(e.g., spatial frequency, orientation, motion, depth) within a local cortical region. With respect to color vision *per se*, the primary processing involves separating color and luminance information, and further separating changes due to the illuminant from those due to visual objects, by lateral interactions over large regions.

To separate luminance and color information, the outputs of $\boldsymbol{P}_e$ cells are combined in two different ways. When their outputs are summed in one way, the luminance components to their responses sum and the color components cancel. Summed in a different combination, the color components sum and the luminance components cancel. Consider a striate cortex cell that combines inputs from one or more $+L_o$ and $+M_o$ cells in a region. The cortical cell would respond to luminance variations but not to color variations, since the neurons providing its inputs both fire to luminance increments in the RF center and to decrements in the surround, but the color organizations of its inputs are opposite to each other (one being L-M and the other M-L). Combined with input from a $+S_o$ cell, this would produce a V1 cell that fires to white (light increments) and inhibits to black (light decrements) but does not respond to pure color variations. This is represented in the top row of Fig. 1C. However, a V1 cell receiving inputs from both $+L_o$ and $-M_o$ cells, or from both $+M_o$ and $-L_o$ cells (columns in Fig. 1C), would respond to color changes but not to luminance variations since their color responses would add, but their luminance RFs, which are opposite to each other, would cancel. This organization by itself would produce L-M color cells that would fire to so-called warm colors (red and yellow) and inhibit to cool colors (blue and green). M-L cells would fire to cool colors and inhibit to warm colors. As shown in Fig. 1C, the further addition of $+S_o$ or $-S_o$ cells can split these classes into separate red and yellow, and separate blue and green systems, respectively.

All of the primary visual information is passed through V1, but subsequent visual areas are partially specialized for the further analysis of various different functional aspects of vision. One later visual area (V4) is crucially involved with color perception. Individuals with localized V4 lesions can still discriminate objects on the basis of their color variations, but they report that the objects now appear to have no hue, as if viewed on a black-white television screen. There is also a report of one case with the reverse loss: a patient who could see colored but not black-white objects.

11. Color Appearance

The appearance of a color can be specified by values along just three perceptual dimensions known as hue, saturation and brightness. Hue refers to the characteristic described by such color names as red, yellow, green, and blue. Saturation refers to the extent to which the stimulus differs perceptually from a purely achromatic (i.e., white, gray, black) axis. The third dimension is brightness or lightness. That our perceptual space is three-dimensional reflects the basic trichromacy of vision.

A normal observer can describe the hue of any light (disregarding surface characteristics) by using one or more of only four color names (red, yellow, green, and blue). These so-called unique hues form two opponent pairs, red–green and blue–yellow. Red and green normally cannot be seen in the same place at the same time; if unique red and unique green lights are added in appropriate proportions, the colors cancel and one sees a neutral gray. Orange can be seen as a mixture of red and yellow, and purple as a mixture of red and blue, but there is no color seen as a red–green mixture (or as a blue–yellow mixture). This perceptual opponency is also reflected in color contrast. Red can induce the appearance of green into neighboring regions, and after staring at a red surface one sees a green after-image. The yellow–blue opponent pair produces similar effects. It was these perceptual characteristics of color that led Ewald Hering in the nineteenth century to propose that the various color systems were not independent but rather that color was processed in a spectrally opponent organization, an idea which has since been amply verified in the presence, discussed above, of spectrally-opponent cells in the path from receptors to the cortex.

See also: Color Vision Theory; Vision, Low-level Theory of; Vision, Psychology of; Visual Perception, Neural Basis of; Visual System in the Brain

Bibliography

De Valois R L, De Valois R L 1988 *Spatial Vision*. Oxford University Press, New York

Hurvich L M 1981 *Color Vision*. Sinauer Press, Sunderland, MA

Kaiser P K, Boynton R M 1996 *Human Color Vision*. Optical Society of America, Washington, DC

Neitz J, Neitz M 1998 Molecular genetics and the biological basis of color vision. In: Backhaus W G S, Kliegl R, Werner J S (eds.) *Color Vision*. Walter de Gruyter, Berlin, pp. 101–19

Spillmann L, Werner J S 1990 *Visual Perception: The Neurophysiological Foundations*. Academic Press, New York

K. K. De Valois and R. L. De Valois

Color Vision Theory

Color vision is the ability to distinguish and identify lights and objects on the basis of their spectral properties. This entry presents several key topics that underlie current theories of human color vision. These are trichromacy, color opponency, adaptation, and color constancy.

1. Introduction

Information about color is transformed as it flows from the stimulus through the initial stages of the human visual system. At each image location, the color stimulus is specified by the amount of power it contains at each wavelength. The classic color matching experiment shows that the normal human visual system is trichromatic: only three dimensions of spectral variation are coded by the visual system. The biological basis of normal trichromacy is that the retina contains three classes of cone photopigment. After the initial encoding of light by the cones, further processing occurs. Two aspects of this processing are particularly important. First, signals from three classes of cones are recombined to form a luminance and two color opponent channels. Second, there is adaptive signal regulation that keeps neural signals within their operating range and stabilizes the appearance of objects across changes of illumination.

