Color constancy: a method for recovering surface spectral reflectance

Laurence T. Maloney* and Brian A. Wandell

Department of Psychology, Stanford University, Stanford, California 94305

Received July 18, 1985; accepted August 9, 1985

Human and machine visual sensing is enhanced when surface properties of objects in scenes, including color, can be reliably estimated despite changes in the ambient lighting conditions. We describe a computational method for estimating surface spectral reflectance when the spectral power distribution of the ambient light is not known.

INTRODUCTION

When evening approaches, and daylight gives way to artificial light, we notice little change in the colors of objects around us. The perceptual ability that permits us to discount spectral variation in the ambient light and assign stable colors to objects is called color constancy. Much of human color-vision research focuses on the adaptational mechanisms underlying color constancy. Yet, given the kinds of information available in the initial stages of biological vision, it is not known how color constancy is even possible. Indeed, without restrictions on the range of lights and surfaces that the visual system will encounter, color constancy is not, in general, possible.[1]

In this paper we describe an algorithm for estimating the surface reflectance functions of objects in a scene with incomplete knowledge of the spectral power distribution of the ambient light. We assume that lights and surfaces present in the environment are constrained in a way that we make explicit below. An image-processing system using this algorithm can assign colors that are constant despite changes in the lighting on the scene. This capability is essential to correct color rendering in photography, in television, and in the construction of artificial visual systems for robotics. We describe how constraints on lights and surfaces in the environment make color constancy possible for a visual system and discuss the implications of the algorithm and these constraints for human color vision.

PRELIMINARY DEFINITIONS

Consider a visual sensing device consisting of a lens that focuses light from a scene onto a planar array of sensors, analogous to a retina. We introduce the following definitions and assumptions. We begin by restricting attention to a region of the scene where the spectral power distribution of the light is constant.[2] At any location in the scene, the ambient light is specified by its spectral power distribution, $E(\lambda)$, which describes the energy per second at each wavelength, λ. The ambient light is reflected from surfaces and focused onto the sensor array. The proportion of light of wavelength λ reflected from an object toward location x on the sensor array is determined by the surface spectral reflectance, $S^x(\lambda)$. The superscript x denotes the spatial position on the two-dimensional sensor array at which the object is imaged.[3] The light arriving at each location x on the sensor array is described by the function $E(\lambda)S^x(\lambda)$.

We assume that there are p distinct classes of sensors at each location x. In human vision, there are four photoreceptor classes (rods and cones), of which three (cones) are known to be active in daylight vision. We denote the relative wavelength sensitivity of the visual color sensors of the kth class as $R_k(\lambda)$

Each of the p sensors at location x records a sensor quantum catch

$$\rho_k{}^x = \int E(\lambda)S^x(\lambda)R_k(\lambda)\mathrm{d}\lambda, \qquad k = 1, 2, \ldots, p, \quad (1)$$

where the integral is taken over the entire spectrum. The information about the scene available to the visual system is contained in the p sensor quantum catches at each location x. The spectral reflectance at each location $S^x(\lambda)$ is assumed to be unknown.

Given only the sensor responses $\rho_k{}^x$, we show how to recover the surface spectral reflectances $S^x(\lambda)$ over a range of possible ambient lights $E(\lambda)$. Knowledge of $S^x(\lambda)$ permits us to compute color descriptors that are independent of the ambient light $E(\lambda)$.

PREVIOUS WORK

In their early and important work on color constancy, Helson and Judd[4] studied and formally modeled the ability of human observers to achieve this goal. Land and McCann[5] proposed a theory of color vision (the retinex theory) and a method for computing color-constant color descriptors given only the kinds of information available in the sensor responses. The retinex algorithm performs this task only under a limited set of physical conditions.[6] Buchsbaum[7] demonstrated that it is possible to compute color descriptors that are completely independent of the ambient light in an image if the average spectral reflectance of the objects in the image is known. Buchsbaum's result is the strongest that has been obtained until now. His result is useful for many applications, but it is of limited utility in visual sensing applications such as photography and satellite remote sensing in which it is not possible to know in advance the average spectral reflectance function of the objects in the image.

