importance in children's lives need to be questioned and thoroughly investigated.

Notes

1. P.C. Glick, Fifty years of family demography: A record of social change, *Journal of Marriage and the Family, 50*, 861–873 (1988); J.A. Seltzer and S.M. Bianchi, Children's contact with absent parents, *Journal of Marriage and the Family, 50*, 663–677 (1988).

2. P. Bronstein and C.P. Cowan, Eds., *Fatherhood Today: Men's Changing Role in the Family* (Wiley, New York, 1988); M.E. Lamb, Ed., *The Role of the Father in Child Development*, rev. ed. (Wiley, New York, 1981); M.E. Lamb, Ed., *The Father's Role: Cross-Cultural Perspectives* (Wiley, New York, 1987).

3. P.J. Caplan, *Don't Blame Mother: Mending the Mother-Daughter Relationship* (Harper & Row, New York, 1989); P.J. Caplan and I. Hall-McCorquodale, Mother-blaming in major clinical journals, *American Journal of Orthopsychiatry, 55*, 345–353 (1985).

4. V. Phares and B.E. Compas, The role of fathers in child and adolescent psychopathology: Make room for daddy, *Psychological Bulletin, 111*, 387–412 (1992).

5. V. Phares, Where's Poppa?: The relative lack of attention to the role of fathers in child and adolescent psychopathology, *American Psychologist, 47*, 656–664 (1992).

6. K. Dodge, Developmental psychopathology in children of depressed mothers, *Developmental Psychology, 26*, 3–6 (1990); G. Downey and J.C. Coyne, Children of depressed parents: An integrative review, *Psychological Bulletin, 108*, 50–76 (1990).

7. T. Jacob, Ed., *Family Interaction and Psychopathology: Theories, Methods, and Findings* (Plenum Press, New York, 1987); R.D. Parke, K.B. MacDonald, A. Beitel, and N. Bhavnagri, The role of the family in the development of peer relationships, in *Social Learning Systems: Approaches to Marriage and the Family*, R. Peters and R.J. McMahan, Eds. (Brunner-Mazel, New York, 1988).

8. L.L. Humphrey, Observed family interactions among subtypes of eating disorders using structural analysis of social behavior, *Journal of Consulting and Clinical Psychology, 57*, 206–214 (1989); T. Jacob, G.L. Krahn, and K. Leonard, Parent-child interactions in families with alcoholic fathers, *Journal of Consulting and Clinical Psychology, 59*, 176–181 (1991).

Color Constancy: From Physics to Appearance

David H. Brainard, Brian A. Wandell, and Eduardo-Jose Chichilnisky

Vision provides information about objects. Computing information about objects is difficult, however, because there is no simple mapping between an object's intrinsic properties and the corresponding retinal image. An object's retinal image confounds the object's properties and extrinsic factors such as the object's location, orientation, and illumination. A visual representation that provides reliable information about objects must compensate for these extrinsic factors.

David H. Brainard is an Assistant Professor of Psychology at the University of California's Santa Barbara campus. He is interested in formulating and testing computational models of human vision. **Brian A. Wandell** has taught at Stanford University, where he is Professor of Psychology, since 1979. His primary research is on human and computer vision. **Eduardo-Jose Chichilnisky** is a graduate student in the neurosciences program at Stanford University. He studies the neural mechanisms involved in chromatic adaptation and color appearance. Address correspondence to David H. Brainard, Department of Psychology, University of California, Santa Barbara, CA 93106.

In the case of object color, the retinal image depends not only on the object's intrinsic surface reflectance but also on the properties of the illumination. The ability of the visual system to maintain an object's color appearance across variations in illumination is called *color constancy*. Color constancy is an example of a larger class of perceptual constancies (e.g., size constancy and shape constancy) that together allow people to perceive a stable physical world. A detailed characterization of color constancy may provide insights that generalize to other perceptual systems.

FORMATION OF THE RETINAL IMAGE

To what extent can an ideal system be color constant? To answer this question, we analyze the physics of retinal image formation. The top of Figure 1 shows an illuminant reflected from a surface toward an observer. The illuminant's spectral power distribution, $E(\lambda)$, specifies the amount of power in the illuminant at each wavelength. The surface's spectral reflectance function, $S(\lambda)$, specifies the fraction of incident power reflected at each wavelength. The spectral power distribution of the light arriving at the observer's cornea, which we call the color signal, is $C(\lambda) = E(\lambda)\,S(\lambda)$.

The retina contains three distinct classes of light-sensitive photoreceptors: the long-wavelength-sensitive (L), middle-wavelength-sensitive (M), and short-wavelength-sensitive (S) cones. The classes are distinguished by the type of photopigment molecules they contain. The amounts of light absorbed by the three types of photopigment (the cone *quantum catches*) provide the information available to the visual system about the color signal.