2. Trichromacy

2.1 Color Matching

The physical property of light relevant for color vision is the spectral power distribution. A light's spectral power distribution specifies the amount of power it contains at each wavelength in the visible spectrum, often taken to lie roughly between 400 and 700 nm. In practice, spectral power distributions are measured at discrete sample wavelengths. Let the measured power values be denoted by $b_1, \ldots, b_{N_\lambda}$ where N_λ denotes the number of sample wavelengths. Then the vector

$$\mathbf{b} = \begin{bmatrix} b_1 \\ \vdots \\ b_{N_\lambda} \end{bmatrix} \tag{1}$$

provides a compact representation of the spectral power distribution. Use of a vector representation for spectral quantities facilitates a variety of colorimetric computations (e.g., Brainard 1995). Wavelength sample spacings between 1 and 10 nm are typical.

Trichromacy is demonstrated by the basic color matching experiment (Wandell 1995, Brainard 1995). In this experiment, an observer views a bipartite field. One side of the field contains a test light. This light is experimentally controlled and can have an arbitrary spectral power distribution. On the other side of the field is the matching light. This consists of the weighted mixture of three primary lights. Each primary has a fixed relative spectral power distribution, but its overall intensity in the mixture can be controlled by the observer. The observer's task is to adjust the primary intensities until the mixture has the same color appearance as the test light. The primaries used in the experiment are chosen to be independent, so that no weighted mixture of any two produces a match to the third.

Because the matching light is constrained to be a weighted mixture of three primaries, it will not generally be possible for the observer to make the test and matching lights physically identical. For many test lights, however, the observer can adjust the matching light so that it appears identical to the test light even though the two differ physically. For some test lights, no choice of primary intensities will afford a match. In these cases one or more of the primaries can be mixed with the test light and primary intensities found so that the primary/test mixture matches the mixture of the remaining primaries. A useful descriptive convention for the color matching experiment is to assign a negative intensity to any primary that must be mixed with the test to make a match. Given this convention, any test light can be matched by a mixture of three independent primaries.

The color matching experiment is an empirical system. Given a test light described by a vector $\mathbf{b}$, the experiment returns a vector

$$\mathbf{t} = \begin{bmatrix} t_1 \\ t_2 \\ t_3 \end{bmatrix} \tag{2}$$

whose entries are the individual primary intensities. When the primaries are scaled by these intensities and mixed, a match to the test light is created. The vector $\mathbf{t}$ specifies what are called the tristimulus coordinates of the light $\mathbf{b}$. A theory of color matching should let us predict $\mathbf{t}$ for any test light $\mathbf{b}$, given the spectral power distributions of the primary lights.

As an empirical generalization, the color matching system is a linear system (e.g., Wyszecki and Stiles 1982, Brainard 1995, Wandell 1995). That is, if we have two test lights $\mathbf{b}_1$ and $\mathbf{b}_2$ with tristimulus coordinates $\mathbf{t}_1$ and $\mathbf{t}_2$, then any weighted mixture $(a_1\mathbf{b}_1 + a_2\mathbf{b}_2)$ of the two test lights has tristimulus coordinates given by the corresponding mixture $(a_1\mathbf{t}_1 + a_2\mathbf{t}_2)$. In these vector expressions, multiplication of a vector (e.g., $\mathbf{b}_1$) by a scalar (e.g., a_1) consists of multiplying each entry of the vector by the scalar, while addition of two vectors (e.g., $a_1\mathbf{b}_1$ and $a_2\mathbf{b}_2$) consists of adding the corresponding entries of the two vectors.

The linearity of color matching makes it possible to predict the match that will be made to any test light on the basis of a relatively small number of measurements. Consider the set of monochromatic lights with unit power. If N_λ wavelength samples are used in the underlying representation, then there are N_λ such lights and we can denote their spectral representations by $\mathbf{c}_1, \mathbf{c}_2, \ldots, \mathbf{c}_{N_\lambda}$. Each of the $\mathbf{c}_i$ has a 1 as its ith entry

and zeros elsewhere. Note that any light **b** may be thought of as a weighted mixture of monochromatic lights, so that $\mathbf{b} = \sum_i b_i \mathbf{c}_i$ where b_i is the ith entry of **b**. Let the vectors $\mathbf{t}_i$ specify the tristimulus coordinates of the monochromatic lights $\mathbf{c}_i$. The linearity of color matching then tells us that the tristimulus coordinates of any light **b** are given by $\mathbf{t} = \sum_i b_i \mathbf{t}_i$.