0740-3232/86/010029-05$02.00

Our purpose here is to improve on Buchsbaum's result. We describe how to recover surface spectral reflectance from an image without knowledge of the average spectral reflectance function.[8]

MODELS OF LIGHTS AND SURFACES REFLECTANCES

We express surface reflectance as a weighted sum of basis spectral reflectance functions $S_j(\lambda)$ as suggested by Sällström,[1] Buchsbaum,[7] and Brill,[9]

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

and term this representation a linear model of surface reflectance. The basis reflectances are fixed. They do not vary with location in the scene and are assumed known. The number of basis elements, n, is referred to as the number of degrees of freedom in the model. Knowledge of the weights $\sigma_j{}^x$ corresponding to a surface reflectance $S^x(\lambda)$ described by the finite linear model amounts to complete knowledge of $S^x(\lambda)$.

Any finite set of surface spectral reflectances can be reproduced by a linear model of this kind if n is large enough. What is surprising is that models with only a few basis reflectances provide excellent approximations to many naturally occurring spectral reflectances. Stiles *et al.*[10] suggest that spectral reflectances may be treated as band-limited functions. A collection of band-limited functions is perfectly captured by a linear model. The number of basis elements is proportional to the band limit.[11] The range of limiting frequencies that they suggest corresponds to three to five basis reflectances in Eq. (2). Buchsbaum and Gottschalk[12] demonstrate that band-limited reflectances generated using as few as three basis reflectances provide metamers to most naturally occurring spectral reflectances.

The Munsell collection includes color chips spanning a wide range of colors. Cohen[13] used a characteristic vector decomposition of the spectral reflectances of 150 Munsell color chips selected at random from a full set of 433 chips to compute the linear models with from 1 to 5 basis reflectances that best approximated the surface reflectances of the Munsell chips.[14] He found that a linear model using as few as three properly chosen basis reflectances captured 99.2% of the overall variance. He states that model reflectances generated by these three basis reflectances provided a good approximation to the spectral reflectances of the Munsell chips. Reflectances generated by these same three basis reflectances also provide good approximations to 337 surface spectral reflectances of naturally occurring objects measured by Krinov.[15]

We also represent the ambient light by a linear model,

$$E(\lambda) = \sum_{i=1}^{m} \epsilon_i E_i(\lambda), \tag{3}$$

with fixed, known basis lights $E_i(\lambda)$. It is natural to inquire how well such a model captures the range of spectral variation of natural lights such as daylight. Judd *et al.*[16] performed a characteristic vector analysis of 622 functions describing the spectral distribution of natural daylight measured over a range of weather conditions and times of day. They determined that, for practical purposes, three to four basis lights provide essentially perfect matches to the spectral distributions measured. The measurements by Judd *et al.* suggest that the number of parameters required to have an adequate linear model of the ambient light may often be small. Dixon and others have independently measured and analyzed spectral power distributions of daylight and drawn similar conclusions.[17]

REFORMULATION OF THE PROBLEM

Next we describe the consequences of assuming linear models of the ambient light and surface reflectances for the problem of color constancy. The m values of ϵ_i in Eq. (3) form a (column) vector ϵ that specifies the light $E(\lambda)$. The n values of σ_j in Eq. (2) form a (column) vector σ that specifies the surface reflectance $S^x(\lambda)$. Substituting Eqs. (2) and (3) into Eq. (1) permits us to express the relationship between the daylight surface reflectances and sensor responses by the matrix equation

$$\rho^x = \Lambda_\epsilon \sigma^x, \tag{4}$$

where ρ^x is a (column) vector formed from the quantum catches of the p sensors at location x. The matrix Λ_ϵ is p by n, and its kjth entry is of the form $\int E(\lambda)S_j(\lambda)R_k(\lambda)d\lambda$. The matrix Λ_ϵ captures the role of the light in transforming surface reflectances at each location x into sensor quantum catches.