Figure 1 illustrates why it is difficult to recover illuminant and sur-

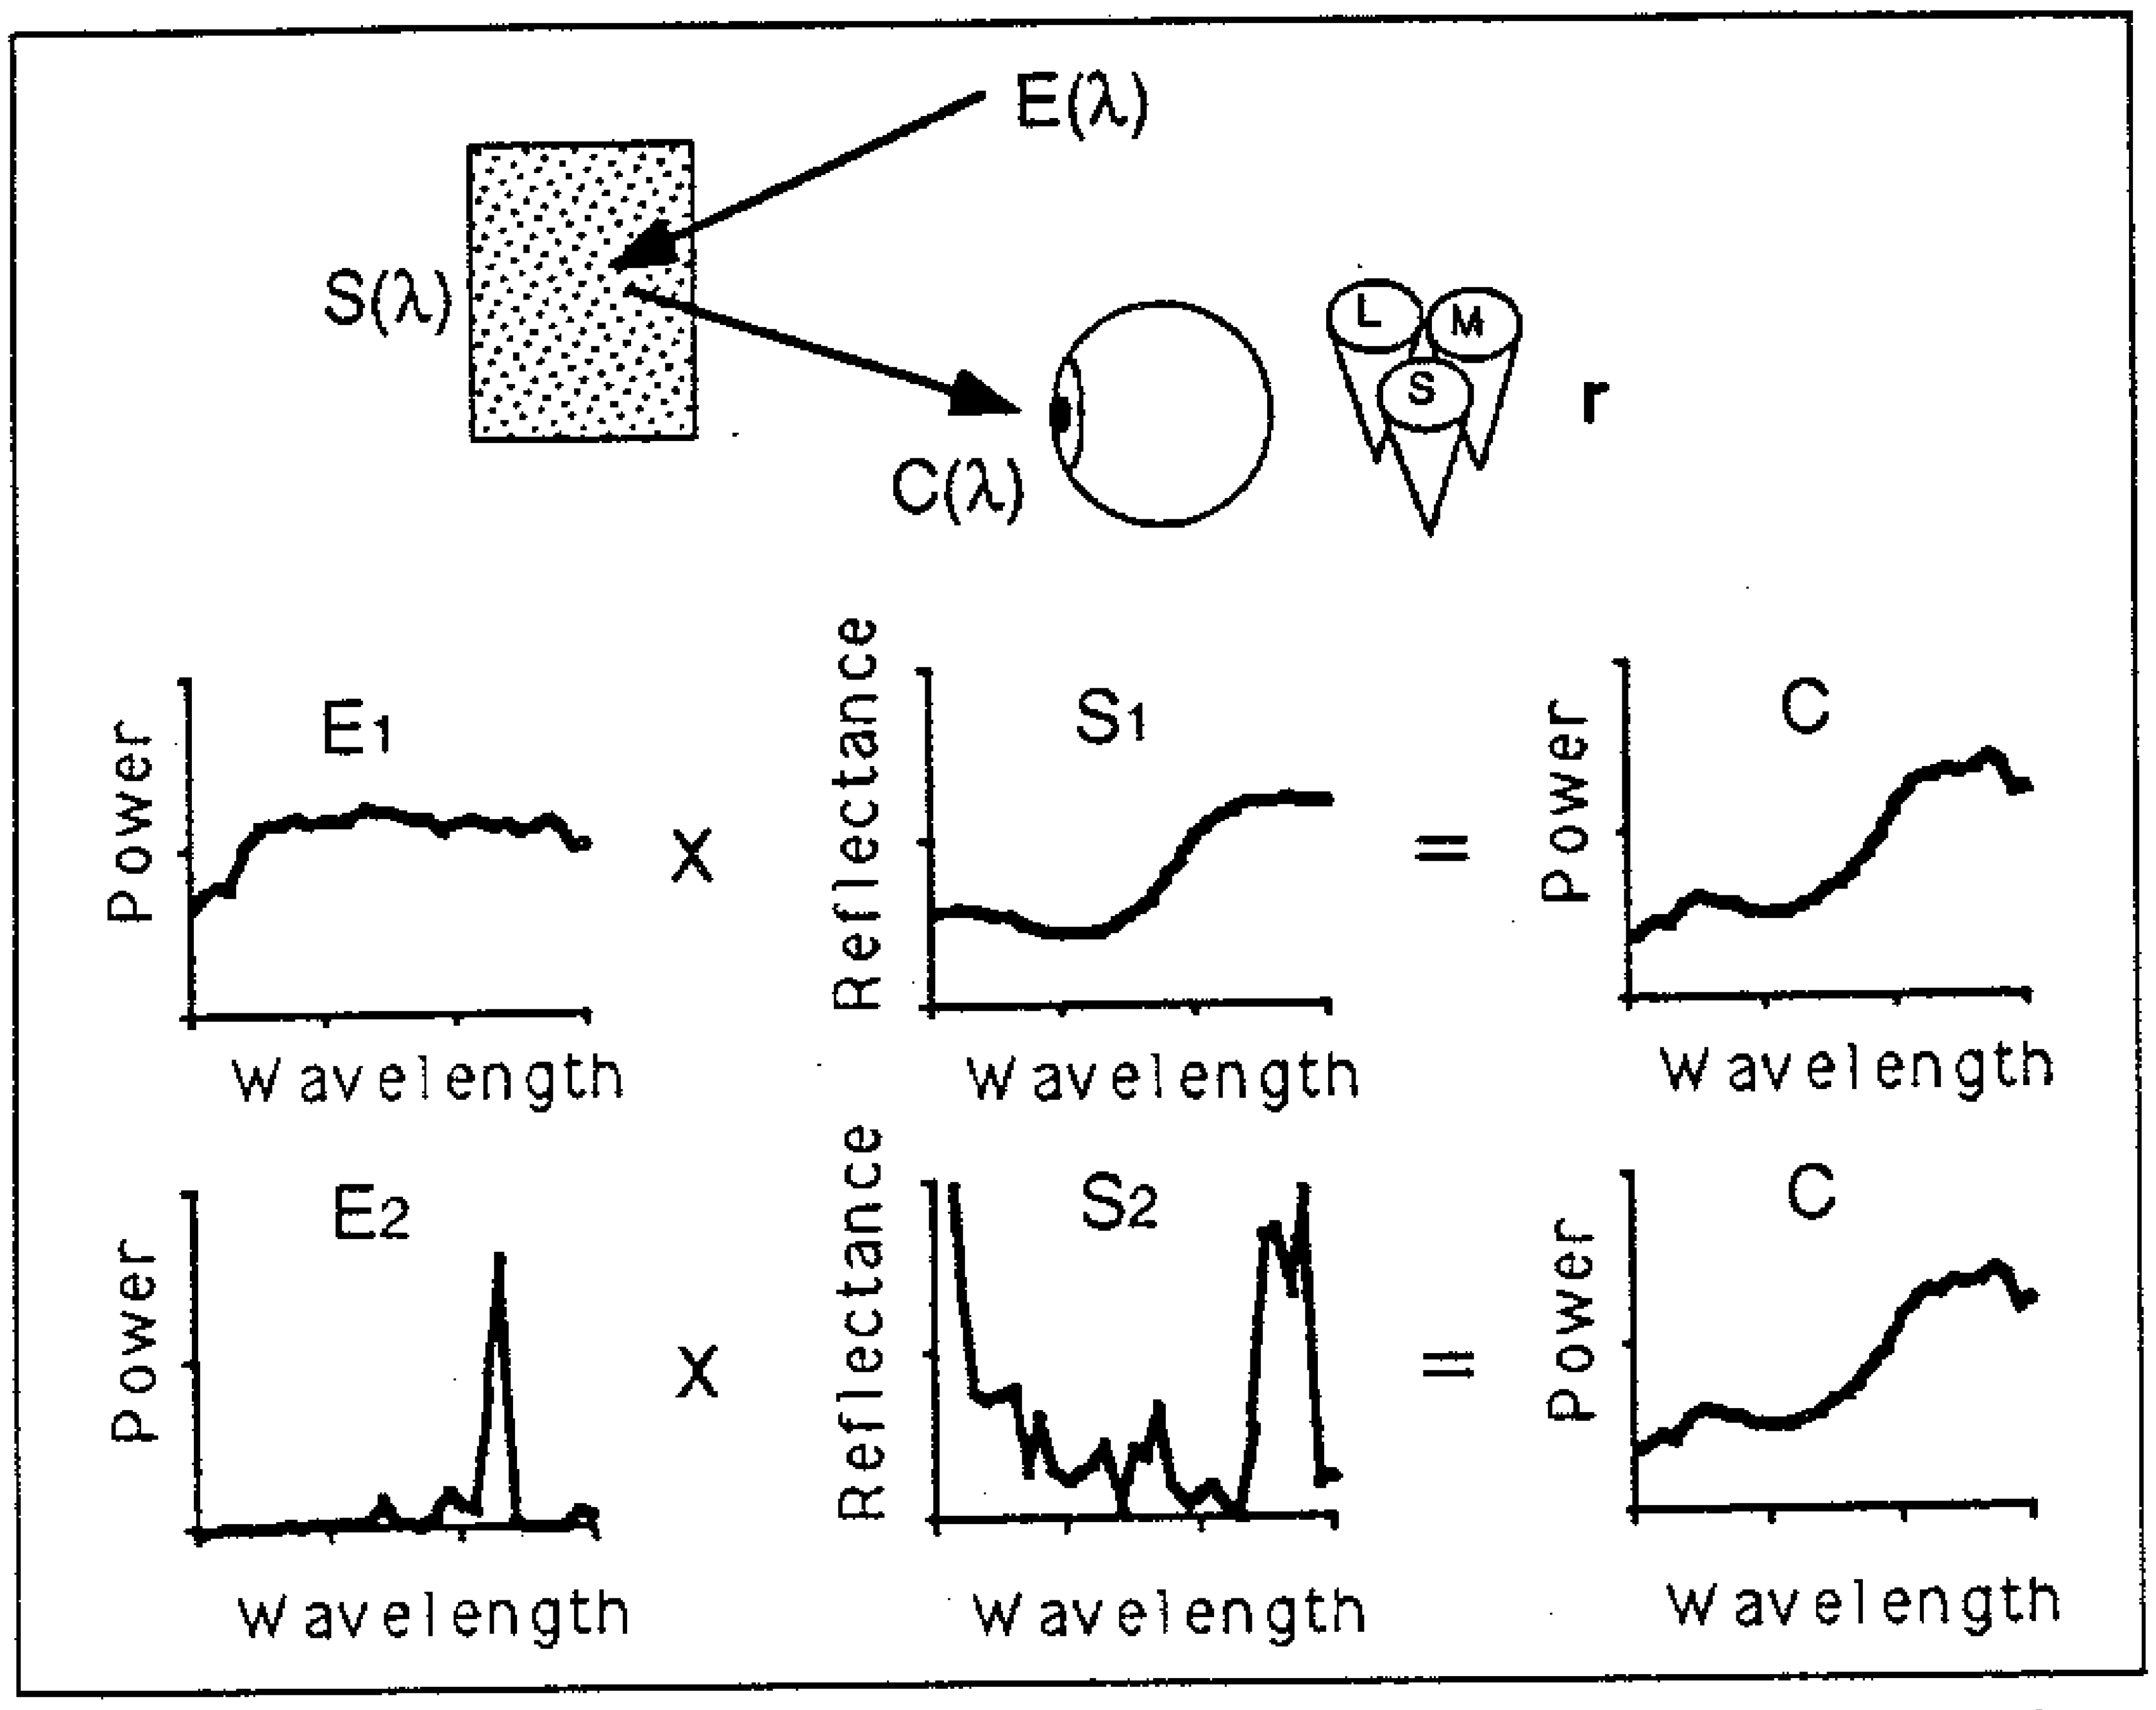

Fig. 1. The physics of image formation. The illuminant, $E(\lambda)$, is reflected from a surface, $S(\lambda)$, to form the color signal, $C(\lambda)$. The information available to the observer about the color signal is the triplet of quantum catches for the three classes of cones. We use the symbol **r** to represent this triplet. The bottom of the figure shows two illuminant–surface pairs that result in the same color signal.

face properties from the retinal image. First, consider the information about a surface available in the color signal at a single image location. In general, surface recovery is not possible: The color signal confounds the illuminant's spectral power distribution and the surface's reflectance function. Any color signal $C(\lambda)$ can be factored into an arbitrary surface reflectance function, $S(\lambda)$, and a compatible illuminant, $E(\lambda) = C(\lambda)/S(\lambda)$.

The ambiguity is shown concretely at the bottom of Figure 1, where we depict two illuminant–surface pairs that result in the same color signal. Even an ideal visual system cannot determine with certainty which pair is actually present. To resolve the ambiguity, the visual system must assume that some prior constraints govern which illuminants and surfaces are likely to occur in natural images. It may be the case that illuminant spectral power distributions similar to $E_2(\lambda)$ or surface reflectance functions similar to $S_2(\lambda)$ are rare. In this case, the visual system would do well to factor the depicted color signal into illuminant $E_1(\lambda)$ and surface $S_1(\lambda)$ rather than into illuminant $E_2(\lambda)$ and surface $S_2(\lambda)$.

LINEAR MODELS OF ILLUMINANTS AND SURFACES

The goal of a growing body of theoretical research is to identify a plausible set of prior constraints that are sufficient to allow the visual system to identify surface reflectance from the cone quantum catches.[1] The idea that unifies this research is that small-dimensional *linear models* may be used to describe which illuminant and surface functions are likely to occur in natural images. Linear models approximate spectral data as weighted sums of a small number of fixed *basis functions*.