A set of tristimulus values $\mathbf{t}_i$ measured for monochromatic lights $\mathbf{c}_i$ is referred to as a set of color matching functions. Although these are often plotted as a function of wavelength, they do not represent the spectral power distributions of lights. The color matching functions may be specified by a single matrix

$$\mathbf{T} = \left[\mathbf{t}_1 \mathbf{t}_2 \mathbf{t}_3 \cdots \mathbf{t}_{N_\lambda} \right] \tag{3}$$

whose N_λ columns consist of the individual tristimulus coordinate vectors $\mathbf{t}_i$. This specification allows computation of tristimulus coordinates from spectral power distributions through simple matrix multiplication:

$$\mathbf{t} = \mathbf{T}\mathbf{b}. \tag{4}$$

Both tristimulus values and color matching functions are defined with respect to the primaries chosen for the underlying color matching experiment. The Commission Internationale de l'Eclairage (CIE) has standardized a system for color representation based on the ideas outlined above. The CIE system is widely used to specify color stimuli and many sources describe it in detail (e.g., Wyszecki and Stiles 1982, Brainard 1995, Kaiser and Boynton 1996).

The advantage of using tristimulus coordinates to describe color stimuli is that they provide a more compact and tractable description than a description in terms of wavelength. Tristimulus coordinates are compact precisely because they do not preserve physical differences that are invisible to the human visual system. The representational simplification afforded by tristimulus coordinates is extremely valuable for studying processing that occurs after the initial encoding of light. On the other hand, it is important to remember that the standard tristimulus representations (e.g., the CIE system) are based on matches made by a typical observer looking directly at a small stimulus at moderate to high light levels. These representations are not necessarily appropriate for applications involving some individual observers, nonhuman color vision, or color cameras (e.g., Wyszecki and Stiles 1982, Brainard 1995).

2.2 Biological Basis of Color Matching

The color matching experiment is agnostic about the biological mechanisms that underlie trichromacy. It is generally accepted, however, that trichromacy typically arises because color vision is mediated by three types of cone photoreceptor. Direct physiological measurements of individual primate cones support this hypothesis (see Wandell 1995, Rodieck 1998). First, the responses of individual cones depend only on the rate at which photopigment molecules are isomerized by the absorption of light quanta; once the intensity of two lights has been adjusted so that they produce the same isomerization rates, the cone response does not distinguish the two lights. This idea is referred to as the principle of univariance. Second, individual cones may be classified into one of three distinct types, each with a characteristic spectral sensitivity. The spectral sensitivity is proportional to the probability that light quanta of different wavelengths will isomerize a molecule of the cone's photopigment. The three types of cones are often referred to as the long- (L), middle- (M), and short- (S) wavelength-sensitive cones. If an observer has only three types of cones, each of which obeys the principle of univariance, two physically distinct lights that produce the same isomerization rates for all three classes of cones will be indistinguishable to the visual system. Quantitative comparison confirms that color matches set by a standard observer (defined as the average of matches set by many individual observers) are well predicted by the equations of isomerization rates in the L-, M-, and S-cones.

As described above, trichromacy occurs for most observers because their retinas contain cones with three classes of photopigments. Genetic considerations, however, indicate that some individuals have retinas containing four classes of cone photopigments (Sharpe et al. 1999). Either these individuals are tetrachromatic (mixture of four primaries required to match any light) or else their trichromacy is mediated by information lost after quantal absorption. In addition, some human observers are dichromatic (only two primaries must be mixed to make a match to any light.) Most cases of dichromacy occur because one photopigment is missing (Sharpe et al. 1999, Neitz and Neitz 2000).

An alternative to using tristimulus coordinates to represent the spectral properties of lights is to use cone coordinates. These are proportional to the isomerization rates of the three classes of cone photopigments. The three dimensional vector

$$\mathbf{q} = \begin{bmatrix} q_L \\ q_M \\ q_S \end{bmatrix} \tag{5}$$

specifies cone coordinates where q_L, q_M, and q_S denote the isomerization rates of the L-, M-, and S-cone photopigments respectively. It can be shown (e.g., Brainard 1995) that cone coordinates and tristimulus coordinates are related by a linear transformation, so that

$$\mathbf{q} = \mathbf{M}_{tq}\mathbf{t} \tag{6}$$

where $\mathbf{M}_{tq}$ is an appropriately chosen 3 by 3 matrix.

Computation of cone coordinates from light spectra requires estimates of the cone spectral sensitivities. For each cone class, these specify the isomerization rates produced by monochromatic lights of unit power. The sensitivities may be specified in matrix form as

$$\mathbf{S} = \begin{bmatrix} \mathbf{s}_L \\ \mathbf{s}_M \\ \mathbf{s}_S \end{bmatrix} \tag{7}$$

where each row of the matrix is a vector whose entries are the spectral sensitivities for one cone class at the sample wavelengths. Given **S**, cone coordinates are computed from the spectral power distribution of a light as

$$\mathbf{q} = \mathbf{Sb} \tag{8}$$

Because Eqns. (4), (6), and (8) hold for any light spectrum **b**, it follows that

$$\mathbf{S} = \mathbf{M}_{tq}\mathbf{T} \tag{9}$$

Current estimates of human cone spectral sensitivities are obtained from color matching data using Eqn. (9) together with a variety of considerations that put constraints on the matrix $\mathbf{M}_{tq}$ (Stockman and Sharpe 1999).

