

9

Color


Edwin Land was one of the great inventors and entrepreneurs in American history; among other achievements, he created instant developing film and founded the Polaroid Corporation. The first instant developing film made black and white reproductions, and after a few years Land decided to create a color version of the film. In order to learn about color appearance, Land returned to his laboratory to experiment with color. He was so surprised by his observations that he decided to write a paper summarizing them. The paper was published in the *Proceedings of the National Academy of Sciences*, and in it Land startled many people by arguing that there are only two, not three, types of cones. He further went on to dismiss the significance of the color-matching experiments. He wrote: "We have come to the conclusion that the classical laws of color mixing conceal great basic laws of color vision" (Land, 1959, p. 115). Land's sharp words, an arrow aimed at the heart of color science, provoked heated rejoinders from two leading scientists, Deane Judd (1960) and Gordon Walls (1960).

What was it that Land, a brilliant man, found so objectionable about color-matching? It seems that Land's reading of the literature led him to believe that the curves we measure in the color-matching experiment could be used to predict color appearance. When he set the textbooks down and began to experiment with color, he was sorely disappointed. He found that the color-matching measurements do not answer many important questions about color appearance.

Land's observation is consistent with our review of color-matching in Chapter 4. The results from the color-matching experiment can be

9.1 COLOR APPEARANCE IN A REGION DEPENDS ON THE SPATIAL PATTERN OF CONE ABSORPTIONS.

The two square regions are physically identical and thus create the same local rate of photopigment absorption. Yet, they appear to have different lightness because of the difference in their relative absorptions compared to nearby areas.



used to predict when two lights will look the same, but they cannot be used to tell us what the two lights look like. As Color Plate 2 illustrates, the experimental results in the color-matching experiment can be explained by the matches between the cone photopigment absorptions at a point, while color appearance forces us to think about the pattern of photopigment absorptions spread across the cone mosaics. Figure 9.1 illustrates this point again. The two squares in the image reflect the same amount of light to your eye. Yet, because the squares are embedded in different surroundings, we interpret the squares very differently, seeing one as light and the other as dark.

Land's paper contained a set of interesting qualitative demonstrations that illustrate these same points. While the limitations of the color-matching experiment were new to Land and the reviewers of his paper, they were not new to most color scientists. For example, Judd (1940) had worked for years trying to understand these effects. Later in this chapter I will review work at Kodak and in academic laboratories, contemporaneous with that of Land, that was designed to elucidate the mechanisms of color appearance. This episode in the history of color science remains important, however, because it reminds us that the phenomena of color appearance are very significant and very compelling, enough so to motivate Land to challenge whether the color establishment had answered the right questions. As to Land's additional and extraordinary claim in those papers, that there are two, not three, types of photoreceptors, well, we all have off days.¹

Color Constancy: Theory

If the absolute rates of photopigment absorptions do not explain color appearance, what does? The illusions in Color Plate 2 and Figure 9.1 both suggest that color appearance is related to the *relative* cone absorption rates. Within an image, bright objects generate more cone absorp-

¹ Of course, even on his off days, Land was worth a billion dollars.

tions than dark objects; red objects create more L-cone absorptions, and blue objects create more S-cone absorptions. Hence, one square in Figure 9.1 appears light because it is associated with more cone absorptions than its neighboring region, while the other appears dark because it is associated with fewer cone absorptions. The relative absorption rate is very closely connected to the idea of the stimulus contrast that has been so important in this book. Color appearance depends more on the local contrast of the cone absorptions than on the absolute level of cone absorptions.

The dependence on relative, rather than absolute, absorption rates is a general phenomenon, not something that is restricted to a few textbook illusions. Consider a simple thought experiment that illustrates the generality of the phenomenon. Suppose you read this book indoors. The white part of the page reflects about 90 percent of the light toward your eye, while the black ink reflects only about 2 percent. Hence, if the ambient illumination inside a reading room is 100 units, the white paper reflects 90 units, and the black ink reflects 2 units. When you take the book outside, the illumination level can be 100 times greater, or 10,000 units. Outside, the black ink reflects 200 units toward your eye, which far exceeds the level of the white paper when you were indoors. Yet, the ink continues to look black. As we walk about the environment, then, we must constantly be inferring the lightness and color of objects by comparing the spatial pattern of cone absorptions.²

This thought experiment also illustrates that the color we perceive informs us mainly about objects. The neural computation of color is structured so that objects retain their color appearance whether we encounter them in shade or sun. When the color appearance of an object changes, we think that the object itself has changed. The defining property of an object is not the absolute amount of light it reflects, but rather how much light it reflects relative to other objects. From our thought experiment, it follows that the color of an object imaged at a point on the retina should be inferred from the relative level of cone absorptions caused by an object. To compute the relative level of cone absorptions, we must take into account the spatial pattern of cone absorptions, not just the cone absorptions at a single point.

On this view, color appearance is a mental explanation of why an object causes relatively more absorptions in one cone type than another object. The physical attribute of an object that describes how well the object reflects light at different wavelengths is called the object's **surface reflectance**. Objects that reflect mainly long-wavelength light usually appear red; objects that reflect mainly short-wavelength light usually appear blue. Yet, as we shall explore in the next few pages, interpreting the cone absorption rates in terms of the surface-reflectance functions is not trivial. How the nervous system makes this interpretation is an essential question in color appearance. A natural place

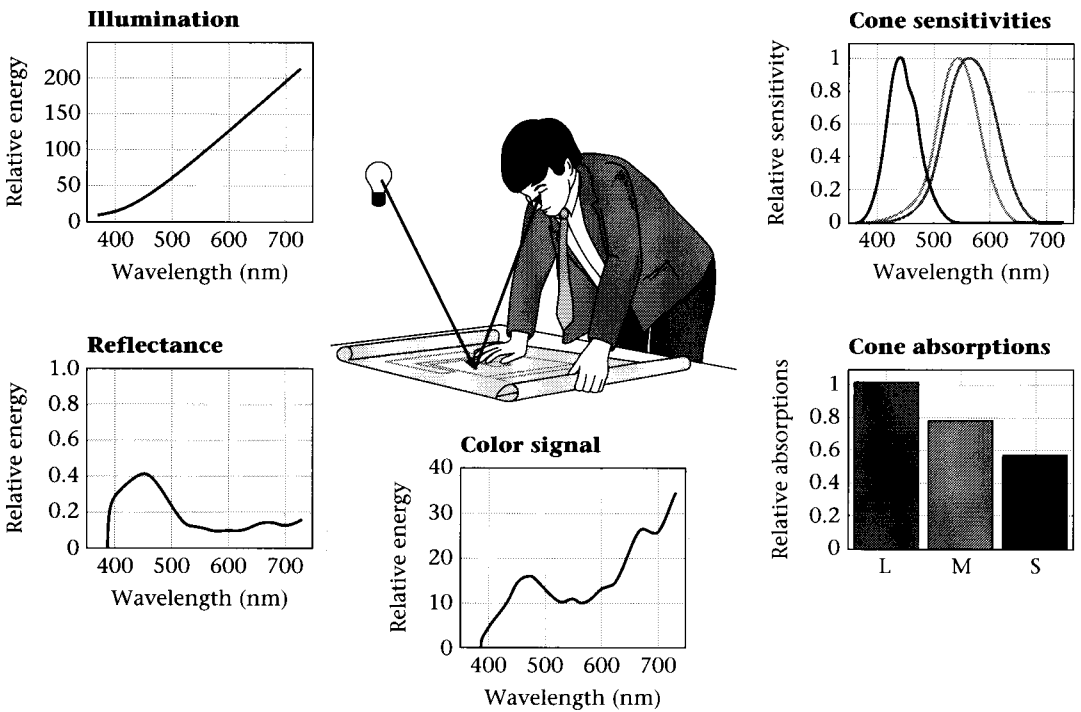
²This example was used by Hering (1964).

to begin our analysis of color appearance, then, is with the question, How can the central nervous system infer an object's surface-reflectance function from the mosaic of cone absorptions?

Spectral Image Formation

To understand the process of inferring surface reflectance from the light incident at our eyes, we must understand a little about how images are formed. The light incident at our corneas and absorbed by our cones depends in part on the properties of the objects that reflect the light and in part on the wavelength composition of the ambient illumination. We must understand each of these components, and how they fit together, to see what information we might extract from the retinal image about the surface-reflectance function. A very simple description of the imaging process is shown in Figure 9.2.


Ordinarily, image formation begins with a light source. We can describe the spectral properties of the light source in terms of the relative



9.2 A DESCRIPTION OF SPECTRAL IMAGE FORMATION. Light from a source arrives at a surface and is reflected toward an observer. The light at the observer is absorbed by the cones and ultimately leads to a perception of color. The functions associated with the imaging process, represented by the graphs, include (counterclockwise from upper left) the spectral power distribution of the light source; the surface-reflectance function of the object; the result of multiplying these two functions to create the color signal incident at the eye; and the cone absorptions caused by the incident signal, which in turn are dependent on the cone sensitivities.

amount of energy emitted at each wavelength, namely the spectral power distribution of the light source (see Chapter 4). The light from the source is either absorbed by the surface or reflected. The fraction of the light reflected by the surface defines the surface-reflectance function. As a first approximation, we can calculate the light reflected toward the eye by multiplying the spectral power distribution and the surface-reflectance function together. We will call this light the **color signal** because it serves as the signal that ultimately leads to the experience of color. The color signal leads to different amounts of absorptions in the three cone classes, and it is the interpretation of these cone absorptions by the nervous system that is the basis of our color perception.