Various limits on surface reflectance recovery are dictated by Eq. (4). We consider the limits on recovery (1) when the light on the scene is assumed to be known and (2) when the light on the scene is unknown.

In the simple case in which the ambient light and (therefore) the lighting matrix Λ_ϵ is known, we see that to recover the n weights that determine the surface reflectance we need merely solve a set of simultaneous linear equations. The recovery procedure reduces to matrix inversion when $p = n$. If p is less that n Eq. (4) is underdetermined and there is no unique solution.

If the ambient light is unknown then it easy to show that we cannot do so well: we cannot in general recover the ambient light vector ϵ or the spectral reflectances even when $p = n$. The matrix Λ_ϵ is square. For any ϵ such that Λ_ϵ is nonsingular there is a set of surface reflectances that satisfy Eq. (4). Any such choice of a light vector ϵ and corresponding surface reflectances σ^x could have produced the observed surface reflectances. No unique solution is possible without additional information concerning lights and surfaces in the scene.

Any solution method must therefore resolve the unfavorable ratio of unknown parameters to observed data points. We do so by assuming that there are more classes of sensors than degrees of freedom in surface reflectances: $p > n$. Suppose that there are $p = n + 1$ linearly independent sensors to sample the image at each location spectrally. In this case from s different spatial locations $s(n + 1)$ data values are obtained on the left-hand side of the equation. The number of unknown parameters is only sn unknowns from the different surface vectors and m unknowns from the light vector. After sampling at $s > m$ locations we finally obtain more data values than unknowns.

In principle, then, the response of the $n + 1$ sensors can contain enough information to permit exact recovery of both the lighting parameter ϵ and the surface reflectance σ^x at each location. Next we outline a method for computing the light ϵ and the n-dimensional surface reflectance vector σ^x given the $n + 1$ dimensional sensor response vector ρ^x at each location.

COMPUTATIONAL METHOD

The ambient light vector ϵ and the surface vectors σ^x contribute to the value of the sensor vectors in different ways. The ambient light vector specifies the value of the light transformation matrix, Λ_ϵ. The matrix Λ_ϵ is a linear transformation from the n-dimensional space of surface reflectances σ^x into the $n + 1$-dimensional space of sensor quantum catches ρ^x. The sensor response to any particular surface, σ^x, is the weighted sum of the n column vectors of Λ_ϵ. Consequently, the sensor responses must fall in a proper subspace of the sensor space determined by Λ_ϵ and therefore by the lighting parameter ϵ.

Figure 1 illustrates the situation when there are three classes of sensors ($p = 3$) and two degrees of freedom in the surface reflectances ($n = 2$). In the particular example shown in Fig. 1, the two-dimensional surface vectors span a plane (passing through the origin) in the three-dimensional sensor space. The light ϵ determines the plane.

We propose a two-step procedure to estimate light and surface reflectances. First, we determine the plane spanning the sensor quantum catches: knowledge of the plane in Fig. 1 permits us to recover the ambient light vector ϵ. Second, once we know the light vector ϵ, we determine the lighting matrix Λ_ϵ and achieve our goal of recovering the surface vectors simply by inverting this transformation. The exact mathematical conditions that must obtain in order for our procedure to yield the unique correct result are analyzed by Maloney.[18]

We have shown that the formal problem of estimating ambient light and surface spectral reflectance from image data may be reduced to a simple computational procedure. Our procedure is also applicable to the construction of automatic sensor systems capable of discounting fluctuations in the ambient light in unusual working environments.

In some natural scenes, the spectral composition of the ambient light varies with spatial location. The computation above can be extended in a straightforward manner to the problem of estimating and discounting a slowly varying (spatial-low-pass) ambient light. We will describe the details of how to do so elsewhere.