3. *Postabsorption Processing*

Color vision does not end with the absorption of light by cone photopigments. Rather, the signals that originate with the absorption of light are transformed as they propagate through neurons in the retina and cortex. Two ideas dominate models of this post-absorption processing. The first is color opponency: signals from different cone types are combined in an antagonistic fashion to produce the visual representation at a more central site. The second idea is adaptation: the relation between the cone coordinates of a light and its central visual representation is not fixed but depends instead on the context in which the light is viewed. Section 3.1 treats opponency, while Sect. 3.2 treats adaptation.

3.1 *Opponency*

Direct physiological measurements of the responses of neurons in the primate retina support the general idea of opponency (e.g., Dacey 2000). These measurements reveal, for example, that some retinal ganglion cells are excited by signals from L-cones and inhibited by signals from M-cones. One suggestion about why this occurs is that it is an effective way to code the cone signals for transmission down the optic nerve (see Wandell 1995).

A possible approach to understanding post-absorption processing is to keep the modeling close to the underlying anatomy and physiology and to characterize what happens to signals at each synapse in the neural chain between photoreceptors and some site in visual cortex. The difficulty is that it is not presently possible to cope with the complexity of actual neural processing. Thus many color theorists have attempted to step back from the details and develop more abstract descriptions of the effect of neural processing. Models of this sort are often called mechanistic models. These models generally specify a transformation between the quantal absorption rates **q** elicited by a stimulus and a corresponding visual representation **u** postulated to exist at some central site. The idea is to choose a transformation so that (a) the color appearance perceived at a location may be obtained directly from the central representation corresponding to that location and (b) the discriminability of two stimuli is predictable from the difference in their central representations.

Most mechanistic models assume that signals from the cones are combined additively to produce signals at three postreceptoral sites. Two of these sites carry opponent signals. These are often referred to as the red-green (RG) and blue-yellow (BY) signals. A third site carries a luminance (LUM) signal, which is not thought to be opponent. If we take

$$\mathbf{u} = \begin{bmatrix} u_{LUM} \\ u_{RG} \\ u_{BY} \end{bmatrix} \tag{10}$$

to be a three-dimensional vector with entries given by the LUM, RG, and BY signals, then the additive relation between cone coordinates **q** and the visual representation **u** may be expressed in matrix form:

$$\mathbf{u} = \mathbf{M}_o\mathbf{q} \tag{11}$$

Many (but not all) detailed models take LUM to be a weighted sum of L- and M-cone signals, RG to be a weighted difference between the L- and M-cone signals, and BY to be a weighted difference between the S-cone signal and a weighted sum of the L- and M-cone signals. In these models $\mathbf{M}_o$ would have the form

$$\mathbf{M}_o = \begin{bmatrix} m_{11} & m_{12} & 0 \\ m_{21} & -m_{22} & 0 \\ -m_{31} & -m_{32} & m_{33} \end{bmatrix} \tag{12}$$

where all of the m_{ij} are positive scalars representing how strongly one cone class contributes to the signal at one post-receptoral site.

Considerable effort has been devoted to establishing whether the linear form for the mapping between **q** and **u** is appropriate, and if so, what values should be

used for the m_{ij}. Several types of experimental evidence have been brought to bear on the question.

As an example, one line of research argues that four color perceptions, those of redness, greenness, blueness, and yellowness, have a special psychological status, in that any color experience may be intuitively described in terms of these four basic perceptions. Thus orange may be naturally described as reddish-yellow and aqua as greenish-blue. In addition, both introspection and color scaling experiments suggest that the percepts of redness and greenness are mutually exclusive so that both are not experienced simultaneously in response to the same stimulus, and similarly for blueness and yellowness (e.g., Hurvich and Jameson 1957, Abramov and Gordon 1994). Given these observations, it is natural to associate the RG signal with the amount of redness or greenness perceived in a light (redness if the signal is positive, greenness if it is negative, and neither red nor green if it is zero) and the BY signal with the amount of blueness or yellowness. Judgments of the four fundamental color perceptions, obtained either through direct scaling (e.g., Abramov and Gordon 1994) or through a hue cancellation procedure (e.g., Hurvich and Jameson 1957), are then used to deduce the appropriate values of the m_{ij} in the second and third rows of $\mathbf{M}_o$. When this framework is used, the entries for the first row of $\mathbf{M}_o$, corresponding to the LUM signal, are typically established through other means such as flicker photometry (e.g., Kaiser and Boynton 1996).

Other approaches to studying the opponent transformation include analyzing measurements of the detection and discrimination of stimuli (e.g., Wyszecki and Stiles 1982, Kaiser and Boynton 1996, Eskew, et al. 1999, Wandell 1999), and measurements of how the color appearance of lights is affected by the context in which they are viewed (e.g., Webster 1996). In part because of a lack of quantitative agreement in the conclusions drawn from different paradigms, there is currently not much consensus about the details of the transformation between $\mathbf{q}$ and $\mathbf{u}$. One of the major open issues in color theory remains how to extend the simple linear model described above so that it accounts for a wider range of results.