In Chapter 4, I reviewed some properties of illuminants and sensors; but this is the first time we have considered the surface-reflectance function. It is worth spending time thinking about some of the properties of how surfaces reflect light. Figure 9.3 shows the reflectance function of four matte papers, that is, papers with a smooth, even surface free from shine or highlights. Because these curves describe the fraction of light reflected, they range between 0 and 1. While it is common to refer to an object as having a surface-reflectance function (as I have just done), in a certain sense, the notion of a surface-reflectance function is a ruse. If you look about the room you are in, you will probably see some surfaces that are shiny or glossy. As you move around, changing the geometrical relationship between yourself, the lighting, and the surface, the light reflected to your eye changes considerably. Hence, the tendency of the surface to reflect light toward your eye does not depend



9.3 THE SURFACE-REFLECTANCE FUNCTION

measures the proportion of light scattered from a surface at each wavelength. The panels show the surface-reflectance functions of various colored papers along with the color name associated with the paper. Surface reflectance is correlated with the color appearance; as Newton (1704) wrote "colours in the object are nothing but a disposition to reflect this or that sort of ray more copiously than the rest."

only on the surface; the light scattered to your eye also depends on the viewing geometry.³

The surface reflectance depends on viewing geometry because the reflected light arises from several different physical processes that occur when light is incident at the surface. Each of these processes contributes simultaneously to the light scattered from a surface, and each has its own unique properties. The full model describing reflectance appears to be complex, but Shafer (1985) has created a simple approximation of the reflection process, called the **dichromatic reflection model**, that captures several important features of surface reflectance. Figure 9.4A sketches the model, which applies to a broad collection of materials called dielectric⁴ surfaces (Shafer, 1985; Klinker et al., 1988; Nayar and Bolle, 1993, Wolff, 1994).

According to the dichromatic reflection model, dielectric material consists of a clear substrate with embedded colorant particles. One way light is scattered from the surface is by a mirror-like reflection at the interface of the surface. This process is called **interface reflection**. A second scattering process takes place when the rays enter the material. These rays are reflected randomly between the colorant particles. A fraction of the incident light is absorbed by the material, heating it up, and part of the light emerges. This process is called **body reflection**.


The spatial distributions of light scattered by these two mechanisms are quite different (Fig. 9.4B). Light scattered by interface reflection is quite restricted in angle, much as a mirror reflects incident rays. Conversely, the light scattered by body reflection emerges equally in all directions. When a surface has no interface reflections, but only body reflections, it is called a matte or Lambertian surface. Interface reflection is commonly called specular reflection and is the reason why some objects appear glossy.

The different geometrical distribution in how body and interface reflections are reflected is the reason why specular highlights on a surface appear much brighter than diffuse reflection. Nearly all of the specular scattering is confined to a small angle; the body reflection is divided among many directions. The interface reflections provide a strong signal, but because they can only be seen from certain angles they are not a reliable source of information. As the object and observer change their geometric relationship the specular highlight moves along the surface of the object, or it may disappear altogether.

For many types of materials, interface reflection is not selective for wavelength. The spectral power distribution of the light scattered at the interface is the same as the spectral power distribution of the incident

³ Also, some types of materials fluoresce, which is to say they absorb light at one wavelength and emit light at another (longer) wavelength. This is also a linear process, but too complex to consider in this discussion.

⁴ Dielectrics are nonconducting materials.



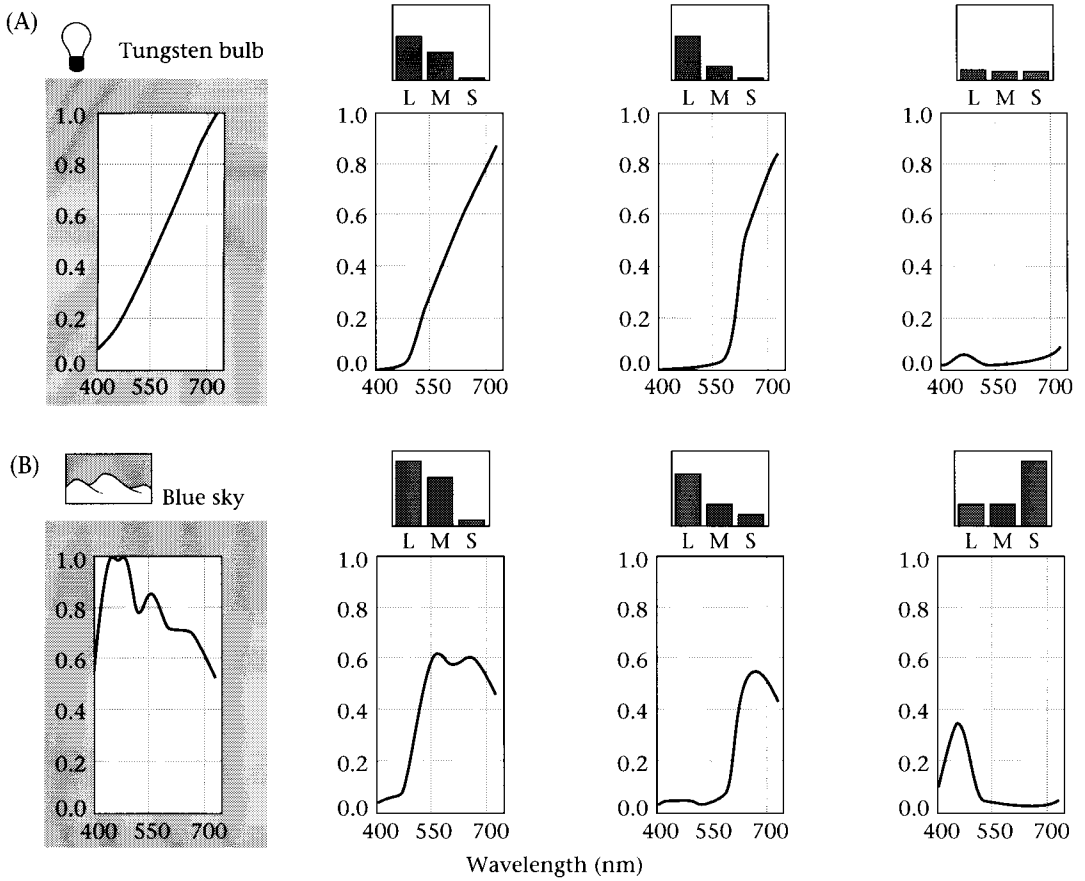
9.4 THE DICHROMATIC REFLECTION MODEL OF SURFACE REFLECTANCE IN INHOMOGENEOUS MATERIALS. (A) Light is scattered from a surface by two different mechanisms. Some incident light is reflected at the interface (interface reflection). Other light enters the material, interacts with the embedded particles, and then emerges as reflected light (body reflection). (B) Interface reflection is likely to be concentrated in one direction. Body reflection emerges with nearly equal likelihood in many different directions. Because interface reflections are concentrated in certain directions, light reflected by this process can be much more intense than light reflected by body reflection.

light. This is why specular highlights take on the color of the illumination source. Body reflection, on the other hand, does not return all wavelengths uniformly. The particles in the medium absorb light selectively, and it is this property that distinguishes objects in terms of their color appearance. Ordinarily, when people refer to the surface reflectance of an object, they mean to refer to the body reflection of the object.

We can describe the reflection of light by a matte surface with a simple mathematical formula. Suppose that the illuminant spectral power distribution is $e(\lambda)$. We suppose that the body reflectance is $s(\lambda)$. Then the color signal, that is, the light arriving at the eye, is

$$c(\lambda) = s(\lambda)e(\lambda) \quad (9.1)$$

Figure 9.5 shows several examples of the light reflected from matte surfaces. The shaded graph in Figure 9.5A is the spectral power distribution of an illuminant similar to a tungsten bulb. The other panels show the



9.5 THE LIGHT REFLECTED FROM OBJECTS CHANGES AS THE ILLUMINANT

CHANGES. (A) The shaded panel on the left shows the spectral power distributions of a light source similar to a tungsten bulb. The three graphs on the right show the light reflected from the yellow, red, and blue papers in Figure 9.3 when illuminated by this source. The bar plots above the graphs show the three cone absorption rates caused by the color signal. (B) When the light source is similar to blue sky, as in the shaded panel on the left, the light reflected from the same papers is quite different.

spectral power distribution of light that would be reflected from the yellow, red, and blue papers in Figure 9.3. The shaded graph in Figure 9.5B shows the spectral power distribution similar to blue sky illumination and the light reflected from the same three papers. Plainly, the spectral composition of the reflected light changes when the illuminant changes.