The number of basis functions used in a linear model is called the *dimension* of the model. This nomenclature emphasizes the fact that basis functions may be interpreted as describing the dimensions along which spectra may vary. The insert in Figure 2 shows the basis functions for a three-dimensional linear model for surfaces. The first basis function reflects fairly evenly across the visible spectrum. By varying the weight assigned to this basis function, we can capture variation in overall reflectance from one surface to another. The second basis function reflects positively at the long-wavelength end of the spectrum and negatively at the short-wavelength end. By assigning a positive weight to this basis function, we capture the fact that some surfaces (e.g., red ones) reflect best at longer wavelengths. By assigning a negative weight to this basis function, on the other hand, we can capture the fact that some surfaces (e.g., green ones) reflect best at the middle and short wavelengths. The third basis func-

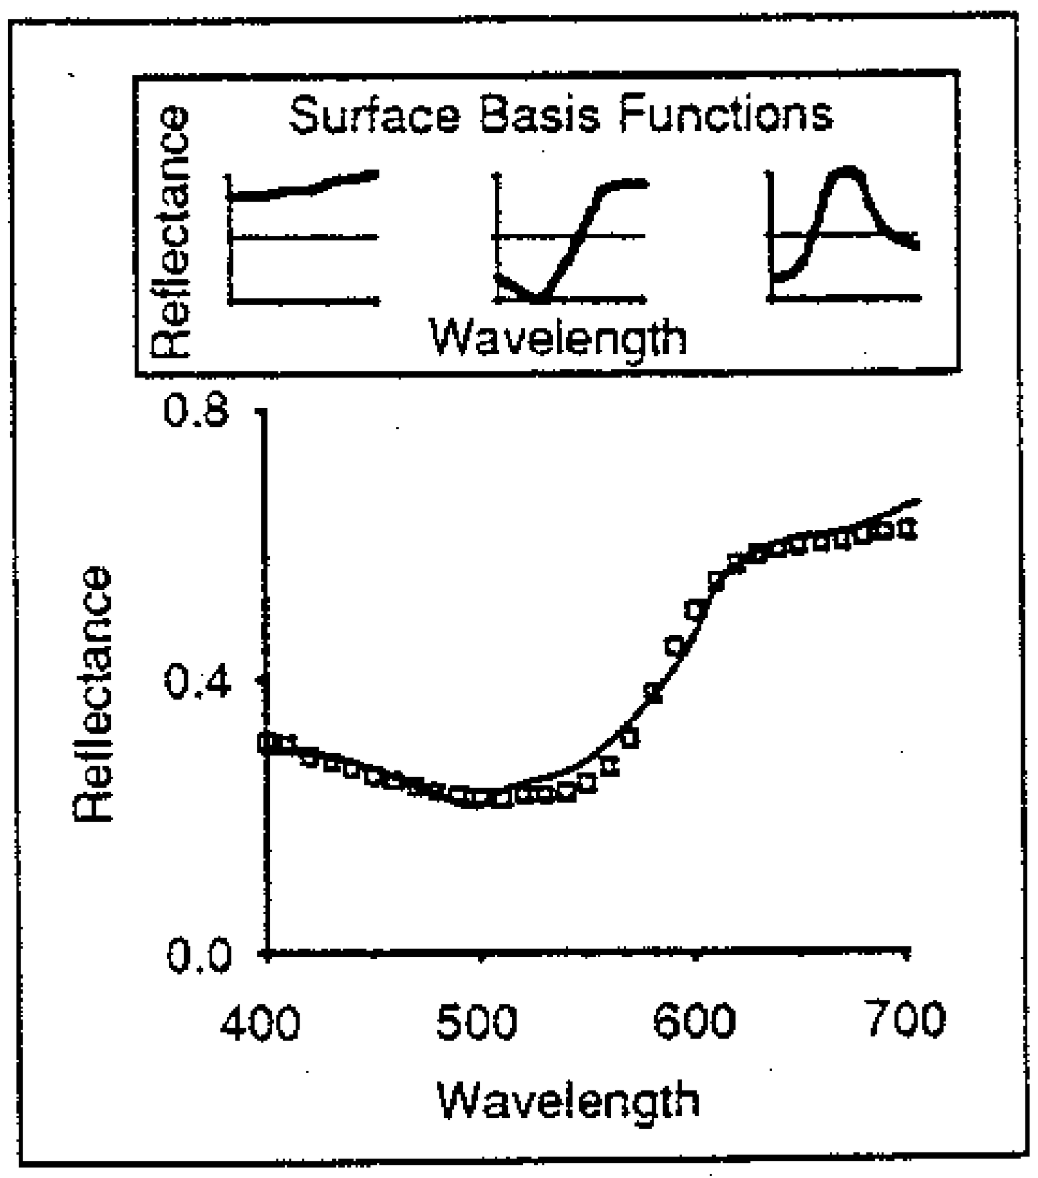

Fig. 2. Linear model approximations. The main plot shows the fit of a linear model approximation (solid line) to a surface reflectance function (open squares). The inset shows the three basis functions of the linear model. These are normalized versions of the functions reported by Cohen.[2] The approximation was obtained by combining the three basis functions with weights of 0.53, 0.14, and −0.06, respectively.

tion, which reflects positively in the middle region of the spectrum and negatively at either end, shows another dimension along which surface reflectances within the model can vary.

The basis functions shown in Figure 2 were obtained by performing a principal components analysis of the reflectance functions of a large set of colored papers.[2] Similar analyses have been performed for spectral power distributions of measured daylight.[3] The results indicate that linear models with a small number of basis functions provide an excellent description of naturally occurring spectra. Figure 2 shows a three-dimensional approximation to a typical surface reflectance function.

Linear models provide a formal description of prior constraints on illuminant and surface functions. Rather than being free to vary arbitrarily at each wavelength in the visible spectrum, functions constrained by linear models can vary only along the dimensions represented by the basis functions. Once the number of dimensions and basis function spectra for a linear model have been determined, an individual spectrum within the model is specified completely by the weights required to form it.

ILLUMINANT AND SURFACE RECOVERY ALGORITHMS

If we make the assumption that linear models describe the illuminant and surface spectral functions, it is possible to estimate these functions from the cone quantum catches. As an example, we describe one estimation algorithm, which we refer to as the *subspace method*.[4] In addition to the linear model constraint, this method is based on the assumption that a collection of surfaces is lit by a common illuminant.

The subspace method operates in two stages. First, it estimates the spectral power distribution of the common illuminant. Then it uses this estimate to deduce the surface reflectance function at each image location. The method is easiest to understand when the surface reflectance functions are constrained by a two-dimensional linear model. In this case, any individual surface may be represented by a point in a two-dimensional space. The coordinates of the point specify the weights required to reconstruct the surface reflectance function from the basis functions. The top of Figure 3 shows such a two-dimensional representation for a hypothetical collection of surfaces.

When a surface is illuminated, the resulting color signal is coded by the visual system as a triplet of quantum catches, one for each class of cone. We can represent these quantum catches using a three-dimensional quantum catch space. The coordinates of a point in this space specify the L, M, and S cone quantum catches. The two plots at the bottom of Figure 3 show the quantum catches from the entire collection under two hypothetical illuminants, E_1 and E_2. In both of these plots, the ensemble of quantum catches from the entire collection of surfaces lies within a plane. But the particular plane that the surfaces are mapped into depends on the illuminant: The planes in the two plots are different.

The subspace method operates by identifying the plane that best fits the quantum catches and using this plane to identify the illuminant. For a fixed illuminant, the plane is invariant with respect to the choice of the surfaces in the image. Informally, one can think of the identification procedure as associating an illuminant with each possible plane in the quantum catch space. Once the illuminant has been identified, the algorithm processes the quantum catches at each location to identify the corresponding surface.