3.2 Adaptation

Figure 1 illustrates a case where the same light has a very different color appearance when seen in two different contexts. The figure shows two disk-annulus stimulus configurations. The central disk is the same in each configuration, but the appearance of the two disks is quite different. To explain this and other context effects, mechanistic models assume that at any given time and image location, the relation between the quantal absorption rates $\mathbf{q}$ and the visual representation $\mathbf{u}$ depends on the quantal absorption rates

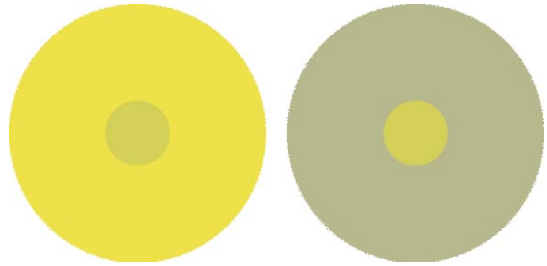

Figure 1
A color context effect. The figure illustrates the color context effect known as simultaneous contrast. The two central disks are physically the same but appear different. The difference in appearance is caused by the fact that each disk is seen in the context of a different annular surround. This figure is best viewed in color. A color version is available in the on-line version of the Encyclopedia

at other locations and at preceding times. To help fix ideas, it is useful to restrict attention to the disk-annulus configuration. For this configuration, the visual representation of the disk may be written as

$$\mathbf{u}_d = f(\mathbf{q}_d; \mathbf{q}_a, \varphi) \tag{13}$$

where $\mathbf{u}_d$ is the visual response to the disk, $\mathbf{q}_d$ and $\mathbf{q}_a$ are the cone coordinates of the disk and annulus respectively, and φ represents other contextual variables such as the size of the disk and annulus and any temporal variation in the stimulus. Clearly, f( ) must incorporate the sort of transformation described by the matrix $\mathbf{M}_o$ in Sect. 3.1 above.

As was the case with the discussion of opponency above, there is not wide agreement about how best to model adaptation. A reasonable point of departure is a cone-specific affine model. In this model, the visual representation $\mathbf{u}$ of a light is related to its cone coordinates $\mathbf{q}$ through an equation of the form

$$\mathbf{u} = \mathbf{M}_o(\mathbf{D}_1\mathbf{q} - \mathbf{q}_1) \tag{14}$$

where $\mathbf{M}_o$ is as in Eqn. (12) and

$$\mathbf{D}_1 = \begin{bmatrix} g_{L1} & 0 & 0 \\ 0 & g_{M1} & 0 \\ 0 & 0 & g_{S1} \end{bmatrix}, \quad \mathbf{q}_1 = \begin{bmatrix} q_{L1} \\ q_{M1} \\ q_{S1} \end{bmatrix} \tag{15}$$

In this formulation, the g's on the diagonals of $\mathbf{D}_1$ characterize multiplicative adaptation that occurs at a cone-specific site in visual processing, before signals from separate cone classes are combined. The entries of the vector $\mathbf{q}_1$ characterize subtractive adaptation. Equation (14) is written in a form that implies that the subtractive adaptation also occurs at a cone-specific

site. The entries of $\mathbf{D}_1$ and $\mathbf{q}_1$ depend on the cone coordinates $\mathbf{q}_a$ of the annulus as well as on spatial and temporal variables characterized by φ. Note that the cone-specific affine model is a generalization of the idea that the visual representation consists of a contrast code.

Asymmetric matching may be used to test the adaptation model of Eqn. (14). In an asymmetric matching experiment, an observer adjusts a match stimulus seen in one context so that it appears to have the same color as a test stimulus seen in another context. More concretely, consider Fig. 1. In the context of this figure, an asymmetric matching experiment could be conducted where the observer was asked to adjust the central disk on the right so that it matched the appearance of the central test disk on the left. Suppose such data are collected for a series of N test disks with cone coordinates $\mathbf{q}_{ti}$. Denote the cone coordinates of the matches by $\mathbf{q}_{mi}$. Within the mechanistic framework, the corresponding visual representations $\mathbf{u}_{ti}$ and $\mathbf{u}_{mi}$ should be equal. If Eqn. (14) provides a good description of performance then

$$\begin{aligned}\mathbf{M}_o(\mathbf{D}_{m1}\mathbf{q}_{mi}-\mathbf{q}_{m1}) &= \mathbf{M}_o(\mathbf{D}_{t1}\mathbf{q}_{ti}-\mathbf{q}_{t1})\\ &\Leftrightarrow \mathbf{q}_{mi} = \mathbf{D}_{m1}^{-1}(\mathbf{D}_{t1}\mathbf{q}_{ti}-\mathbf{q}_{t1}+\mathbf{q}_{m1})\\ &\Leftrightarrow \mathbf{q}_{mi} = \mathbf{D}_{tm}\mathbf{q}_{ti}-\mathbf{q}_{tm}\end{aligned} \quad (16)$$