We can calculate the cone absorption rates caused by each of these color signals (see Chapter 4). These rates are shown in the bar plots above the individual graphs of reflected light in Figure 9.5. By comparing the bar graphs in parts A and B of the figure, you can see that the cone absorption rates from each surface change dramatically when the illumination changes. This observation defines a central problem in

understanding color appearance. Color is a property of objects; but the reflected light and thus the cone absorptions vary with the illumination. If color must describe a property of an object, the nervous system must *interpret* the mosaic of photopigment absorptions and estimate something about the surface-reflectance function. This is an estimation problem. How can the nervous system use the information in the cone absorptions to infer the surface-reflectance function?

Surface-Reflectance Estimation

There are some very strong limitations on what we can achieve when we set out to estimate surface reflectance from cone absorptions. First, notice that the color signal depends on two spectral functions that are continuous functions of wavelength: the spectral power distribution of the ambient illumination and the surface-reflectance function. The light incident at the eye is the product of these two functions; so any illuminant and surface combination that produces this same light will be indistinguishable to the eye.

One easy mathematical way to see why this is so is to consider the color signal. Recall that the color signal is equal to the product of the illuminant spectral power distribution $e(\lambda)$ and the surface-reflectance function, $s(\lambda)$:

$$c(\lambda) = s(\lambda)e(\lambda) \quad (9.2)$$

Suppose we replace the illuminant with a new illuminant, $f(\lambda)e(\lambda)$, and replace all of the surfaces with new functions, $s(\lambda)/f(\lambda)$. This change has no effect on the color signal,

$$c(\lambda) = [s(\lambda)/f(\lambda)][f(\lambda)e(\lambda)] = s(\lambda)e(\lambda) \quad (9.3)$$

and thus no effect on the photopigment absorption rates. Hence, there is no way for the visual system to discriminate between these two illuminant and surface pairs.

Now, consider a second limitation to the estimation problem. The visual system does not measure the spectral power distribution directly. Rather, the visual system only encodes the absorption rates of the three different cone types, so the nervous system cannot be certain which of the many metameric spectral power distributions is responsible for causing the observed cone absorption rates (see Chapter 4). This creates even more uncertainty for the estimation problem.

In the introduction to this part of the book, I quoted Helmholtz's suggestion: the visual system imagines those objects being present that could give rise to the retinal image. We now find that the difficulty we have in following this advice is not that there are no solutions, but

rather than there are too many. We encode so little about the color signal that many different objects could all have given rise to the retinal image.

Which of the many possible solutions should we select? The general strategy we should adopt is straightforward: pick the most likely one. Will this be a helpful estimate, or are there so many likely signals that encoding the most likely one is hardly better than guessing with no information?

Perhaps the most important point that we have learned from color-constancy calculations over the last ten years is this: the set of surface and illuminant functions we encounter is not so diverse as to make estimation from the cone catches useless. Some surface-reflectance functions and some illuminants are much more likely than others, and even with only three types of cones, it is possible to make educated guesses that do more good than harm.

The surface-reflectance estimation algorithms we will review are all based on this principle. They differ only in the set of assumptions they make concerning what the observer knows and what we mean by “most likely.” I review them now because some of the tools are useful and interesting, and some of the summaries of the data are very helpful in making practical calculations and experiments.

Linear Models

To estimate which lights and surfaces are more probable, we need to do two things. First, we need to measure the spectral data from lights and surfaces. Second, we need a way to represent the likelihood of observing different surface and illuminant functions.

Since the early part of the 1980s, linear models of surface and illuminant functions have been used widely to represent our best guess about the most likely surface and illuminant functions. A linear model of a set of spectral functions, such as surface reflectances, is a method of efficiently approximating the measurements. There are several ways to build linear models, including **principal-components analysis**, **centroid analysis**, or **one-mode analysis**. These methods have much in common, but they differ slightly in their linear model formulation and error measures (Judd et al., 1964; Cohen, 1964; Maloney, 1986; Marimont and Wandell, 1993).

As an example of how to build a linear model, we will review a classic paper by Judd, Macadam, and Wyszecki (1964). These authors built a linear model of an important source of illumination, daylight spectral power distributions. They collected more than 600 measurements of the spectral power distribution of daylight at different times of day, under different weather conditions, and on different continents. To measure the spectral power distribution of daylight we place an object with a known reflectance outside. It is common to use blocks of pressed magnesium oxide as a standard object because such blocks reflect light of all wavelengths nearly equally. Moreover, the material is essentially a pure


diffuser: a quantum of light incident on the surface from any angle is reflected back with equal probability in all other directions above the surface.

Since they were interested in the relative spectral composition, not the absolute level, Judd et al. normalized their measurements so that they were all equal to the value 100 at 560 nm. Their main interest was in the wavelength regime visible to the human eye, so they made measurements roughly from 400 nm to 700 nm. Their measurements were spaced every 10 nm. Hence, they could represent each daylight measurement by a set of 31 numbers. Three representative daylight spectral power distributions, normalized to coincide at 560 nm, are shown in Figure 9.6.

The data plotted in Figure 9.6 show that measured daylight relative spectral power distributions can differ depending on the time of day and the weather conditions. But, after examining many daylight functions, Judd et al. found that the curves do not vary wildly and unpredictably; the data are fairly regular. Judd et al. captured the regularities in the data by building a linear model of the observed spectral power distributions. They designed their linear model, a principal-components model, using the following logic.

First, they decided that they wanted to approximate their observations in the **squared-error** sense. That is, suppose $e(\lambda)$ is a measurement, and $\hat{e}(\lambda)$ is the linear-model estimate of the measurement. Then, they decided to select the approximation in order to minimize the squared error

$$\sum_{\lambda} [e(\lambda) - \hat{e}(\lambda)]^2$$



9.6 THE RELATIVE SPECTRAL POWER DISTRIBUTION OF DAYLIGHT UNDER THREE TYPICAL CONDITIONS. The curves drawn here are typical daylight measurements; the curves are normalized to coincide at 560 nm. Source: Judd et al., 1964.

When we consider the collection of observations as a whole, the function that approximates the entire data set with the smallest squared error is the mean. The mean observation from Judd et al.'s data set, $e_0(\lambda)$, is the solid curve in Figure 9.7.

Once we know the mean, we need only approximate the difference between the mean and each individual measurement. We build the linear model to explain these differences as follows. First, we select a fixed set of **basis functions**. Basis functions, like the mean, are descriptions of the measurements. We approximate a measurement's difference from the mean as the weighted sum of the basis functions. For example, suppose $\Delta e_j(\lambda)$ is the difference between the j th daylight measurement and the mean daylight. Further, suppose we select a set of N basis functions, $E_i(\lambda)$. Then, we approximate the differences from the mean as the weighted sum of these basis functions, namely,

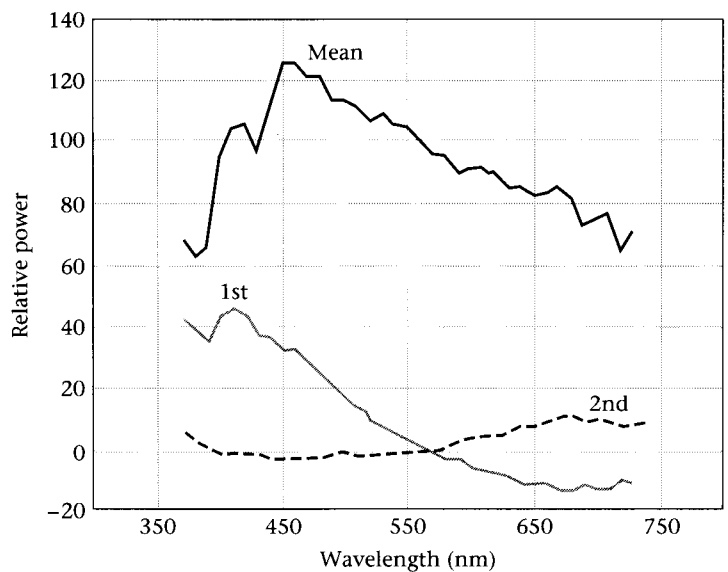
$$\Delta e_j(\lambda) \approx \sum_{i=1}^{i=N} \omega_i E_i(\lambda) \quad (9.4)$$

The basis functions are chosen to make the sum of the squared errors between the *collection* of measurements and their approximations as small as possible. The number of basis functions, N , is called the **dimension** of the linear model. The values ω_i are called the linear-model **weights** or **coefficients**. They are chosen to make the squared error between an *individual* measurement and its approximation as small as possible. They serve to describe the properties of the specific measurement.