The recovery procedure that we developed is exact when the actual physical conditions fall within the bounds defined by the finite-dimensional models of lights and of surface reflectances. The method is readily extended to the case in which the finite-dimensional models only approximate actual lights and surface reflectances. We can no longer guarantee, for example, that the sensor quantum catches in Fig. 1 will lie exactly in the plane determined by the light. Under these circumstances, we estimate the plane that best fits the sensor quantum catches in the least-squares sense and derive an estimate of the light vector $\hat{\epsilon}$. Given $\Lambda_{\hat{\epsilon}}$, we continue by computing the best estimates of surface reflectance σ^x in the least-squares sense.[19] Small deviations from the assumptions of the models produce small errors in estimation.

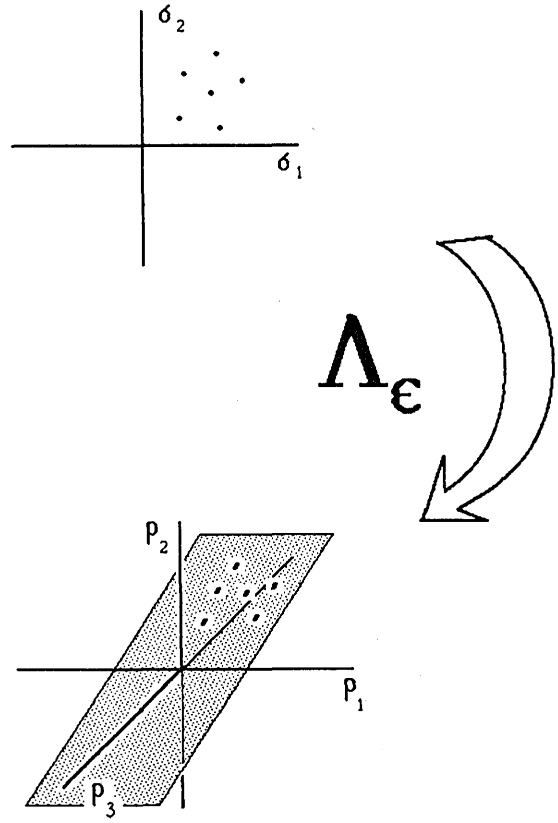

Fig. 1. Outline of the solution method in the case when there are three classes of sensors ($p = 3$) and two degrees of freedom in the model of surface reflectances ($n = 2$). The sensor quantum catches lie on a plane through the origin in the three-dimensional space of sensor quantum catches below. Knowledge of the plane determines the light ϵ and the matrix Λ_ϵ. The matrix Λ_ϵ is inverted to recover the surface reflectances σ^x above.

By increasing the dimensionality of the finite-dimensional approximation, exact solutions may be approached to within any desired degree of precision. A trichromatic visual system can approximate color constancy when reflectances in the visual environment are predominantly captured by a linear model with two degrees of freedom.

IMPLICATIONS

Our calculation has several implications for human color vision. First, it suggests that perfect color constancy is possible only for ranges of lights and reflectances that can be described by a small number of parameters. The human visual system is known to have better color constancy over some ranges of lights than others. Our formulation provides a framework that may be used to determine, experimentally, the range of ambient lights and surfaces over which human color constancy succeeds and over which it fails.

We can test whether a linear model characterizes the range of lights permitting essentially perfect color constancy in human vision as follows. Suppose that the color appearance of a set of surfaces is preserved when measured with two ambient lights, say $E(\lambda)$ and $E'(\lambda)$. If the lights for which human color constancy succeeds form a linear model, then color appearance should also be preserved when the surfaces are viewed under weighted mixtures of the ambient lights $E(\lambda)$ and $E'(\lambda)$.

Second, our results show how the number of classes of photoreceptors active in color vision limits the number of degrees of freedom in the surface reflectances that can be recovered. With three classes of photoreceptors, we can exactly recover surface reflectances drawn from a fixed model of surface reflectance with at most two degrees of freedom. Additional degrees of freedom in the surface reflectance will, in general, introduce error into the estimates of surface reflectance obtained, precluding perfect color constancy.