Why does the method work? The

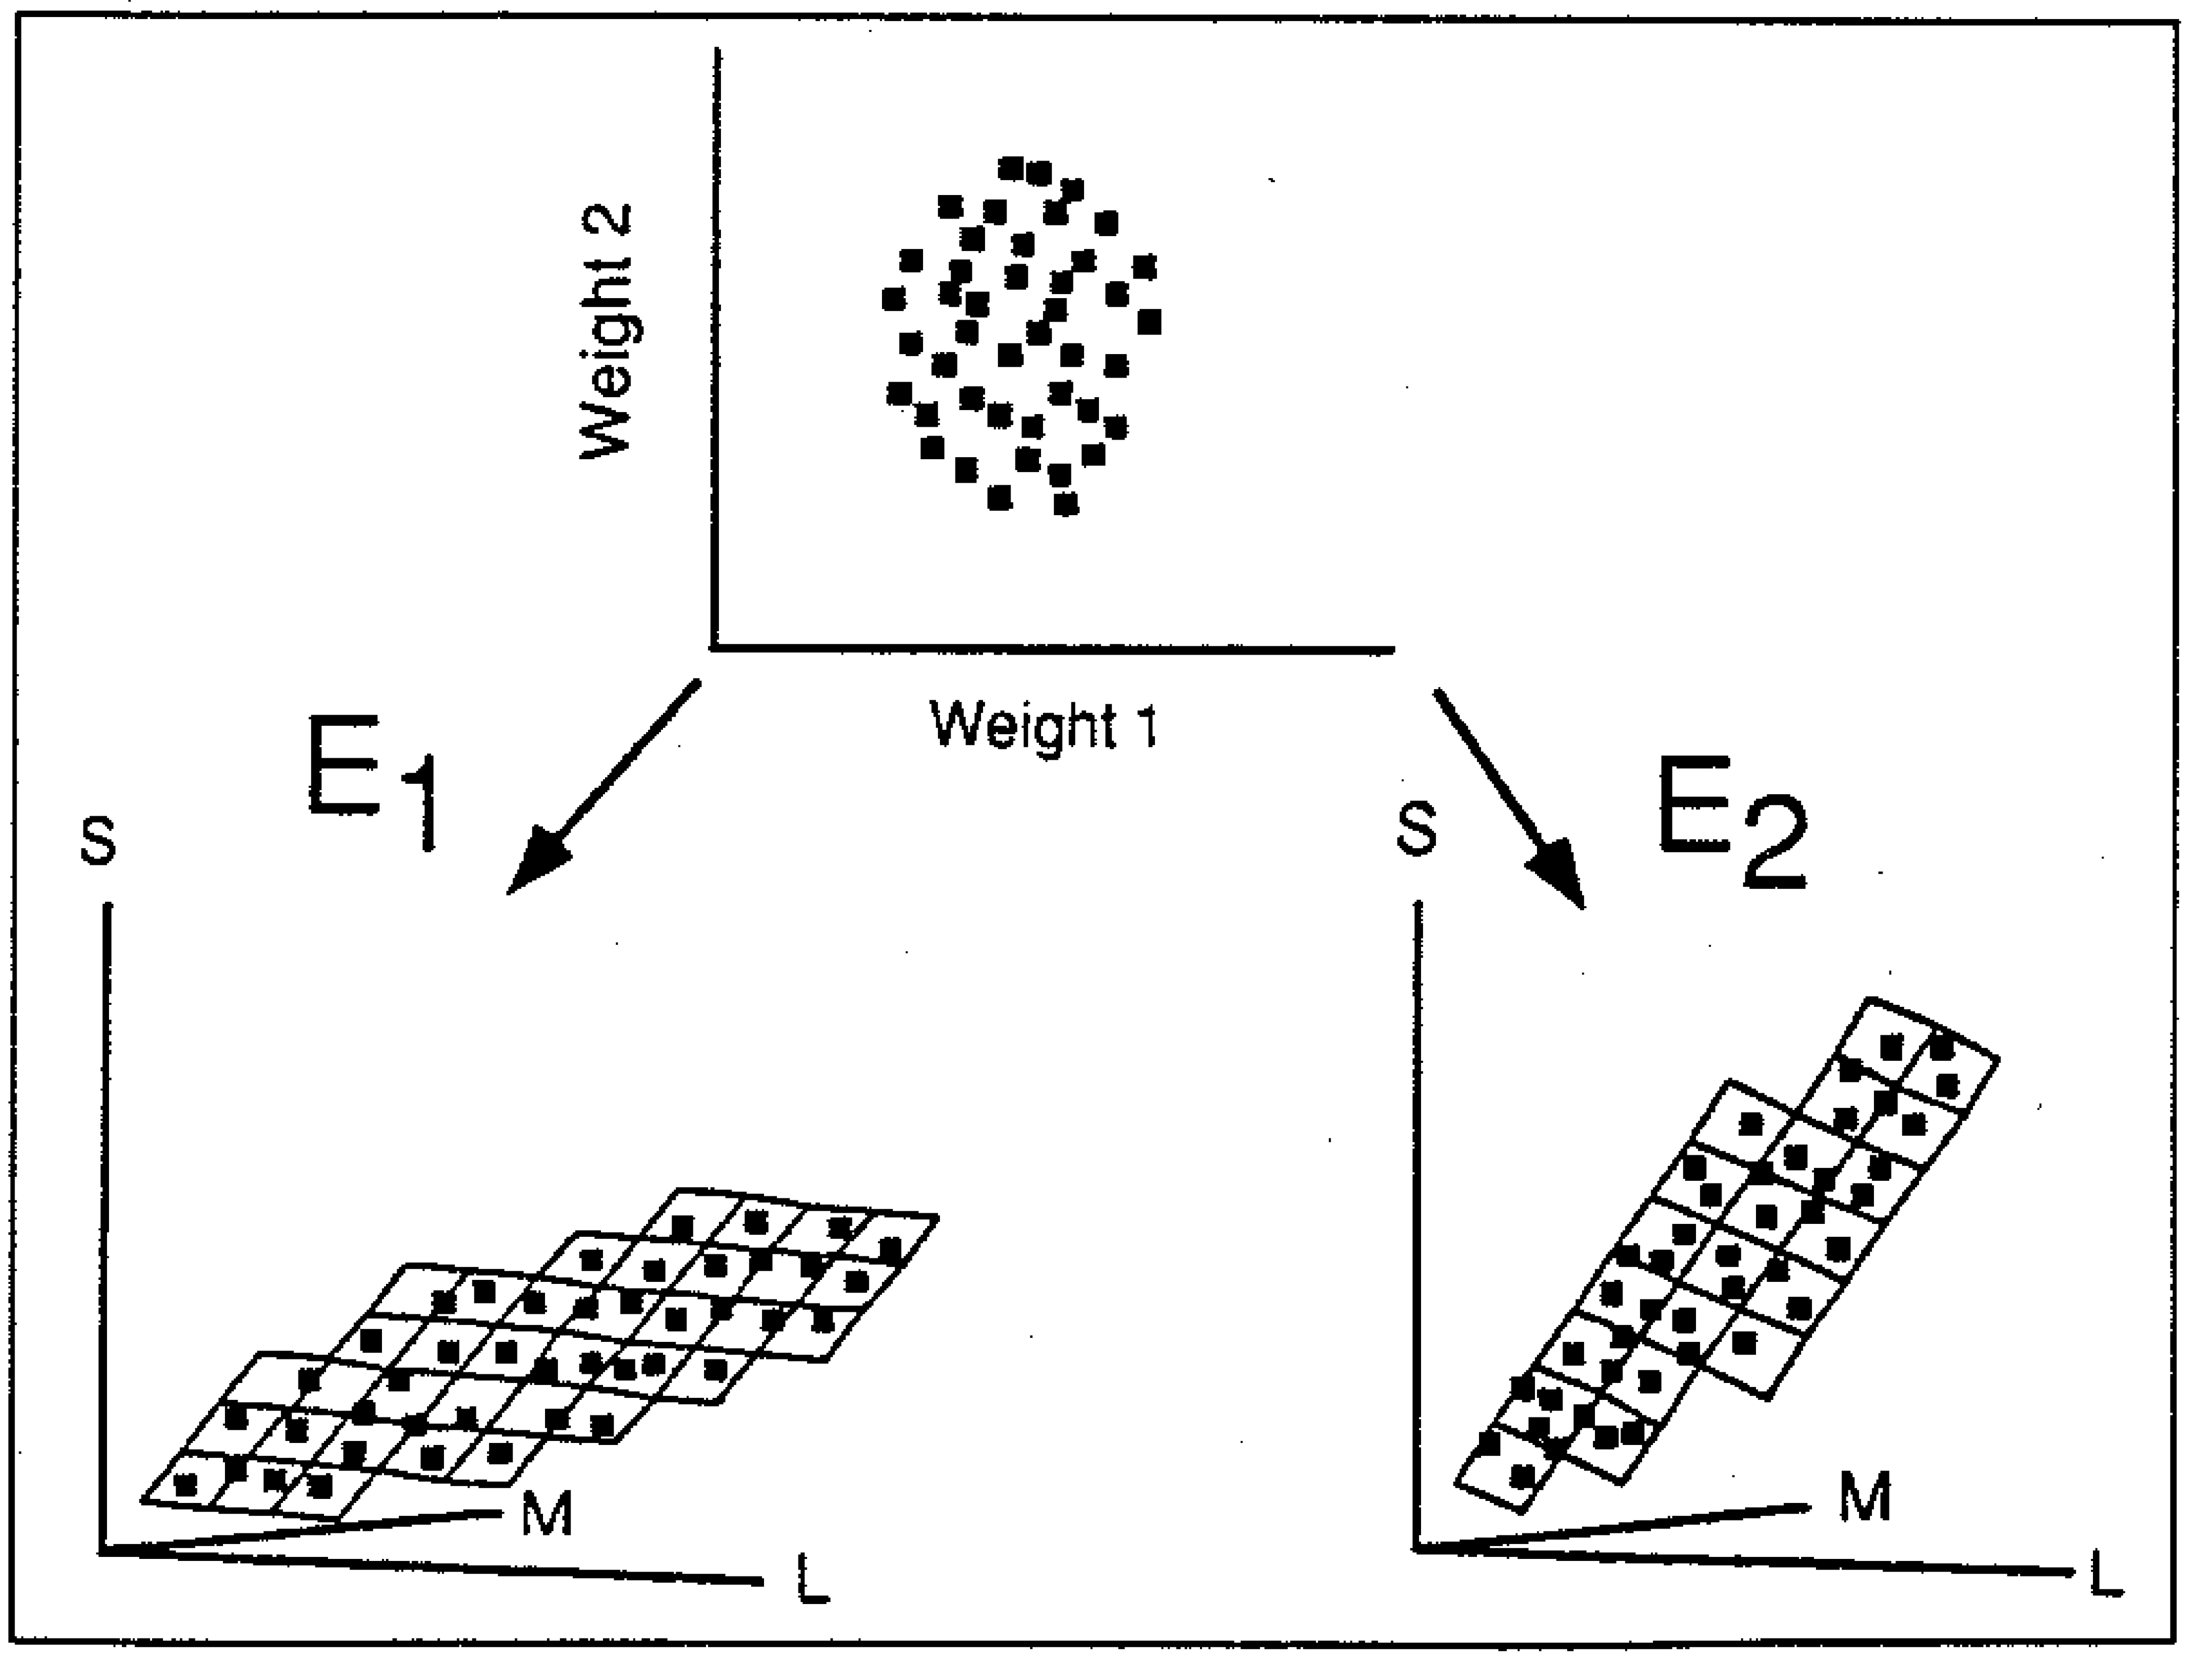

Fig. 3. The idea underlying the subspace method. The plot at the top represents a collection of surfaces within a two-dimensional linear model. The two plots at the bottom represent the quantum catches from the collection of surfaces under two different illuminants. The figure is illustrative and does not represent the results of actual calculations. After Maloney and Wandell.[4]

key observation is that each illuminant maps the collection of surfaces into a plane in the quantum catch space. This follows because there is considerable structure in the relation between illuminants, surfaces, and cone quantum catches. In particular, varying either the illuminant spectral power distribution or the surface reflectance has a linear effect on the cone quantum catches, a property called *bilinearity* (see Fig. 4). Because of bilinearity, the mapping between surfaces and quantum catches is linear when the illuminant is held fixed. Since linear transformations always map planes into planes, the two-dimensional collection of surfaces must end up as a plane in the quantum catch space. Bilinearity is also central to a rigorous analysis of when it is guaranteed that distinct illuminants map the collection of surfaces to distinct planes.[5]

For a trichromatic visual system, our analysis applies rigorously when the surfaces are constrained to lie within a two-dimensional linear model. But the principle of identifying the illuminant by examining how the quantum catch data cluster is quite general. When the data do not lie exactly in a plane, it may be possible to use other measures of how the data cluster to estimate the illuminant. Early color constancy algorithms estimated the illuminant on the basis of the mean quantum catch from a collection of surfaces, but these methods are not robust with respect to the choice of surfaces in the image.[6] The subspace method demonstrates that important information about the illuminant is carried by the higher order statistics of the ensemble of quantum catches (e.g., the best-fitting plane). We believe that this basic insight will lead to algorithms that operate effectively even when the surfaces are not strictly constrained to a two-dimensional linear model. Similar insight has already proved useful for solving other computational vision problems.[7]