As the dimension of the linear model increases, the precision of the

9.7 A LINEAR MODEL FOR DAYLIGHT SPECTRAL POWER DISTRIBUTIONS.

The curve labeled "Mean" is the mean spectral power distribution of a set of daylights whose spectral power distributions were normalized to a value of 100 at 560 nm. The curves labeled "1st" and "2nd" show the two basis curves used to define a linear model of daylights. By adding together the mean and weighted sums of the two basis functions, one can generate examples of typical relative spectral power distributions of daylight. Source: Judd et al., 1964.



linear-model approximation improves. The dimension one chooses for an application depends on the required precision.⁵

Judd et al. found an excellent approximation of the daylight measurements using the mean and two basis functions. The linear-model representation expresses the measurements efficiently. The mean and two basis functions are fixed, side conditions of the model. Their linear-model approximation of each measurement uses only two weights. The empirical results have been confirmed by other investigators, and the results have been adopted by the international color standards organization to create a model of daylights (Sastri and Das, 1965; Dixon, 1978). The mean and two basis terms from the international standard are plotted in Figure 9.7.

Because daylight varies in absolute spectral power distributions, not just in the relative distribution, we should extend Judd et al.'s linear model to a three-dimensional linear model that includes absolute intensity. A three-dimensional linear model we might use consists of the mean and the two derived curves. In this case the three-dimensional linear model approximation becomes:⁶

$$e(\lambda) = \sum_{i=1}^3 \omega_i E_i(\lambda) \quad (9.5)$$

We can express the linear model in Equation 9.5 as a matrix equation, $\mathbf{e} = \mathbf{B}_e \boldsymbol{\omega}$ in which \mathbf{e} is a vector representing the illuminant spectral power distribution, the three columns of \mathbf{B}_e contain the basis functions $E_i(\lambda)$, and $\boldsymbol{\omega}$ is a three-dimensional vector containing the linear-model coefficients ω_i .

We can see why linear models are efficient by writing Equation 9.4 in matrix form:

$$\begin{pmatrix} \cdot \\ \cdot \\ \cdot \\ e(\lambda) \\ \cdot \\ \cdot \\ \cdot \end{pmatrix} = \begin{pmatrix} \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot \\ E_1(\lambda) & E_2(\lambda) & E_3(\lambda) \\ \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot \\ \cdot & \cdot & \cdot \end{pmatrix} \begin{pmatrix} \omega_1 \\ \omega_2 \\ \omega_3 \end{pmatrix}$$

⁵The basis functions that minimize the squared error can be found in several ways, most of which are explained in widely available statistical packages. If the data are in the columns of a matrix, one can apply the singular-value decomposition to the data matrix and use the left singular vectors. Equivalently, one can find the eigenvectors of the covariance matrix of the data.

⁶It is possible to improve on this model slightly, but as a practical matter these three curves do quite well as basis functions.

The single spectral power distribution, the vector on the left, consists of measurements at many different wavelengths. The linear model summarizes each measurement as the weighted sum of the basis functions, which are the same for all measurements, and a few weights, $\boldsymbol{\omega}$, which are unique to each measurement. The linear model is efficient because we represent each additional measurement using only three weights, $\boldsymbol{\omega}$, rather than the full spectral power distribution.

Simple Illuminant Estimation

Efficiency is useful, but if efficiency were our only objective we could find more efficient algorithms. The linear models are also important because they lead to very simple estimation algorithms. As an example, consider how we might use a device with three color sensors, like the eye, to estimate the spectral power distribution of daylight. Such a device is vastly simpler than the spectroradiometer Judd et al. needed to make many measurements of the light.

Suppose we have a device with three color sensors, whose spectral responsivities are, say, $R_i(\lambda)$, $i = 1, 2, 3$. The three sensor responses will be:

$$\begin{aligned} r_1 &= \sum_{\lambda} R_1(\lambda)e(\lambda) \\ r_2 &= \sum_{\lambda} R_2(\lambda)e(\lambda) \\ r_3 &= \sum_{\lambda} R_3(\lambda)e(\lambda) \end{aligned} \quad (9.6)$$

We can group these three linear equations into a single matrix equation

$$\mathbf{r} = \mathbf{S}\mathbf{e} \quad (9.7)$$

where the column vector \mathbf{r} contains the sensor responses, the rows of the matrix \mathbf{S} are the sensor spectral responsivities, and \mathbf{e} is the illuminant spectral power distribution.

Before Judd et al.'s study, one might have thought that three sensor responses are insufficient to estimate the illumination. But from their data we have learned that we can approximate \mathbf{e} with a three-dimensional linear model, $\mathbf{e} \approx \mathbf{B}_e\boldsymbol{\omega}$. This reduces the equation to

$$\mathbf{r} \approx (\mathbf{S}\mathbf{B}_e)\boldsymbol{\omega} \quad (9.8)$$

The matrix $\mathbf{S}\mathbf{B}_e$ is 3×3 , and its entries are all known. The sensor responses, \mathbf{r} , are also known. The only unknown is $\boldsymbol{\omega}$. Hence, we can estimate, $\boldsymbol{\omega}$, and use these weights to calculate the spectral power distribution, $\mathbf{B}_e\boldsymbol{\omega}$.

This calculation illustrates two aspects of the role of linear models. First, linear models represent a priori knowledge about the likely set of inputs. Using this information permits us to convert underdetermined linear equations (Equation 9.7) into equations we can solve (Equation 9.8). Linear models are a blunt but useful tool for representing probabilities. Using linear models, it becomes possible to use measurements from only three color sensors to estimate the full relative spectral power distribution of daylight illumination.

Second, linear models work smoothly with the imaging equations. Since the imaging equations are linear, the estimation methods remain linear and simple.

Surface Reflectance Models

Daylight represents an important class of signals for vision. For most of the history of the earth, daylight was the only important light source. There is no similar set of surface-reflectance functions. I was reminded of this once by the brilliant color scientist, G. Wyszecki. When I was just beginning my study of these issues, I asked him why he had not undertaken a study of surfaces similar to those in the daylight study. He shrugged at me and answered, "How do you sample the universe?"

Wyszecki was right, of course. The daylight measurement study could begin and end in a single paper. There is no specific set of surfaces that is of equal importance to daylight, so we have no way to perform a similar analysis on surfaces. But, there are two related questions we can investigate. First, we can ask what the properties are of certain collections of surfaces that are of specific interest to us, say, for practical applications. Second, we can ask what the visual system, with only three types of cones, can infer about surfaces.

Over the years linear models for special sets of materials have been used in many applications. Printer and scanner manufacturers may be interested in the reflectance functions of inks. Computer graphics programmers may be interested in the reflectance factors of geological materials, or tea pots. Color scientists have repeatedly measured the reflectance functions of standard color samples used in industry.⁷ These cases can be of practical value and interest in printing and scanning applications (e.g., Drew and Funt, 1992; Marimont and Wandell, 1992; Farrell and Wandell, 1993; Vrhel and Trussell, 1994). To discover the regularities in surface functions, then, we should measure the body reflection terms. From the studies that have taken place over the last several years, it has become increasingly clear that in the visible wavelength region the surface-reflectance functions tend to be quite smooth and thus exhibit a great deal of regularity. Hence, linear models serve to describe the reflectance functions quite well.

⁷ Such as the Munsell chips, which are described later in this chapter.

For example, Cohen (1964), Maloney (1986), and Parkkinen et al. (1989) studied the reflectance properties of a variety of surfaces including special samples and some natural objects. For each of the sets studied by these authors, the data can be modeled nearly perfectly using a linear model with fewer than six dimensions. Excellent approximations, though not quite perfect, can be obtained by three-dimensional approximations.

As an example, I have built a linear model to approximate a small collection of surface-reflectance functions for a color target, the **Macbeth ColorChecker**, which is used widely in industrial applications. The target consists of 24 square patches laid out in a 4×6 array. The surfaces in this target were selected to have reflectance functions similar to a range of naturally occurring surfaces. They include reflectances similar to human skin, flora, and other materials (McCamy et al., 1976). To create the linear model, I measured the surface-reflectance functions of these patches with a spectral radiometer in my laboratory. The original data set, then, consisted of measurements from 370 nm to 730 nm in 1-nm steps for each of the 24 patches. Then, using conventional statistical packages, I calculated a three-dimensional linear model to fit all of these surface-reflectance functions. The linear model basis functions, $S_i(\lambda)$, were selected to minimize the squared error:⁸