Third, the parameter counting argument indicates that a minimum number of distinct surface reflectances must be present in the scene to permit recovery. In the illustration of the solution above, at least two distinct surface reflectances are needed to specify the plane (which must pass through the origin) that determines the light. In general, at least $p - 1$ distinct surfaces are needed in order to determine uniquely the light vector ϵ. In the presence of small deviations from the linear models of light and surface reflectances, an increase in the number of distinct surface reflectances will, in general, improve the estimate of the light and the corresponding surface reflectance estimates. Our analysis suggests that color constancy should improve with the number of distinct surfaces in a scene.[20]

Fourth, our algorithm may explain the reduced spatial sampling of the short-wavelength receptors in human vision.[21] Note that incorporating an additional sensor class reduces the spatial sampling density within any single type of sensor class. This creates a conflict between two goals of the visual system. Better color correction for the ambient light can be obtained by including more sensor classes. But including additional sensor classes reduces the spatial sampling density within individual sensor classes. A reasonable trade-off between these two goals may be obtained by the following observation.

Once the light vector ϵ has been estimated, we may calculate the inverse of the lighting matrix Λ_ϵ. Once this matrix is known, the $n + 1$ quantum catches at each location are redundant: only n sensor values are needed to compute the value of σ^x. It follows that if the ambient light varies slowly across the scene, the $n + 1$th sensor class need not be present at as high a sampling density as the other sensor classes.

In a three-sensor system there is no need to place the third sensor with as high a spatial sampling density as the first two sensors. It follows that reducing the retinal space occupied by the short-wavelength sensor class permits the system to obtain a higher spatial resolution of the estimated surface spectral reflectance function with virtually no deterioration in its ability to correct for variation in the spectral power distribution of the ambient light.

ACKNOWLEDGMENTS

This research was supported by grant no. 2 RO1 EY03164 from the National Eye Institute and NASA grant NCC-2-44. We thank M. Pavel, R. N. Shepard, and D. Varner for their advice and suggestions. Stanford University has applied for a patent directed to the matter described herein.

* Lawrence T. Maloney is now at the Human Performance Center, University of Michigan, 330 Packard Road, Ann Arbor, Michigan 48104.