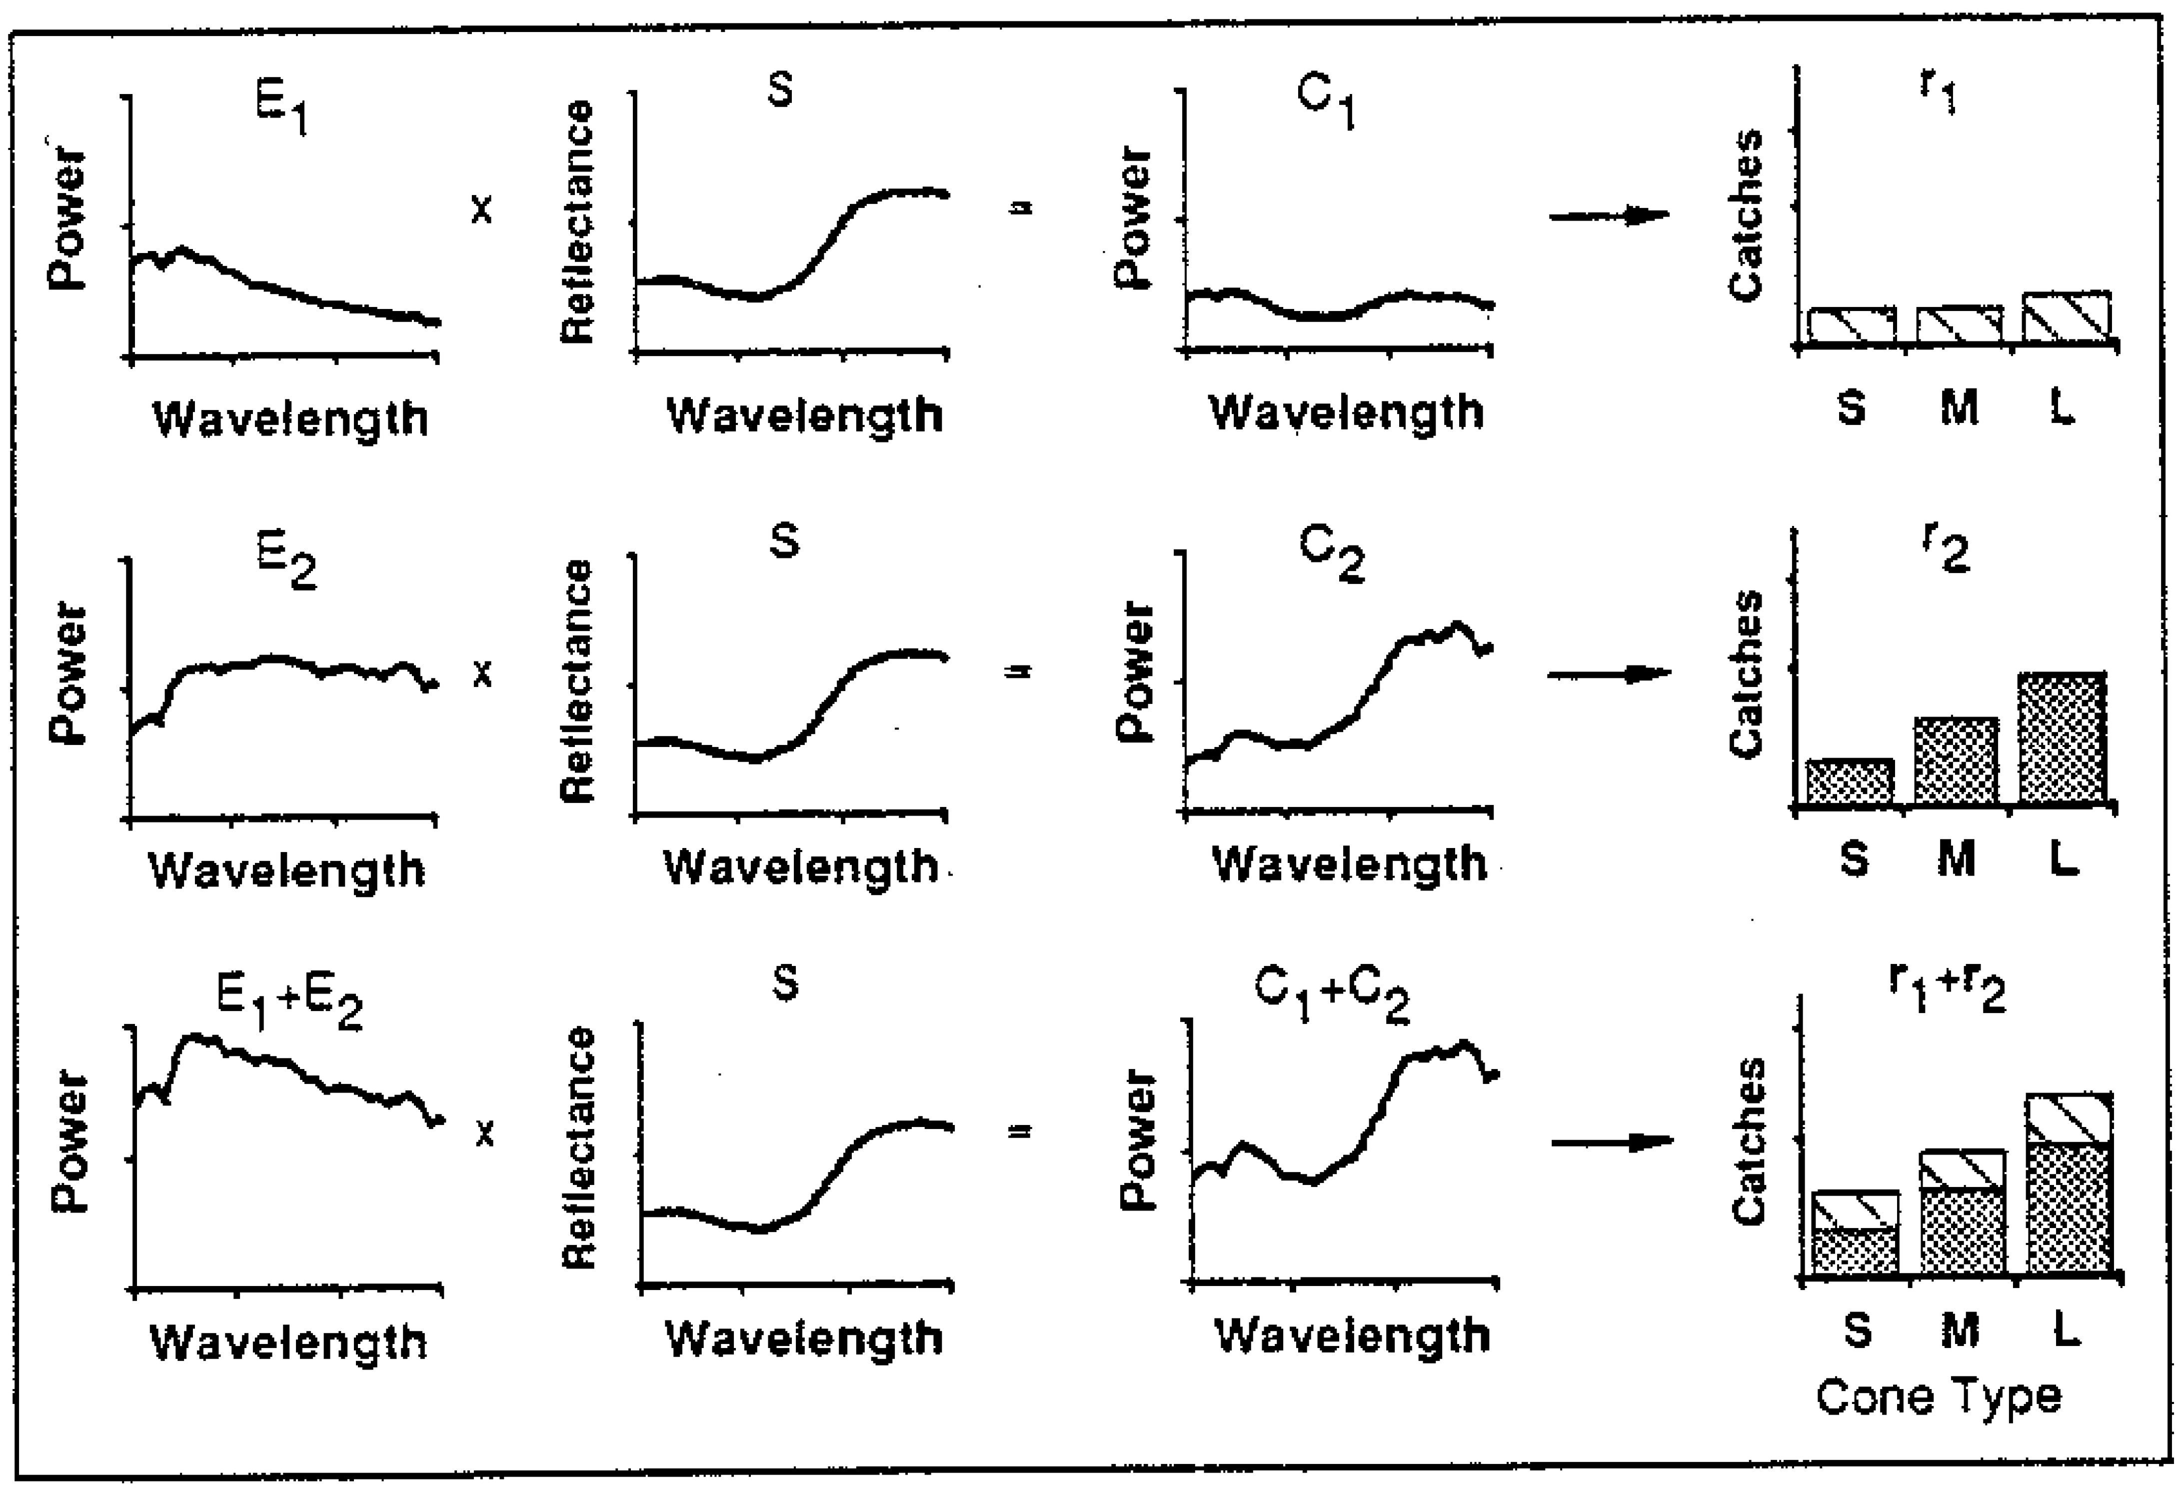

Fig. 4. The bilinear relation between illuminants, surfaces, and quantum catches. The top row shows the color signal and cone quantum catches when an illuminant E_1 reflects from a surface S. The middle row shows the color signal and cone quantum catches when a second illuminant, E_2, reflects from the same surface. The bottom row illustrates that when the sum $E_1 + E_2$ reflects from the surface, the color signal is the sum $C_1 + C_2$. Because photopigment absorption is a linear process, we can extend this statement to the level of the cone quantum catches. The quantum catches under $E_1 + E_2$ are the sum $\mathbf{r}_1 + \mathbf{r}_2$ of the quantum catches under E_1 and E_2 individually. The symmetric roles of illuminants and surfaces in image formation permit us to show that the cone quantum catches are also linear with respect to sums of surface reflectance functions.

COLOR APPEARANCE EXPERIMENTS

The theoretical work on color constancy has important implications for the study of human color vision. Indeed, we have organized our experimental work using three ideas taken from the theory. Two of these ideas indicate how to choose experimental conditions that allow color constancy: (a) The illuminant and surface functions should be constrained by small-dimensional linear models, and (b) the images should contain multiple surfaces viewed under a common illuminant. The third idea is the bilinear nature of the physical relation between illuminants, surfaces, and quantum catches. We analyzed our data to test whether the psychological relation between illuminants, surfaces, and color appearance is consistent with bilinearity.