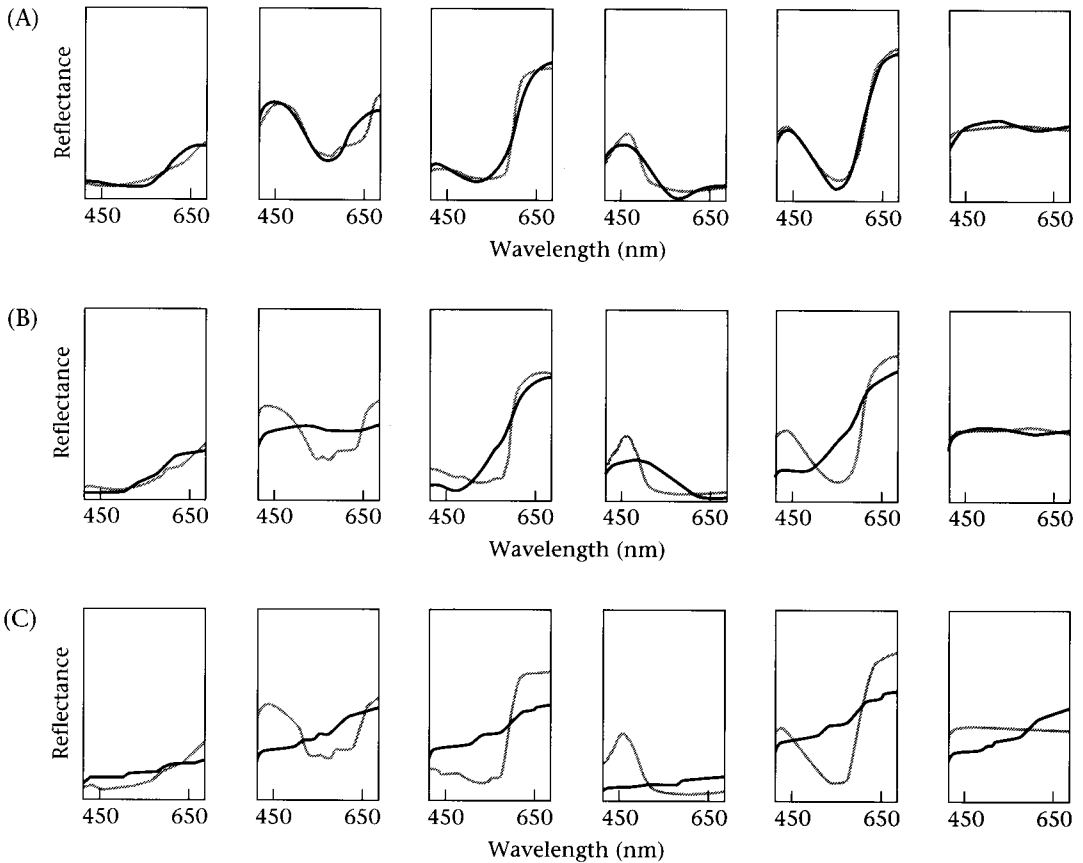
$$\left[s(\lambda) - \sum_{i=1}^N \sigma_i S_i(\lambda) \right]^2 \quad (9.9)$$

The values σ_i are called the **surface coefficients**, and we will represent them as a vector, $\boldsymbol{\sigma} = (\sigma_1, \dots, \sigma_N)$. There are fewer surface coefficients than data measurements. If we create a matrix whose columns are the basis functions, \mathbf{B}_s , then we can express the linear model approximation as $\mathbf{B}_s \boldsymbol{\sigma}$.

The shaded lines in Figure 9.8 show the reflectance functions of six of the 24 surfaces. The solid curves within each row of the figure contain the approximations using linear models with different dimensionality. The bottom row shows a one-dimensional linear model; in this case the approximations are scaled copies of one another. As we increase the dimensionality of the linear model the approximations become very similar to the originals. In the three-dimensional model, the approximations are quite close to the true functions.

The low-dimensional linear model approximates these surface-reflectance functions because the functions vary smoothly as a function of wavelength. The linear model consists of a few, slowly varying basis functions, as shown in Figure 9.9. The first basis function captures the light–dark variation of the surfaces. The second basis function captures a red–green variation, and the third captures a blue–yellow variation.

⁸I calculated the singular value decomposition of the matrix whose columns consist of the surface reflectance vectors. I used the left singular vectors as the basis functions.




9.8 A LINEAR MODEL TO APPROXIMATE THE SURFACE REFLECTANCES IN THE MACBETH COLORCHECKER. The panels in each row of this figure show the surface-reflectance functions of six colored surfaces (shaded lines) and the approximation to these functions using a linear model (solid lines). The approximations using linear models with (A) three, (B) two, and (C) one dimension are shown.

Although the approximations are quite good, there are differences between the surface-reflectance functions and the three-dimensional linear model. We might ask whether these differences are visually salient. This question is answered by the renderings of these surface approximations in Color Plate 4. The one-dimensional model looks like a collection of surfaces in various shades of gray. For the blue-sky illumination used in the rendering, linear models with three or more dimensions appear very similar to a rendering using the complete set of data.

Sensor-Based Error Measures

I have described linear models that minimize the squared error between the approximation and the original spectral function. As the last analysis showed, however, when we choose linear models to minimize the



9.9 BASIS FUNCTIONS OF THE LINEAR MODEL FOR THE MACBETH

COLORCHECKER. The surface-reflectance functions in the collection vary smoothly with wavelength, as do the basis functions. The first basis function is all positive and explains the most variance in the surface-reflectance functions. The basis functions are ordered in terms of their relative significance for reducing the error in the linear-model approximation to the surfaces.


spectral error we are not always certain whether we have done a good job in minimizing the visual error. In some applications, the spectral error may not be the objective function that we care most about minimizing. When the final consumer of the image is a human, we may only need to capture that part of the reflectance function that is seen by the sensors.

If we are modeling reflectance functions for a computer graphics application, for example, there is no point in modeling the reflectance function at 300 nm, since the human visual system cannot sense light in that part of the spectrum anyway. For these applications, we should be careful to model accurately those parts of the function that are most significant for vision. In these cases, one should select a linear model of the surface reflectance by minimizing a different error measure, one that takes into account the selectivity of the eye.

The procedure for creating linear models that are appropriate for the eye is described in Marimont and Wandell (1992). That study considers how to define linear models that minimize the root-mean-squared




1 In its regenerated form, the rod photopigment rhodopsin appears purple. (Top) The pigment can be seen easily in light reflected from the eyes of an alligator that had been kept in the dark for several hours. (Bottom) After exposure to light, rhodopsin bleaches and becomes transparent; only the alligator's reflective white tapetum is visible after light exposure. From Rushton, 1962.




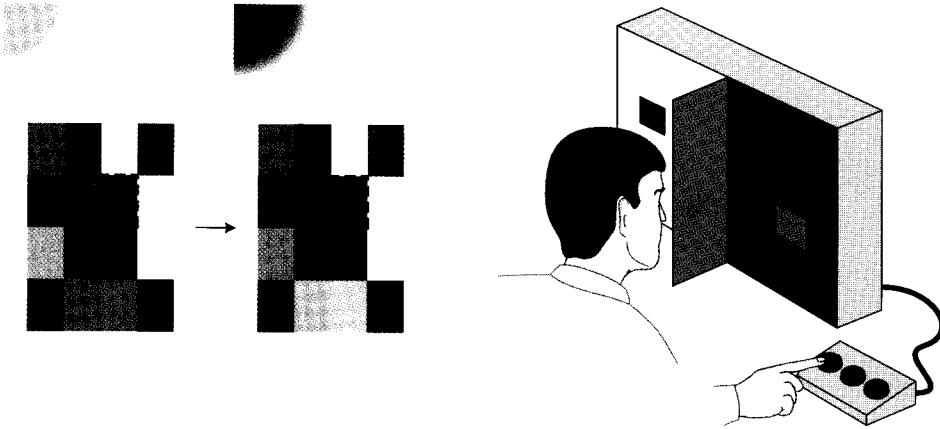
2 This collage by Josef Albers shows that the results of the color-matching experiment cannot explain color appearance. Although the light reflected from the two "X" shapes have the same spectral power distribution (as can be seen at the top of the figure where the X's join), and thus have the same effect on the photopigments, their appearance differs. This illustration demonstrates that the photopigment responses in a small region do not determine color appearance in that region; rather, color appearance depends on the spatial structure of the image as a whole. From Albers, 1975.

3 Our visual experience of color depends on the spatiotemporal frequency range of the stimulus. The graph divides the spatiotemporal plane into three regions corresponding to three basic color appearance regimes. At low spatiotemporal frequencies, color perception is trichromatic: we perceive blue–yellow, red–green, and light–dark variations in contrast. At moderate spatiotemporal frequencies, color perception is dichromatic: we fail to perceive blue–yellow contrast variations. At high spatiotemporal frequencies, color perception is monochromatic: we perceive all image contrast as light–dark modulations of the mean color.




4 Linear-model approximations to the Macbeth ColorChecker are shown rendered under a blue-sky illumination; the linear model dimension of each approximation is listed above the images. In the one-dimensional approximation, the surfaces appear monochromatic, varying only in lightness. As the dimension of the linear-model increases, the approximations appear more similar to the original. For these surfaces, this illuminant, and this method of creating a linear model, the original is visually indistinguishable from the approximations when we use four or more linear-model dimensions.








5 The memory-matching (left) and haploscopic (right) methods for measuring asymmetric color matches. In the memory-matching method, the subject sees a target under one illumination, remembers its appearance, and then identifies a matching target after adapting to a second illumination. In the haploscopic method, the observer adapts each eye separately and makes a simultaneous appearance match. An important feature of the haploscopic apparatus depicted is that the bipartite test field is divided by a perpendicular, two-sided mirror that enlarges the size of the background presented to each eye. If a black dividing screen is used instead of a mirror, shifting the gaze between right and left visual fields causes significant changes in the mean illumination.