where $\mathbf{D}_{tm} = \mathbf{D}_{m1}^{-1}\mathbf{D}_{t1}$ and $\mathbf{q}_{tm} = \mathbf{D}_{m1}^{-1}(\mathbf{q}_{t1}-\mathbf{q}_{m1})$. This prediction may be checked by finding the diagonal matrix $\mathbf{D}_{tm}$ and vector $\mathbf{q}_{tm}$ that provide the best fit to the data and evaluating the quality of the fit. Tests of this sort indicate that the cone specific affine model accounts for much of the variance in asymmetric matching data, both for the disk annulus configuration (Wandell 1995, 1999) and for more complex stimuli (Brainard and Wandell 1992). Nonetheless, there are clear instances for which Eqn. (16) does not give a complete account of asymmetric matching (e.g., Delahunt and Brainard 2000) and other color appearance data (e.g., Webster 1996, Mausfeld 1998, D'Zmura and Singer 1999).

The cone-specific affine model may also be tested against psychophysical data on the detection and discrimination of colored lights. Here again the model provides a reasonable point of departure but fails in detail (e.g., Eskew et al. 1999).

To extend the cone specific affine model, various theorists have suggested the need for adaptation at a second site (after signals from separate cone classes have been combined) and for the inclusion of nonlinearities in the relation between $\mathbf{q}$ and $\mathbf{u}$ (see references cited in the previous two paragraphs). An additional open question concerns how the entries of $\mathbf{D}_1$ and $\mathbf{q}_1$ are determined by the viewing context (e.g., Brainard and Wandell 1992, Delahunt and Brainard 2000).

4. Color Constancy

The discussion so far has focussed on how the visual system represents and processes the spectrum of light that enters the eye. This is natural, since light is the proximal stimulus that initiates color vision. On the other hand, we use color primarily to name objects. The spectrum of the light reflected to the eye from an object depends both on an intrinsic property of the object, its surface reflectance function, and on extrinsic factors, including the spectral power distribution of the illuminant and how the object is oriented relative to the observer.

Given that the light reflected to the eye varies with the illuminant and viewing geometry, how is it that color is a useful psychological property of objects? The answer is that the visual system processes the retinal image to stabilize the color appearance of objects across changes extrinsic to the object (e.g., changes in the spectrum of the illuminant). This stabilization process is called color constancy.

Color constancy is closely linked to the phenomenon of adaptation described above (Maloney 1999). Indeed, quantitative models of color constancy generally incorporate the same idea that underlies mechanistic models of visual processing: at some central site there is a visual representation $\mathbf{u}$ that correlates with color appearance. To stabilize this representation against changes in illumination, it is supposed that the relation between the quantal absorption rates $\mathbf{q}$ elicited by the light reflected from an object and the visual representation $\mathbf{u}$ depends on the scene in which the object is viewed. In the case of color constancy, the emphasis has been on how the visual system processes the retinal image so that the transformation between $\mathbf{q}$ and $\mathbf{u}$ has the effect of compensating for the variation in the light reflected to the eye caused by changes of illumination and viewing geometry. Psychophysical data on the color appearance of objects viewed under different illuminants are often well-modeled by transformations consistent with Eqn. (14) (e.g., Brainard and Wandell 1992).

The central theoretical question of color constancy is how the visual system can start with image data and factor it into an accurate representation of the surfaces and illuminants in the scene. This question has received extensive treatment, at least for simple scenes. A brief introduction to this literature on computational color constancy follows.

4.1 Computational Color Constancy

Consider a scene consisting of diffusely illuminated flat, matte surfaces. For such scenes, the spectrum $\mathbf{b}$ reflected to the eye from each surface is given by the wavelength-by-wavelength product of the spectral power distribution of the illuminant $\mathbf{e}$ and the surface reflectance function $\mathbf{s}$. The surface reflectance function specifies, at each sample wavelength, the fraction of

incident light reflected to the eye. The information about **b** coded by the visual system is its cone coordinates, which may be computed as

$$\mathbf{q} = \mathbf{S}\mathbf{b} = \mathbf{S}\ \mathrm{diag}(\mathbf{e})\mathbf{s} \tag{17}$$

where the function **diag**( ) returns a square diagonal matrix with its argument placed along the diagonal. Clearly **e** and **s** are not uniquely determined from knowledge of **q**: without additional constraints the color constancy problem is underdetermined. Fortunately the spectra of naturally occurring illuminants and surfaces are not arbitrary. Although the physical processes that constrain these spectra are not well understood, analyses of measurements of both illuminants and surfaces shows that their spectra are well described by small-dimensional linear models (see Brainard 1995, Maloney 1999).

Consider surface reflectances. It is possible to define three fixed basis functions so that naturally occurring surface reflectances are reasonably well approximated by a linear combination of these basis functions. Thus for any surface, we have

$$\mathbf{s} \approx w_{s1}\mathbf{b}_{s1} + w_{s2}\mathbf{b}_{s2} + w_{s3}\mathbf{b}_{s3} \tag{18}$$

where $\mathbf{b}_{s1}$, $\mathbf{b}_{s2}$, and $\mathbf{b}_{s3}$ are the spectra of the basis functions and w_{s1}, w_{s2}, and w_{s3} are scalar weights that provide the best approximation of **s** within the linear model. Eqn. (18) may be rewritten as

$$\mathbf{s} \approx \mathbf{B}_s\mathbf{w}_s \tag{19}$$

where the three columns of matrix $\mathbf{B}_s$ contain the basis functions and the vector w_s contains the scalar weights.