Cross-Illuminant Color Matching

We measured how color appearance depends on the illuminant using an asymmetric color-matching task, which we have described in detail elsewhere.[8] Subjects viewed a computer display consisting of a simulation of matte surfaces under a standard illuminant typical of daylight. In each session, we trained subjects to remember the color appearance of a single test surface under the standard illuminant. We then had subjects set color matches to the test surface from memory. On control trials, subjects set matches while the illuminant was left unchanged. On test trials, subjects set matches under a different illuminant. The

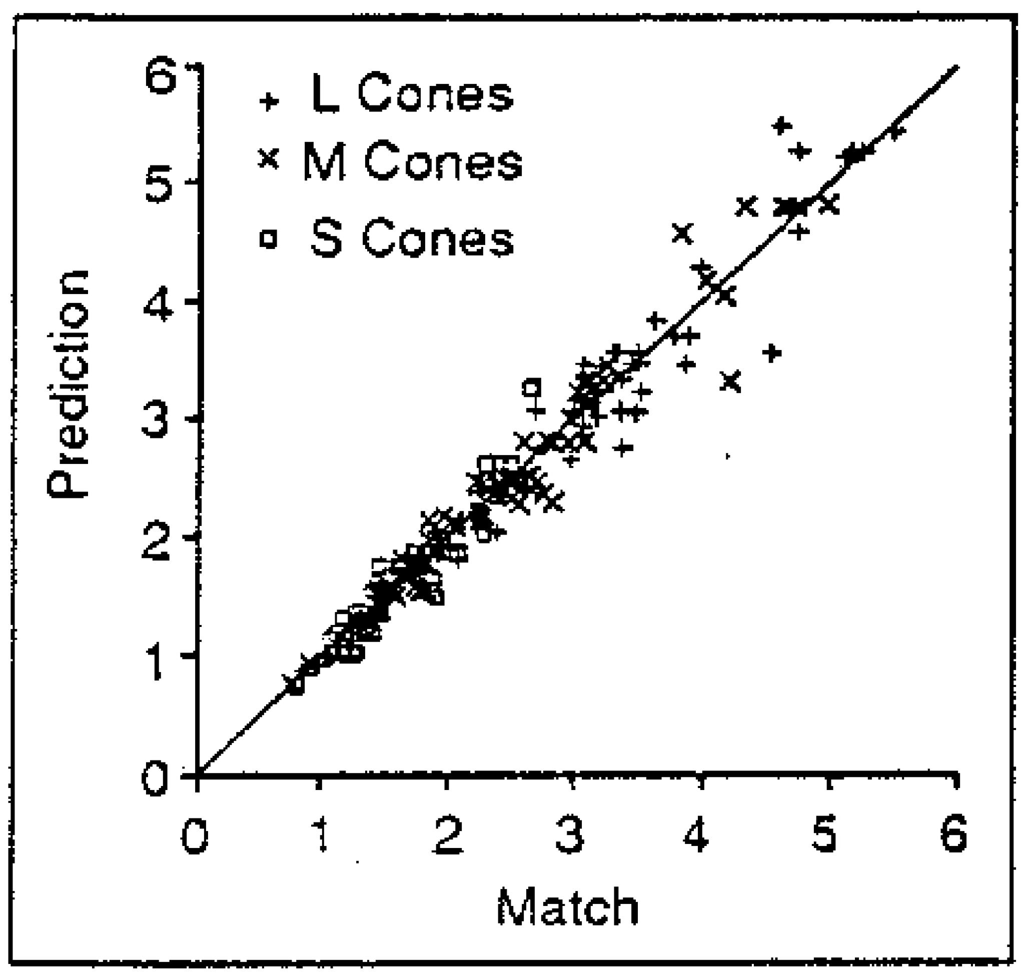

Fig. 5. Test of bilinearity as a model of human performance. The predicted cone quantum catches for a collection of asymmetric matches, computed under the constraint of bilinearity, are plotted against the corresponding measured cone quantum catches. The data were collected using our cross-illuminant matching procedure. The three symbol types indicate the results for the L, M, and S cones. If the bilinear model held perfectly, all the points would lie along the diagonal.

control matches measured the subjects' ability to remember the color of the test surface, while the test matches measured the effect of the illumination change. Our experimental design incorporated the ideas taken from the theoretical analysis. We used many surfaces in our displays, and we simulated illuminants and surfaces whose spectral functions are well described by small-dimensional linear models.