6 These two square-wave patterns were printed using the same ink, and hence the colored bars reflect the same light toward your eyes. The color appearance differs, however, because they are printed at different spatial frequencies. At many viewing distances, the thin bars on the left appear relatively green while the thick bars on the right appear relatively blue. You can also influence the color appearance of the bars by modifying the temporal frequencies of the two patterns: move the book up and down rapidly while viewing the patterns. From Wandell, 1993.



7 The apparent spatial sharpness, or focus, of a color image depends mainly on the light–dark components of the image rather than the opponent-colors components. Using a digital representation, the image at top left was decomposed into three separate red–green, blue–yellow, and light–dark contrast images. Three new versions of the original were made by blurring each of the three contrast images in turn and then combining the contrast images. The results of blurring the red–green (top right), blue–yellow (bottom left), or light–dark (bottom right) contrast image are shown. Although the physical defocus of each contrast image prior to recombination was the same, only the reproduction with a blurred light–dark contrast image appears defocused. Courtesy of Hagit Hel-Or.

error in the photopigment absorption rates, rather than the root-mean-squared error of the spectral curves. The method is called one-mode analysis. For many applications, the error measure minimized by one-mode analysis is superior to root-mean-squared error of the spectral curves.

Surface and Illuminant Estimation Algorithms

There is much regularity in daylight and surface functions, so it makes sense to evaluate how well we estimate spectral functions from sensor responses. Estimation algorithms rely on two essential components.

First, we need a method of representing our knowledge about the likely surface and illuminant functions. For example, linear models can be used to encode our a priori knowledge.⁹ Second, all modern estimation methods assume that the illumination varies either slowly or not at all across the image. This assumption is important because it means that the illumination adds very few new parameters to estimate.

Consider an example of an image with p points. We expect to obtain three cone absorption rates at each image point, so there are $3p$ measurements. If we can use a surface three-dimensional model for the surfaces, then there are $3p$ unknown surface coefficients; and there is a linear relationship between the measurements and the unknown quantities. If the illuminant is known, then the problem is straightforward to solve.

Additional unknown illuminant parameters make the problem a challenge. If the illuminant can vary from point to point, there will be $6p$ unknown parameters and the mismatch between known and unknown parameters will be very great. However, if the illuminant is constant across the image, we only have three additional parameters. In this case, by making some modest assumptions about the image, we can find ways to infer these three parameters and then proceed to estimate the surface parameters.

Modern estimation algorithms work by finding a method to overcome the mismatch between the measurements and the unknowns. We can divide existing estimation algorithms into two groups. The majority of estimation algorithms infer the illumination parameters by making one additional assumption about the image contents. For example, suppose we know the reflectance function of one surface. Then, we can use the sensor responses from that surface to measure the illuminant parameters. Knowing the reflectance function of one surface in the image compensates for the three unknown illuminant parameters. There are several implementations of this principle. The most important is the assumption that the average of all the surfaces in the image is gray, which

⁹ More sophisticated methods are based on using Bayesian estimation as part of the calculation. For example, see D'Zmura and Iverson (1993) and Freeman and Brainard (1994).

is called the **gray-world assumption** (Buchsbbaum, 1980; Land, 1986). Other algorithms are based on the assumption that the brightest surface in the image is a **uniform perfect reflector** (McCann et al., 1977). An interesting variant on both of these assumptions is the idea that we can identify **specularities** or glossy surfaces in the image. Since specularities reflect the illumination directly, often without altering the illuminant's spectral power distribution, the sensor responses to glossy surfaces provide information about the illuminant (D'Zmura and Lennie, 1986; Lee, 1986; Tominaga and Wandell, 1989, 1990).

A second group of estimation algorithms compensates for the mismatch in measurements and parameters by suggesting ways to acquire more data. For example, Maloney and Wandell (1986; Wandell, 1987) showed that by adding a fourth sensor (at three spatial locations), one can also estimate the surface and illuminants. D'Zmura and Iverson (1993a, b) explored an interesting variant of this idea. They asked what information is available if we observe the same surface under several illuminants. Pooling information about the same object seen under different illuminants is much like acquiring additional information from extra sensors (e.g., see Wandell, 1987).

Illuminant Correction: An Example Calculation

Before returning to experimental measurements of color appearance, we will perform a calculation that is of some practical interest as well as of some interest in understanding how the visual system might compensate for illumination changes. By working this example, we will start to consider what neural operations might permit the visual pathways to compensate for changes in the ambient illumination.

First, we will consider the general expression that shows how the surface and illuminant functions combine to yield the cone absorption rates. The three equations for the three cone types (L, M, and S), are:

$$\begin{aligned} r_1 &= \sum_{\lambda} R_1(\lambda)E(\lambda)S(\lambda) \\ r_2 &= \sum_{\lambda} R_2(\lambda)E(\lambda)S(\lambda) \\ r_3 &= \sum_{\lambda} R_3(\lambda)E(\lambda)S(\lambda) \end{aligned} \quad (9.10)$$

Next, we replace the illuminant and surface functions in Equation 9.10 with their linear-model approximations. This yields a new relationship between the coefficients of the surface-reflectance linear model, σ , and the three-dimensional vector of cone absorptions, \mathbf{r} ,

$$\mathbf{r} \approx \Lambda_c \sigma \quad (9.11)$$

We call the matrix Λ_e that relates these two vectors the **lighting matrix**. The entries of this matrix depend upon the illumination, e . The ij th entry of the lighting matrix is

$$\sum_{\lambda} \left(\sum_k \omega_k E_k(\lambda) \right) R_i(\lambda) S_j(\lambda) \quad (9.12)$$

We can compute two lighting matrices (Equation 9.11) from the spectral curves we have been using as examples. For one lighting matrix, I used the mean daylight spectral power distribution, and for the other I used the spectral power distribution of a tungsten bulb. I used the linear model of the Macbeth ColorChecker for the surface basis functions and the Stockman and MacLeod cone absorption functions (see Fig. 3.3). The lighting matrix for the blue-sky illumination is

$$\begin{pmatrix} 591.48 & 223.05 & -643.05 \\ 477.62 & 376.93 & -564.48 \\ 267.61 & 487.46 & 350.39 \end{pmatrix}$$

and the lighting matrix for the tungsten bulb is

$$\begin{pmatrix} 593.45 & 168.37 & -646.71 \\ 445.79 & 312.79 & -564.33 \\ 152.46 & 278.51 & 185.47 \end{pmatrix}$$

Notice that the largest differences between the matrices are in the third row. This is the row that describes the effect of each surface coefficient on the S-cone absorptions. The blue-sky lighting matrix contains much larger values than the matrix for the tungsten bulb. This makes sense because that is the region of the spectrum where these two illuminant spectral power distributions differ the most (see Fig. 9.5).

Imagine the following way in which the visual system might compensate for illumination changes. Suppose that the cortical analysis of color is based upon the assumption that the illumination is always that of a blue sky. When the illumination is different from the blue sky, the retina must try to provide a neural signal that is similar to the one that would have been observed under a blue sky. What computation does the retina need to perform in order to transform the cone absorptions obtained under the tungsten-bulb illuminant into the cone absorptions that would have occurred under the blue sky?

The cone absorptions from a single surface under the two illuminants can be written as:

$$\mathbf{r} \approx \Lambda_e \boldsymbol{\sigma}$$

$$\mathbf{r}' \approx \Lambda'_e \boldsymbol{\sigma}$$

By inverting the lighting matrices and recombining terms, we find that the cone absorptions under the two illuminants should be related linearly as,

$$\mathbf{r} \approx \Lambda_e \Lambda_e'^{-1} \mathbf{r}' \quad (9.13)$$

Hence, we can transform the cone absorptions from a surface illuminated by the tungsten bulb into the cone absorptions of the same surface illuminated by the blue sky by the following:

$$\mathbf{r} = \begin{pmatrix} 0.8119 & 0.2271 & 0.0550 \\ -0.0803 & 1.1344 & 0.1282 \\ 0.0429 & -0.0755 & 1.8091 \end{pmatrix} \mathbf{r}' \quad (9.14)$$

The retina can compensate for the illumination change by linearly transforming the observed cone absorptions, \mathbf{r}' , into a new signal, \mathbf{r} .

Now, what does it mean for the retina to compute a linear transformation? The linear transformation consists of a simple series of multiplications and additions. For example, consider the third row of the matrix in Equation 9.14. This row defines how the new S-cone absorptions should be computed from the observed absorptions. When we write this transformation as a single linear equation we see that the observed and transformed signals are related as

$$S = 0.0429 L' + -0.0755 M' + 1.8091 S' \quad (9.15)$$

The transformed S-cone absorption is mainly a scaled version of the observed absorptions. Because the tungsten bulb emits much less energy in the short-wavelength part of the spectrum, the scale factor is larger than 1. In addition, to be absolutely precise, we should add in a small amount of the observed L-cone signal and subtract out a small amount of the observed M-cone signal. But, these contributions are relatively small compared to the contribution from the S-cones. In general, for each of the cone types, the largest contributions to the transformed signal are scaled copies of the same signal type.