REFERENCES AND NOTES

1. See P. Sällström, "Colour and physics: some remarks concerning the physical aspects of human colour vision," Institute of Physics Rep. 73-09 (University of Stockholm, Stockholm, 1973).
2. The computational method that we develop requires only that the ambient light be approximately constant over small local patches of the image. The method is easier to explain if we restrict attention to a region of the image across which the ambient light does not change.
3. In general the surface reflectance function may depend on the geometry of the scene, the angle of incidence of the light on the surface, and the angle between the surface and the line of sight. We are concerned here with the analysis of a single image drawn from a scene with fixed geometric relations among objects, light sources, and the visual sensor array. $S^x(\lambda)$ refers to the proportion of light returned from the object toward the sensor array within that fixed geometrical framework.
4. D. B. Judd, "Hue saturation and lightness of surface colors with chromatic illumination," J. Opt. Soc. Am. **30**, 2 (1940); 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. **23**, 439 (1938).
5. E. H. Land and J. J. McCann, "Lightness and retinex theory," J. Opt. Soc. Am. **61**, 1 (1971); E. H. Land, "Recent advances in retinex theory and some implications for cortical computations: color vision and the natural image," Proc. Nat. Acad. Sci. U.S. **80**, 5163 (1983); E. H. Land, D. H. Hubel, M. Livingston, S. Perry, and M. Burns, "Colour-generating interactions across the corpus callosum," Nature **303**, 616 (1983).
6. D. Brainard and B. Wandell, "An analysis of the retinex theory of color vision," Stanford Applied Psychology Lab. Tech. Rep. 1985-04 (Stanford University, Stanford, Calif., 1985).
7. G. Buchsbaum, "A spatial processor model for object colour perception," J. Franklin Inst. **310**, 1 (1980).
8. It is not possible to recover $E(\lambda)$ better than to within a multiplicative constant given only the sensor quantum catches. If, for example, the intensity of the light is doubled to $2E$ but all reflectances are halved to $\frac{1}{2}S^x(\lambda)$, it is easy to verify that the sensor quantum catches in Eq. (1) are unchanged. When we speak of recovering the ambient light and surface reflectances, we mean recovery up to this unknown mutiplicative constant.
9. M. H. Brill, "A device performing illuminant-invariant assessment of chromatic relations," J. Theor. Biol. **71**, 473 (1978).
10. W. S. Stiles, G. Wyszecki, and N. Ohta, "Counting metameric object-color stimuli using frequency-limited spectral reflectance functions," J. Opt. Soc. Am. **67**, 779 (1977).
11. For a discussion of band-limited functions, see R. Bracewell, *The Fourier Transform and Its Application*, 2nd ed. (McGraw-Hill, New York, 1978), Chap. 10.
12. G. Buchsbaum and A. Gottschalk, "Chromaticity coordinates of frequency-limited functions," J. Opt. Soc. Am. A **1**, 885 (1984).
13. J. Cohen, "Dependency of the spectral reflectance curves of the Munsell color chips," Psychonomic Sci. **1**, 369 (1964).
14. See K. V. Mardia, J. T. Kent, and J. M. Bibby, *Multivariate Analysis* (Academic, London, 1979), Chap. 8, for a discussion of characteristic vector analysis (also known as principal-components analysis or the Karhunen–Loève decomposition).
15. E. Krinov, *Spectral Reflectance Properties of Natural Formations*, Technical translation TT-439 (National Research Council of Canada, Ottawa, 1947); details of the fit of Cohen's characteristic vectors to the Munsell surface reflectances are given in L. Maloney, "Computational approaches to color constancy," Stanford Applied Psychology Lab. Tech. Rep. 1985-01 (Stanford University, Stanford, Calif., 1985).
16. D. B. Judd, D. L. MacAdam, and G. Wyszecki, "Spectral distribution of typical daylight as a function of correlated color temperature," J. Opt. Soc. Am. **54**, 1031 (1964).
17. E. R. Dixon, "Spectral distribution of Australian daylight," J. Opt. Soc. Am. **68**, 437 (1978); G. T. Winch, M. C. Boshoff, C. J. Kok, and A. G. du Toit, "Spectroradiometric and colorimetric characteristics of daylight in the southern hemisphere: Pre-

toria, South Africa," J. Opt. Soc. Am. **56,** 456 (1966); S. R. Das and V. D. P. Sastri, "Spectral distribution and color of tropical daylight," J. Opt. Soc. Am. **55,** 319 (1965); V. D. P. Sastri and S. R. Das, "Spectral distribution and color of north sky at Delhi," J. Opt. Soc. Am. **56,** 829 (1966); "Typical spectral distributions and color for tropical daylight," J. Opt. Soc. Am. **58,** 391 (1968).

18. L. Maloney, "Computational approaches to color constancy," Stanford Applied Psychology Lab. Tech. Rep. 1985-01 (Stanford University, Stanford, Calif., 1985).
19. See Ref. 18, Chap. 4, for details.
20. R. B. MacLeod, "An experimental investigation of brightness constancy," Arch. Psychol. **135,** 1 (1932), reports that human brightness constancy does.
21. E. N. Willmer and W. D. Wright, "Colour sensitivity of the fovea centralis," Nature **156,** 119 (1945); D. R. Williams, D. I. A. MacLeod, and M. Hayhoe, "Punctate sensitivity of the blue-sensitive mechanism," Vision Res. **21,** 1357 (1981); F. M. de Monasterio, S. J. Schein, and E. P. McCrane, "Staining of blue-sensitive cones of the macaque retina by a fluorescent dye," Science **213,** 1278 (1981).