The Bilinear Structure of Cross-Illuminant Color Matches

We were particularly interested in whether color appearance judgments exhibit the same bilinear structure that emerged from the analysis of illuminant and surface recovery. To see how bilinearity might be reflected in the color appearance data, consider the following experimental conditions. Suppose the observer makes matches for two separate illuminant changes ΔE_1 and ΔE_2. If color appearance follows the bilinear structure, we should be able to predict the match to the sum of the illuminant changes, $\Delta E_1 + \Delta E_2$. The bilinear model also makes predictions for the effect of varying the surface reflectance when the illuminant change is held fixed.[8,9] We used our cross-illuminant matching data to check whether a bilinear model describes human performance. The results, summarized in Figure 5, support bilinearity.

The bilinearity of cross-illuminant matching data has an important consequence for theories of color appearance. Suppose the surfaces and illuminants are described by small-dimensional linear models. By measuring the effect of variation along the illuminant linear model dimensions on the color appearance of the surface linear model bases, we can determine the parameters of the bilinear transformation. Now suppose we wish to know the color appearance of a test surface under some illuminant. We can use the bilinear transformation to find a surface under a standard, canonical illuminant whose color appearance matches that of the test surface. Studies of color naming and ordering properties need be carried out only under the standard canonical illuminant. Bilinearity permits a great reduction in the number of measurements necessary to characterize color appearance.

How Color Constant Are Observers?

We can use our cross-illuminant matching data to evaluate how well observers discount the illuminant change. Suppose we assume that when the subject sets a cross-illuminant match, he or she believes that the surface reflectance function of the simulated object is identical to the surface reflectance function of the original. Under this assumption, we can estimate from the matching data the illuminant change that the