In this matrix, and in many practical examples, the only additive term that is not negligible is the contribution of the M-cone response to the transformed L signal. As a rough rule, because the diagonal terms are much larger than the off-diagonal terms, we can obtain good first-order approximation to the proper transformation by simply scaling the observed cone absorptions (see Foster and Nascimento, 1994).

Compensating for the illumination change by a purely diagonal scaling of the cone absorptions is called the **von Kries coefficient law**. The correction is not as precise as the best linear correction, but it frequently provides a good approximation. As we shall see in the next section, the von Kries coefficient law describes certain aspects of human color-appearance measurements as well.

Color Constancy: Experiments

In woven and embroidered stuffs the appearance of colors is profoundly affected by their juxtaposition with one another (purple, for instance, appears different on white than on black wool), and also by differences of illumination. Thus embroiderers say that they often make mistakes in their colors when they work by lamplight, and use the wrong ones.

Aristotle, *Meteorologica*

Asymmetric Color-Matching Experiments

To formulate some ideas about how the visual pathways compute color appearance, we adopted the view that color appearance is a psychological estimate of the surface-reflectance function (body reflectance). By thinking about computational methods of estimating body reflectance, we have discovered how the cone absorptions from an object vary with illuminant changes. Finally, we have seen that it is possible to compensate approximately for these changes by applying a linear transformation to the cone absorptions.

From the computational analysis, we have discovered a good question to examine in experimental studies of color appearance: What is the relationship between the cone absorptions of objects that appear the same under different illuminants? The computational analysis suggests that the cone absorptions of lights with the same color appearance, but seen under different illuminants, are related by a linear transformation.

It is up to the experimentalist, then, to find a method to use for measuring the cone absorptions that correspond to the same color appearance when seen in different viewing contexts. Color Plate 5 illustrates two methods of making such measurements. Panel A shows a **memory-matching** method. In this method, the subject studies the color of a target that is presented under one illumination source and then must select a target that looks the same under a second illumination source. These measurements identify the stimuli, and thus the cone absorptions, of targets that appear the same under the two illuminants. The drawback of the method is that making such matches is very time-consuming because the subject must adapt to the two illumination sources completely, a process which can take two minutes or more.

Color Plate 5B shows a second method called **dichoptic matching**. In this experiment the observer views the two scenes simultaneously in different eyes. One eye is exposed to a large neutral surface illuminated by, say, a daylight lamp. The other eye is exposed to an equivalent large neutral surface illuminated by, say, a tungsten lamp. These surfaces define a background stimulus that is different for each eye. The two backgrounds fuse in appearance and appear as a single large background. To

establish the asymmetric color matches, the experimenter places a test object on top of the standard background seen by one eye. The observer selects a matching object from an array of choices seen by the second eye. The color-appearance mapping is defined by measuring the cone absorptions of the test and matching objects, usually small colored papers, seen under their respective illuminants.

The dichoptic method has the advantage that the matches may be set quickly, avoiding the tedious delays required for visual adaptation in memory matches. The method has the disadvantage of making the assumption that adaptation occurs independently, prior to binocular combination.¹⁰

The experimental methods shown in Color Plate 5 generalize the conventional color-matching experiment. These methods are called **asymmetric color-matching** because, unlike standard color-matching, in these experiments the matches are set between stimuli presented in different contexts. As we have already seen, because color appearance discounts estimated changes of the illumination, matches set in the asymmetric color-matching experiment are *not* cone-absorption matches. Rather, the observer is establishing a match at a more central site following the correction for the properties of the scene.

The asymmetric color-matching experiment is directly relevant to the questions raised by our computational analysis of color appearance. Moreover, the experiment has a central place in the study of color appearance simply for practical experimental reasons. There are many general questions we might ask about color appearance. For example, we would like to be able to measure which colors are similar to one another; which colors have a common hue, saturation, or brightness; and so forth. If we had to study these questions separately under each illuminant, the task would be overwhelming. By beginning with the asymmetric color-matching experiment, we can divide color appearance measurements into two parts and reduce our experimental burden. The asymmetric matches define a mapping between the cone absorptions of objects with the same color appearance seen under different illuminants. From these experiments, we learn how to convert a target seen under one illuminant into an equivalent target under a standard illuminant. This transformation saves a great deal of experimental effort since we can focus most of our questions about color appearance on studies using just one standard illuminant.

¹⁰ This binocular method makes sense if one accepts the view that the adjustment for the illumination is mediated primarily before the signals from the two eyes are combined, in the superficial layers of area V1. The coherence of the experimental method can be tested psychophysically by examining transitivity. The observer matches a test on backgrounds L and R_1 , and then on backgrounds L and R_2 . The experimenter then places R_1 in the left eye and R_2 in the right eye and verifies that the matching lights match one another. There is no guarantee, of course, that these measurements are governed by precisely the same visual mechanisms that govern adaptation under normal viewing conditions.

The Linearity of Asymmetric Color Matches

We have seen measurements of superposition to test linearity throughout this volume. Tests of linearity in asymmetric color-matching appear very early in the color-appearance literature. When von Kries (1902, 1905) introduced the coefficient law, he listed several testable empirical results. Among the predictions, he listed the basic test of linearity, namely,

. . . there exist several very simple laws, which also appear to be specially adapted for experimental test. Namely, it must be that if L_1 on one retinal region causes the same result as L_2 on another, and similarly M_1 , working on the first, causes the same effect as M_2 on the other, in every case also $L_1 + M_1$ must have here the same effect as $L_2 + M_2$ there.

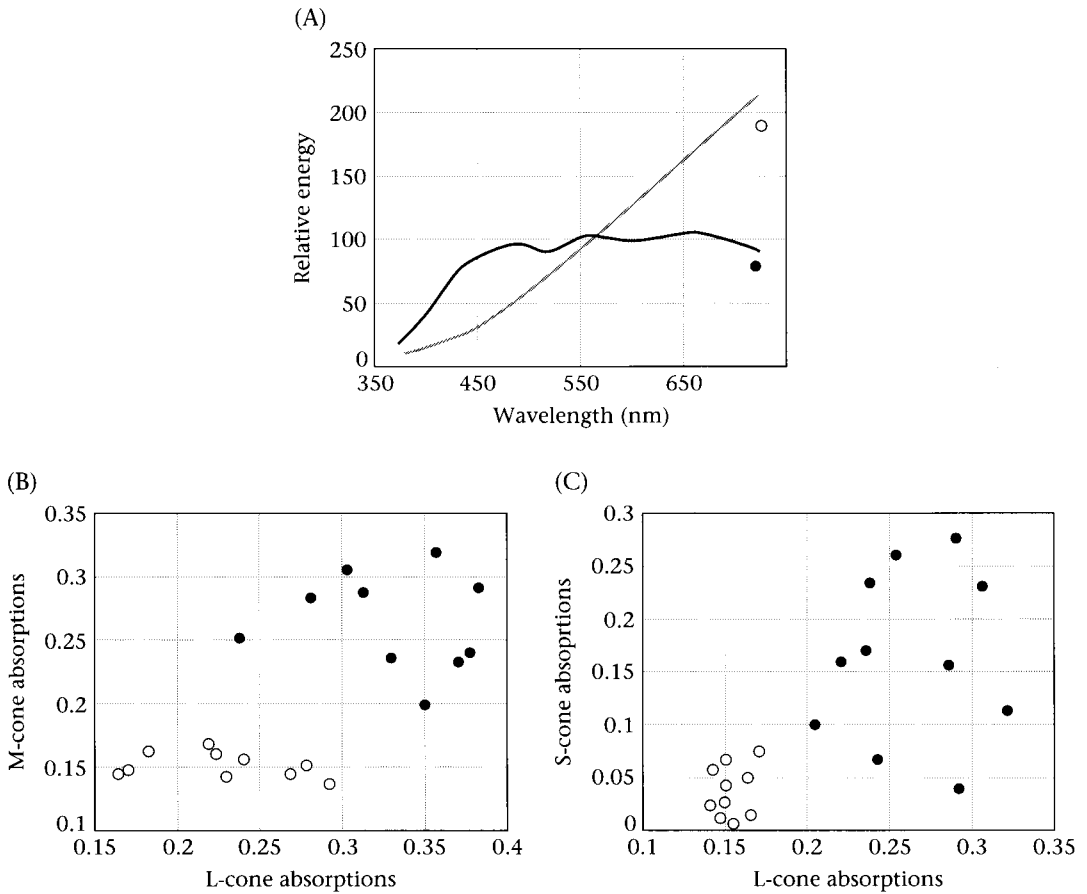
. . . The extended studies of Wirth (1900–1903) show that the law can be considered as nearly valid for reacting lights that are not too weak.

Evidently, von Kries not only raised the question of linearity of asymmetric color-matching, but by 1905 he considered it answered affirmatively.