When the surface reflectance functions lie within a three-dimensional linear model Eqn. (17) may inverted, once an estimate $\hat{\mathbf{e}}$ of the illuminant has been obtained (see below for discussion of illuminant estimation.) Start by rewriting Eqn. (17) as:

$$\mathbf{q} = [\mathbf{S}\ \mathrm{diag}(\hat{\mathbf{e}})\mathbf{B}_s]\ \mathbf{w}_s = \mathbf{L}_{\hat{\mathbf{e}}}\mathbf{w}_s \tag{20}$$

where $\mathbf{L}_{\hat{\mathbf{e}}}$ is a three-by-three matrix that depends on the illuminant estimate. This matrix may be inverted using standard methods to yield an estimate of $\mathbf{w}_s$:

$$\hat{\mathbf{w}}_s = \mathbf{L}_{\hat{e}}^{-1}\ \mathbf{q} \tag{21}$$

The estimate may then be used together with Eqn. (19) to estimate the surface reflectance function.

Many computational color constancy algorithms assume a linear model constraint for surface reflectance functions. This reduces the constancy problem to finding an estimate of the illuminant to plug into Eqn. (20). For illustrative purposes, an algorithm due to Buchsbaum (1980) is described here. In Buchsbaum's algorithm, two additional assumptions are added. The first is that a three-dimensional linear model also describes illuminant spectral power distributions, so that

$$\mathbf{e} \approx \mathbf{B}_e\mathbf{w}_e \tag{22}$$

The second is that the spatial average of the surface reflectance functions ($\bar{\mathbf{s}}$) is the same in all scenes and known. These additional constraints imply that

$$\bar{\mathbf{q}} = [\mathbf{S}\ \mathrm{diag}(\bar{\mathbf{s}})\mathbf{B}_e]\hat{\mathbf{w}}_e = \mathbf{L}_{\bar{s}}\hat{\mathbf{w}}_e \tag{23}$$

where $\bar{\mathbf{q}}$ is the spatial average of the quantal absorption rates and $\mathbf{L}_{\bar{\mathbf{s}}}$ is a known three-by-three matrix. Inverting Eqn. 23 yields an estimate for the illuminant $\hat{\mathbf{e}} = \mathbf{B}_e\mathbf{w}_{\hat{e}}$. This estimate is then used to provide the matrix $\mathbf{L}_{\hat{e}}^{-1}$ to be used in Eqn. (21).

Buchsbaum's algorithm shows how the addition of appropriate assumptions allows solution of the computational color constancy problem. The difficulty with Buchsbaum's algorithm is that its assumptions are too restrictive. In particular, it seems unlikely that the spatial average of surface reflectances is constant across scenes, nor do real scenes consist of diffusely illuminated flat, matte surfaces. Subsequent work has focused on ways to provide reasonable estimates of the illuminant and surface reflectances under other sets of assumptions (e.g., Maloney 1999).

4.2 Computational Color Constancy and Human Performance

How does the computational work relate to human performance? This question has not yet been resolved, but it seems appropriate to close with a few observations. First, the estimated linear model weights of Eqn. (21) may be associated with the mechanism responses **u** discussed in Sect. 2. In both types of theory, these quantities represent the visual response computed from the quantal absorption rates, and both are meant to allow direct prediction of appearance. In the mechanistic approach, one considers a series of transformations whose form is derived from experiments with simple stimulus configurations. In the computational approach, the form of the transformation is derived from consideration of the problem color vision is trying to solve. In both cases, however, the emphasis is on finding the appropriate parametric form of the transformation and on understanding how the parameters are set as a function of the image data. The connection between the two approaches is discussed in more detail by Maloney (1999).

The value of the computational approach to understanding human vision depends on how accurately the transformations it posits may be used to predict the appearance of stimuli measured in psychophysical experiments. There have been only a few empirical comparisons of this sort to date. These comparisons do, however, indicate that the computational approach shows promise for advancing our understand-

ing of human color vision (Bloj, Kersten, and Hurlbert 1999, Brainard, Kraft, and Longre 2001).