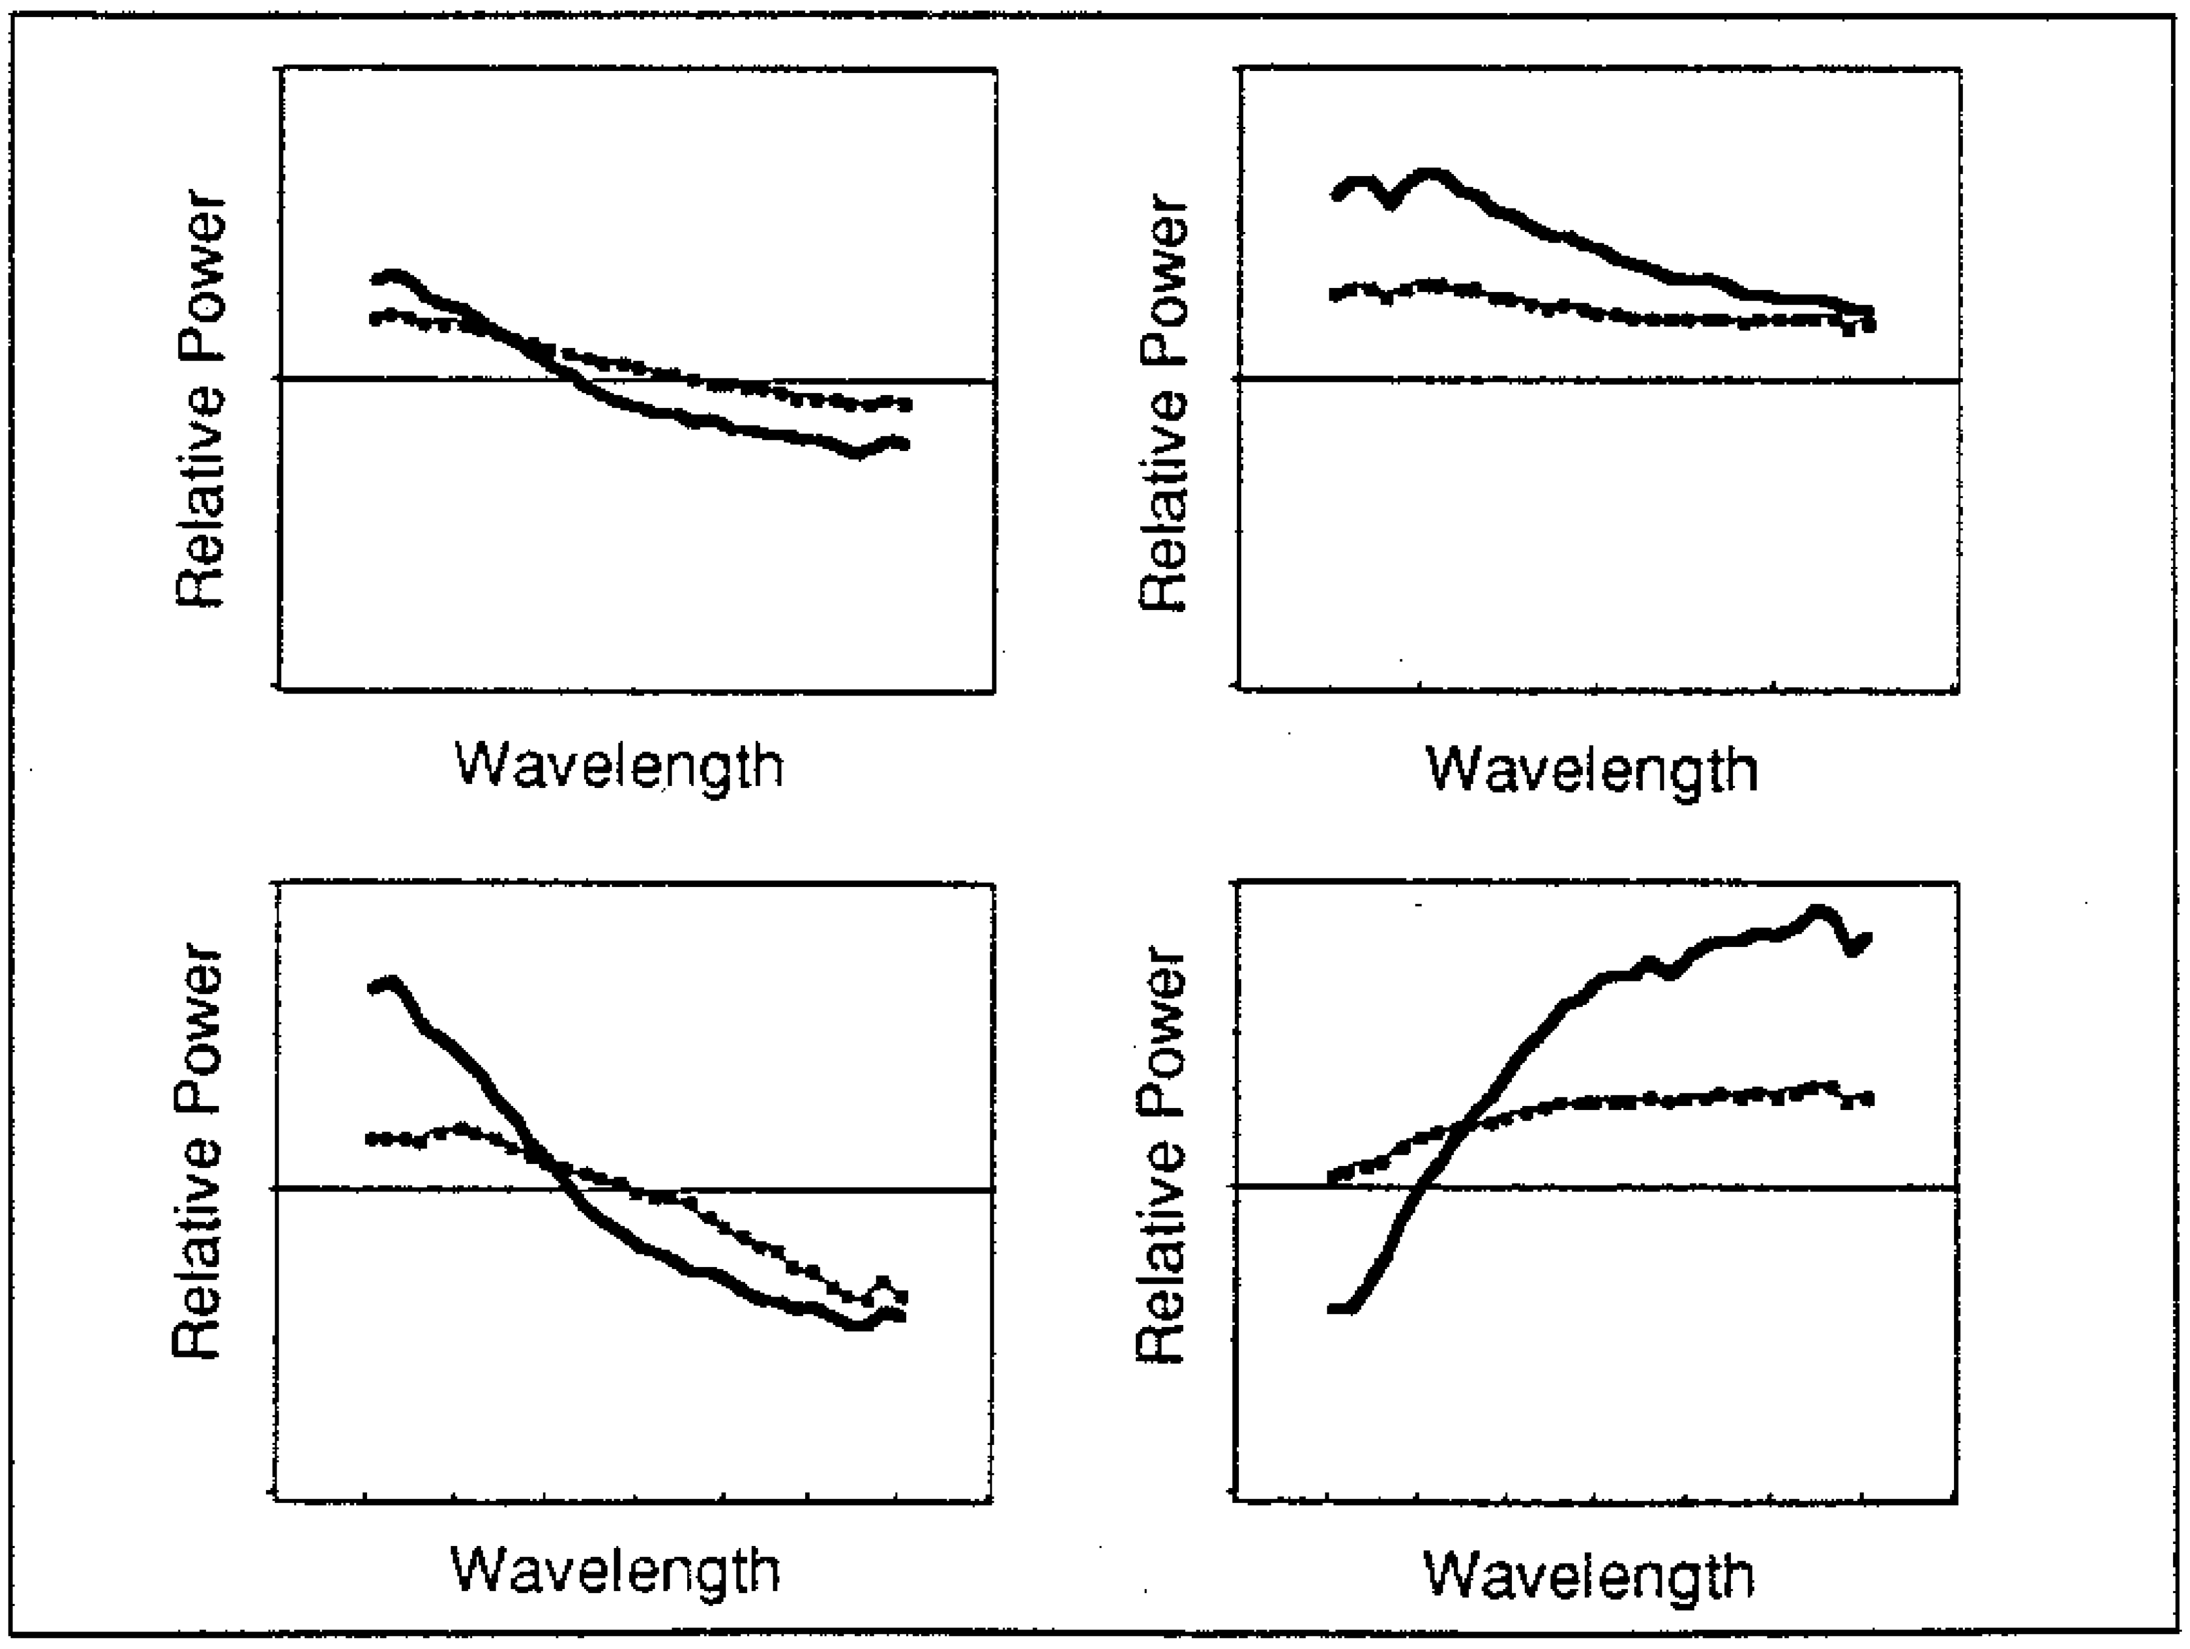

Fig. 6. Equivalent illuminant representation of our cross-illuminant matching data. The solid lines plot the spectral power distribution of four different experimental illuminant changes. The connected points show the spectral power distribution of the corresponding equivalent illuminant change. Each equivalent illuminant change is computed from a single cross-illuminant match. If subjects were perfectly color constant, the two spectral power distributions would coincide.

subject must have perceived. We call the subject's perceived change the *equivalent illuminant change.* Figure 6 plots the experimental and equivalent illuminant change for four experimental conditions. The figure illustrates that the equivalent illuminant change is similar in relative spectral power distribution to the experimental illuminant change, but that it is not as large: Our subjects show partial color constancy.[9]

Perhaps observers are always only partially color constant. But our experiments measured performance for rather simple scenes. These scenes do contain sufficient information for complete color constancy, but in natural images observers may use other kinds of information as well. Our approach can be used to organize the study of color appearance in richer scenes. The key hypothesis is that the illuminant controls the relation between quantum catches and color appearance.

Neural Mechanisms of Color Appearance

Because the visual system adjusts to changes in illumination, our cross-illuminant color matches were not photopigment absorption matches. Rather, our experiments measured equivalence at a later point in the visual pathways. We can explain our results by assuming that the only effect of changing the illumination is to scale the signals initiated by the cone quantum catches. This is often referred to as von Kries adaptation.[10] What is new about our results is how these scale factors are related to the illuminant: Our results suggest that the sensitivity of each cone class is inversely proportional to a linear function of the illuminant.[8] This in turn suggests that early gain control mechanisms share a common quantitative structure with illuminant and surface recovery algorithms. The study of physiological gain control may benefit from understanding its role in determining color appearance.

Acknowledgments—We thank M. D'Zmura, L. Maloney, D. Marimont, A. Poirson, and D. Varner for many useful discussions. The editor and referees also provided helpful comments. This work was supported by Grant 2 RO1 EY03164 from the National Eye Institute, Grant NCC 2-307 from the NASA-Ames Research Center, and a National Science Foundation graduate fellowship.

Notes

1. L.T. Maloney, Color constancy and color perception: The linear models framework, in *Attention and Performance XIV: Synergies in Experimental Psychology, Artificial Intelligence, and Cognitive Neuroscience,* D.E. Meyer and S.E. Kornblum, Eds. (MIT Press, Cambridge, MA, 1992); D.H. Marimont and B.A. Wandell, Linear models of surface and illuminant spectra, *Journal of the Optical Society of America A, 9,* 1905–1913 (1992).

2. J. Cohen, Dependency of the spectral reflectance curves of the Munsell color chips, *Psychonomic Science, 1,* 369–370 (1964).

3. D.B. Judd, D.L. MacAdam, and G.W. Wyszecki, Spectral distribution of typical daylight as a function of correlated color temperature, *Journal of the Optical Society of America, 54,* 1031 (1964).

4. L.T. Maloney and B.A. Wandell, Color constancy: A method for recovering surface spectral reflectances, *Journal of the Optical Society of America A, 3,* 29–33 (1986); B.A. Wandel, The synthesis and analysis of color images, *IEEE Transactions on Pattern Analysis and Machine Intelligence, PAMI-9,* 2–13 (1987).

5. M. D'Zmura and G. Iverson, *Color Structure From Chromatic Motion: I. Basic Theory,* Technical Report MBS 92-25 (Institute for Mathematical Behavioral Sciences, University of California, Irvine, 1992).

6. D.H. Brainard and B.A. Wandell, Analysis of the retinex theory of color vision, *Journal of the Optical Society of America A, 3,* 1651–1661 (1986).

7. D.J. Heeger and A.D. Jepson, Subspace methods for recovering rigid motion: I. Algorithm and implementation, *International Journal of Computer Vision, 7,* 95–117 (1992).

8. D.H. Brainard and B.A. Wandell, Asymmetric color-matching: How color appearance depends on the illuminant, *Journal of the Optical Society of America A, 9,* 1433–1448 (1992).

9. D.H. Brainard and B.A. Wandell, A bilinear model of the illuminant's effect on color appearance, in *Computational Models of Visual Processing,* M.S. Landy and J.A. Movshon, Eds. (MIT Press, Cambridge, MA, 1991).

10. J. von Kries, Influence of adaptation on the effects produced by luminous stimuli, in *Sources of Color Vision,* D.L. MacAdam, Ed. (MIT Press, Cambridge, MA, 1970). (Original work published 1905)

Published by Cambridge University Press