, an equivalent definition is that a black illuminant is an illuminant that generates a zero response vector when it illuminates any surface, or

$$0 = \boldsymbol{RSe} = \boldsymbol{L}_s \boldsymbol{e}. \tag{14}$$

A vector $\boldsymbol{e}$ satisfies (14) if and only if it is in the null space of the matrix $\boldsymbol{L}_s$. Whenever $N_\lambda > P$, $\boldsymbol{L}_s$ will have a nontrivial null space.

Black illuminants for an ensemble of surfaces must simultaneously be black illuminants for each surface encountered by the visual system. When the surfaces are described by an N-dimensional linear model with basis vectors s_j, $j = 1, \cdots, N$, the vector of sensor responses to any surface $\boldsymbol{s}$ when illuminated by $\boldsymbol{e}$ is the weighted sum of the responses to each of the basis vectors $\boldsymbol{s}_j$. Thus, if $\boldsymbol{e}$ is simultaneously in the null space of the N matrices $\boldsymbol{L}_{s_1}$ through $\boldsymbol{L}_{s_N}$, it will be a black illuminant with respect to all of the surfaces in the linear model. Let the matrix $\boldsymbol{L}_{\{s\}}$ be the PN by N_λ matrix formed by stacking the rows of the N matrices $\boldsymbol{L}_{s_1}$ through $\boldsymbol{L}_{s_N}$. An illuminant $\boldsymbol{e}$ in the null space of $\boldsymbol{L}_{\{s\}}$ is in the null space of $\boldsymbol{L}_{s_1}$ through $\boldsymbol{L}_{s_N}$ and is thus a black illuminant. Whenever $N_\lambda > PN$, $\boldsymbol{L}_{\{s\}}$ will have a nontrivial null space. If we take $P = 3$, $N = 3$, and $N_\lambda = 31$, then the dimension of the null space of $\boldsymbol{L}_{\{s\}}$ is at least $31 - (3 \times 3) = 22$.

In our computations, we used the singular value decomposition [29], [30] to find an orthonormal set of basis vectors for the null space of the matrix $\boldsymbol{L}_{\{s\}}$. The singular value decomposition also provides an orthonormal set of basis vectors for the subspace of visible illuminants. Once we have the two sets of orthonormal basis vectors, any illuminant can be decomposed into a black and visible term by projecting it onto each of the two subspaces.

References

[1] E. H. Land and J. J. McCann, "Lightness and retinex theory," *J. Opt. Soc. Amer.*, vol. 61, pp. 1-11, 1971.

[2] M. H. Brill, "A device performing illuminant-invariant assessment of chromatic relations," *J. Theor. Biol.*, vol. 71, pp. 473-478, 1978.

[3] G. Buchsbaum, "A spatial processor model for object color perception," *J. Franklin Inst.*, vol. 310, no. 1, pp. 1-26, 1980.

[4] L. T. Maloney and B. Wandell, "A computational model of color constancy," *J. Opt. Soc. Amer.*, vol. 1, no. 3, pp. 29-33, 1986.

[5] B. A. Wandell, "The synthesis and analysis of color images," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. PAMI-9, pp. 2-13, 1987.

[6] M. D'Zmura and P. Lennie, "Mechanisms of color constancy," *J. Opt. Soc. Amer.*, vol. 3, no. 10, pp. 1662-1672, 1986.

[7] H. Lee, "Method for computing the scene-illuminant chromaticity from specular highlights," *J. Opt. Soc. Amer.*, vol. 3, no. 10, pp. 1694-1699, 1986.

[8] G. J. Klinker, S. A. Shafer, and T. Kanade, "The measurement of highlights in color images," *Int. J. Comput. Vision*, to be published.

[9] M. H. Brill and G. West, "Chromatic adaptation and color constancy: A possible dichotomy," *Color Res. Appl.*, vol. 11, pp. 196-204, 1986.

[10] P. Sallstrom, "Color and physics: Some remarks concerning the physical aspects of human color vision," Inst. Physics Rep. 73-09, University of Stockholm, Stockholm, Sweden, 1973.