See also: Color Vision; Psychophysical Theory and Laws, History of; Psychophysics; Vision, Low-level Theory of; Vision, Psychology of; Visual Perception, Neural Basis of; Visual System in the Brain

Bibliography

Abramov I, Gordon J 1994 Color appearance: on seeing red—or yellow, or green, or blue. *Annual Review of Psychology* **45**: 451–85

Bloj M G, Kersten D, Hurlbert A C 1999 Perception of three-dimensional shape influences colour perception through mutual illumination. *Nature* **402**: 877–9

Brainard D H 1995 Colorimetry. In: Bass M (ed.) *Handbook of Optics: Volume 1. Fundamentals, Techniques, and Design*. McGraw-Hill, New York, pp. 26.1–26.54

Brainard D H, Brunt W A, Speigle J M 1997 Color constancy in the nearly natural image. 1. Asymmetric matches. *Journal of the Optical Society of America A* **14**: 2091–110

Brainard D H, Kraft J M, Longère P 2001 Color constancy: developing empirical tests of computational models. In: Mausfeld R, Heyer D (eds.) *Colour Perception: From Light to Object*. Oxford University Press, Oxford, UK

Brainard D H, Wandell B A 1992 Asymmetric color-matching: How color appearance depends on the illuminant. *Journal of the Optical Society of America A* **9**(9): 1433–48

Buchsbaum G 1980 A spatial processor model for object colour perception. *Journal of the Franklin Institute* **310**: 1–26

D'Zmura M, Singer B 1999 Contrast gain control. In: Gegenfurtner K, Sharpe L T (eds.) *Color Vision: From Genes to Perception*. Cambridge University Press, Cambridge, UK, pp. 369–85

Dacey D M 2000 Parallel pathways for spectral coding in primate retina. *Annual Review of Neuroscience* **23**: 743–75

Delahunt P B, Brainard D H 2000 Control of chromatic adaptation: Signals from separate cone classes interact. *Vision Research* **40**: 2885–903

Eskew R T, McLellan J S, Giulianini F 1999 Chromatic detection and discrimination. In: Gegenfurtner K, Sharpe L T (eds.) *Color Vision: From Genes to Perception*. Cambridge University Press, Cambridge, UK, pp. 345–68

Hurvich L M, Jameson D 1957 An opponent-process theory of color vision. *Psychological Review* **64**(6): 384–404

Kaiser P K, Boynton R M 1996 *Human Color Vision*, 2nd edn. Optical Society of America, Washington, DC

Maloney L T 1999 Physics-based approaches to modeling surface color perception. In: Gegenfurtner K, Sharpe L T (eds.) *Color Vision: From Genes to Perception*. Cambridge University Press, Cambridge, UK, pp. 387–416

Mausfeld R 1998 Color perception: From Grassman codes to a dual code for object and illumination colors. In: Backhaus W G K, Kliegl R, Werner J S (eds.) *Color Vision—Perspectives from Different Disciplines*. Walter de Gruyter, Berlin, pp. 219–50

Neitz M, Neitz J 2000 Molecular genetics of color vision and color vision defects. *Archives of Ophthalmology* **118**: 691–700

Rodieck R W 1998 *The First Steps in Seeing*. Sinauer, Sunderland, MA

Sharpe L T, Stockman A, Jagle H, Nathans J 1999 Opsin genes, cone photopigments, color vision, and color blindness. In: Gegenfurtner K, Sharpe L T (eds.) *Color Vision*: *From Genes to Perception*. Cambridge University Press, Cambridge, UK, pp. 3–51

Stockman A, Sharpe L T 1999 Cone spectral sensitivities and color matching. In: Gegenfurtner K, Sharpe L T (eds.) *Color Vision: From Genes to Perception*. Cambridge University Press, Cambridge, UK, pp. 53–87

Wandell B A 1995 *Foundations of Vision*. Sinauer, Sunderland, MA

Wandell B A 1999 Computational neuroimaging: color representations and processing. In: Gazzaniga M (ed.) *The New Cognitive Neurosciences*, 2nd edn. MIT Press, Cambridge, MA, pp. 291–303

Webster M A 1996 Human colour perception and its adaptation. *Network: Computation in Neural Systems* **7**: 587–634

Wyszecki G, Stiles W S 1982 *Color Science—Concepts and Methods. Quantitative Data and Formulae*, 2nd edn. John Wiley, New York

D. H. Brainard

Combinatorial Data Analysis

Combinatorial data analysis (CDA) refers to a class of methods for the study of relevant data sets in which the arrangement of a collection of objects is the absolutely central concept. Characteristically, CDA is involved with either: (a) the identification of arrangements that are optimal for a specific representation of a given data set, and where such an exploratory process is typically carried out according to some specific loss or merit function that guides a combinatorial search over a domain of possible structures constructed from the constraints imposed by the particular representation selected; or (b) a confirmatory determination as to whether a specific object arrangement given *a priori* reflects the observed data, and where such a confirmatory process is typically operationalized by comparing the empirically observed degree of correspondence between some given data set and the specific structure conjectured for it, to a reference distribution constructed from the collection of all possible structures of the same form that could have been conjectured.

The boundaries of what CDA might encompass are somewhat open but generally we would exclude methods based on the postulation of strong stochastic models and their specific unknown parametric structures as underlying a given data set. Although CDA might use or empirically construct various weighting functions, the weights so obtained are not to be interpreted as parameter estimates in some presumed stochastic model viewed in turn as responsible for generating the data. Manifest data are emphasized solely, and the traditional concern for an assumed relationship between the data and a restrictively parameterized stochastic model is avoided. For

International Encyclopedia of the Social & Behavioral Sciences ISBN: 0-08-043076-7