[11] L. T. Maloney, "Evaluation of linear models of surface spectral reflectance with small numbers of parameters," *J. Opt. Soc. Amer.*, vol. 3, no. 10, pp. 1673-1683, 1986.

[12] J. Cohen, "Dependency of the spectral reflectance curves of the Munsell color chips," *Psych. Sci.*, vol. 1, pp. 369-370, 1964.

[13] L. T. Maloney, "Computational approaches to color vision," Ph.D. dissertation, Stanford University, Stanford, CA, 1984.

[14] E. L. Krinov, "Surface reflectance properties of natural formations," *Nat. Res. Council Canada: Tech. Translation*, vol. TT-439, 1947.

[15] D. Nickerson, Spectrophotometric data for a collection of Munsell samples, U.S. Dep. Agriculture, Washington, D.C., 1957.

[16] 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.

[17] W. B. Cowan and N. Rowell, "A fast inexpensive spectroradiometer," Personal commun., 1988.

[18] R. M. Boynton, *Human Color Vision*. New York: Holt, Rinehart, and Winston, 1979.

[19] H. J. A. Dartnall, "The interpretation of spectral sensitivity curves," *Brit. Med. Bull.*, vol. 9, pp. 24-30, 1953.

[20] H. J. A. Dartnall and J. N. Lythgoe, "The clustering of fish visual pigments around discrete spectral positions, and its bearing on chemical structure," in *CIBA Foundation Symposium on Physiology and Experimental Psychology*, J. Knight, Ed. London: Churchill, 1965, pp. 3-21.

[21] —, "The spectral clustering of visual pigments," *Vision Res.*, vol. 5, pp. 81-100, 1965.

[22] G. H. Jacobs and J. Neitz, "Spectral positioning of mammalian cone pigments," *J. Opt. Soc. Amer.*, vol. 2, p. 23(a), 1985.

[23] G. Wyszecki, "Evaluation of metameric colors," *J. Opt. Soc. Amer.*, vol. 48, p. 451, 1958.

[24] W. S. Stiles and G. Wyszecki, "Counting metameric object colors," *J. Opt. Soc. Amer.*, vol. 52, p. 313, 1962.

[25] 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, pp. 779-784, 1977.

[26] G. West and M. H. Brill, "Necessary and sufficient conditions for von Kries chromatic adaptation to give color constancy," *J. Theor. Biol.*, vol. 15, p. 249, 1982.

[27] J. B. Cohen and W. E. Kappauf, "Metameric color stimuli, fundamental metemers, and Wyszecki's metameric blacks theory, algebra, geometry, application," *Amer. J. Psych.*, vol. 95, pp. 537-564, 1982.

[28] —, "Color mixture and fundamental metamers: Theory, algebra, geometry, application," *Amer. J. Psych.*, vol. 98, no. 2, pp. 171-259, 1985.

[29] G. H. Golub and C. F. van Loan, *Matrix Computations*. Baltimore, MD: Johns Hopkins University, 1983.

[30] D. H. Brainard and B. A. Wandell, "The color analysis package," to be published.

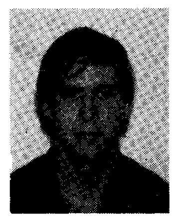

David H. Brainard is studying for the Ph.D. degree in psychology and the M.S. degree in electrical engineering at Stanford University, Stanford, CA.

His main research interests are in human and machine color vision. His research in human color vision centers on studying the dependence of color appearance on viewing context and on investigating the role of color in pattern recognition. His research in machine color vision centers on understanding how color image data can be used to identify objects under conditions of uncertain illumination.

Brian A. Wandell is a Professor of Psychology at Stanford University, Stanford, CA. His recent interests have been in studying human color perception and pattern recognition algorithms that make use of wavelength information. He is an Editor of *Vision Research*, and he shared the 1986 Troland Research Award from the U.S. National Academy of Sciences for his work in color vision.

William B. Cowan is a Research Officer in the Division of Physics at the National Research Council of Canada, Ottawa, Ont., and an Adjunct Professor in the Computer Graphics Laboratory at the University of Waterloo, Waterloo, Ont. He consults extensively on color usage in the human-computer interface and has taught short courses on color perception and color measurement for SIG-GRAPH and the Optical Society of America.