While von Kries considered the question settled, not everyone was persuaded. Over the years, there have been many separate experimental tests of linearity in the asymmetric color-matching experiment. I am particularly impressed by a series of papers by Elaine Wassef, working first in London and then at the University College for Girls in Cairo. Wassef wrote at roughly the same time E. H. Land was working at Polaroid. In her papers, she reports on new studies and provides a review of the experimental tests of asymmetric color-matching linearity (Wassef, 1952, 1958, 1959).¹¹ Like von Kries, Wassef asked whether one could predict asymmetric color matches using the principle of superposition. Like von Kries, she concluded that the weight of the experimental evidence supported the linearity hypothesis: when the illumination changes, the cone absorptions of the test and matching lights are related by a linear transformation.

I have replotted some of Wassef's data to illustrate the nature of the measurements and the size of the effect (Figure 9.10). The illuminant spectral power distributions she used in her dichoptic matching experiment are plotted in Figure 9.10A. To plot her results, I have converted Wassef's reported measurements into cone absorptions. I have plotted the cone absorptions of the surfaces that matched in color appearance when seen under the two illuminants. Figure 9.10B shows the L- and M-cone absorptions of the surfaces under the two illuminants, and Figure 9.10C shows the L- and S-cone absorptions. The cone absorptions for objects seen under a tungsten illuminant are plotted as open circles;

¹¹ Interestingly, one of the largest sets of data she reviewed was a series of experiments performed at the Kodak research laboratories, Polaroid's competitor.



9.10 DATA FROM AN ASYMMETRIC COLOR-MATCHING EXPERIMENT USING THE DICHOPTIC METHOD. The test and matching lights are viewed in different contexts and appear identical; but the two lights have very different cone absorption rates. Hence, appearance matches made across an illuminant change are not cone absorption matches. (A) The spectral power distributions of two illuminants, one approximating mean daylight and the other a tungsten illuminant are shown. (B) Cone absorptions for targets that appear identical to one another in these two contexts are shown as scatterplots for the L- and M-cones, and (C) for the L- and S-cones. The points plotted as open circles are cone absorptions for tests seen under the first illuminant; matches seen under the second illuminant are plotted as filled circles. The stimuli represented by the absorptions have the same color appearance, but they correspond to very different cone absorptions. Source: Wasef, 1959.

the cone absorptions for objects seen under a blue-sky illumination are plotted as filled circles. The size of the effect is quite substantial. Two sets of points show the cone absorptions of targets that look identical in their respective contexts, yet the cone absorption values from the surfaces under these illuminants do not even overlap in their values.

Taken together, the color matching and asymmetric color matching results show the following. When the objects are in the same context, equating the cone absorptions equates appearance, but when the two objects are seen under different illuminants, equating cone absorptions


does not equate appearance. Within each context the observer uses the pattern of cone absorptions to infer color appearance, probably by comparing the relative cone absorption rates. Color appearance is an interpretation of the physical properties of the objects in the image.

The von Kries Coefficient Law: Experiments

Through his coefficient law, von Kries sought to explain these asymmetric color matches by a simple physiological mechanism. He suggested that the visual pathways adjust to the illumination by scaling the signals from the individual cone classes. This hypothesis has a simple experimental prediction: if we plot, say, the S-cone absorptions of the test and match surfaces on a single graph, the data should fall along a straight line through the origin. The slope of the predicted line is the scale factor for the illuminant change.

Neither von Kries nor Wassef knew the cone spectral curves; hence, they could not create the graph they needed to test the von Kries coefficient law directly. However, using an indirect measurement based on estimation of the eigenvectors of the measured linear transformations, Burnham et al. (1957) and Wassef (1959) rejected von Kries scaling. Despite this rejection, von Kries's hypothesis continued to be used widely to explain how color appearance varies with illumination. Among theorists, for example, E. H. Land relied entirely on von Kries scaling as the foundation of his retinex theory (Brewer, 1954; Brainard and Wandell, 1986; Land, 1986).

Today, we have good estimates of the spectral sensitivities of the cone photopigments, and it is possible to convert Wassef's data into cone absorptions and to analyze the von Kries coefficient law directly. Figure 9.11 shows a graphical evaluation of von Kries's hypothesis for the data in Figure 9.10. Each panel plots the cone absorptions of corresponding test and match targets for one of the three cone types. As



9.11 GRAPHICAL EVALUATION OF VON KRIES'S HYPOTHESIS for the data in Figure 9.10. These appearance matches were made by presenting the test and match objects to different eyes. The illuminant for one eye was similar to a tungsten bulb, and the other eye was blue skylight. As predicted by von Kries, the cone absorptions of the test and match surfaces approximate a straight line. However, the von Kries coefficient law predicts that the line should pass through the origin of the graph; while not precisely correct, the rule is a helpful starting point. Source: Wassef, 1959.

predicted by von Kries, the cone absorptions of the test and match targets fall along a line. Moreover, the slopes of the lines relating the cone absorptions also make sense. The slope is largest for the S-cones, where illuminant change has its largest effect. The data are not perfectly consistent with von Kries scaling, however, because the lines through the data do not pass through the origin, as required by the theory.¹²

There is an emerging consensus in many branches of color science that the von Kries coefficient law explains much about how color appearance depends on the illumination. Von Kries's simple hypothesis is important partly because of its practical utility and partly because of its implications for the representation of color appearance within the brain. The hypothesis explains the major adjustments for color constancy in terms of the photoreceptor signal and how the photoreceptor signals combine within the retina (Chapter 4). Hence, von Kries's hypothesis implies that either (a) the main adjustment takes place very early, or (b) the photoreceptor signals can be separated in the central representation. This topic will come up again later, when we review some of the phenomena concerning color appearance in the central nervous system.


How Color-Constant Are We?

Finally, we will consider how well the visual pathways correct for illumination change. On this point there is some consensus: the asymmetric color matches do not compensate completely for the illumination change (Helson, 1938; Judd, 1940).

Brainard and Wandell (1991b, 1992) described this phenomenon using an apparatus consisting of simulated surfaces and illuminants and an asymmetric color-matching experiment based on memory-matches. Subjects were presented with images of simulated colored papers, rendered under a diffuse daylight illuminant, on a CRT display. The subjects memorized the color appearance of one of the surfaces. Next, the simulated illuminant was changed, slowly over a period of two minutes, giving subjects a chance to adapt to the new illuminant. Then, the subject adjusted the appearance of a simulated surface to match the color appearance of the surface he or she had memorized.

We can represent the difference between the two simulated illuminants by plotting the illuminant change. The filled symbols in Figure 9.12 show the illuminant changes in two experimental conditions. The top panel shows an illuminant change that increased the short-wavelength light and decreased the long-wavelength light. The bottom panel shows an illuminant change that increased the energy at all wavelengths.

¹² Indeed, this is equally a failure of the simple linearity that Wassef uses to summarize the data, and more in line with some of the conclusions that Burnham et al. (1957) drew about their data.



9.12 A COMPARISON OF THE ILLUMINANT CHANGE AND THE SUBJECTIVE ILLUMINANT CHANGE, as inferred from an asymmetric matching experiment. The simulated illuminant change and subjective illuminant changes are shown by the filled and open circles, respectively. Subjects behave as if the illuminant change is about half of the true illuminant change. Source: Brainard and Wandell, 1991b.

Suppose that subjects equated the perceived surface reflectance, but that the illuminant change they estimated was different from the true illuminant change. In that case, we can use the observed matches to infer the subjects' illuminant estimates, which are plotted as the open symbols in two panels of Figure 9.12. Subjects are acting as if the illuminant change they are correcting for is similar to the simulated illuminant change, but smaller. Subjects' performance is conservative, correcting for about half the true illuminant change.

The experiments in Brainard and Wandell (1991b, 1992) were conducted on display monitors, and the images were far less complex than full natural scenes. It is possible that, given additional clues, subjects may come closer to true illuminant estimation. But in most laboratory experiments to date, subjects do not compensate fully for changes in the illumination. When the illumination changes, color appearance changes less than it might if color were defined by the cone absorptions, but it changes more than it would if the nervous system used the best possible computational algorithms.

The Perceptual Organization of Color

In this section, I will review some of the methods for describing the perceptual organization of color appearance. Specifically, we will consider the relationship between different colors and some of the systems for describing color appearance. In addition to the implications this organization has for understanding the neural representation of color appearance, there are also many practical needs for descriptive systems of color appearance. Artists and designers need ways to identify and specify the

color appearance of a design. Further, they need ways of organizing colors and finding interesting color harmonies. Engineers need to assess the appearance and discriminability of colors used for highway signs, for labeling parts and packaging, and for designing software icons.

Language provides us with a useful start for organizing color appearance. Spoken English in the United States consists of 11 color terms that are widely and consistently used.¹³ While the number of terms used differs across cultures, there is a remarkable hierarchical organization to the order in which color names appear. Cultures with a small number of basic color names always include white, black and red. Color terms such as purple and pink enter later (Berlin and Kay, 1969; Boynton and Olsen, 1987).

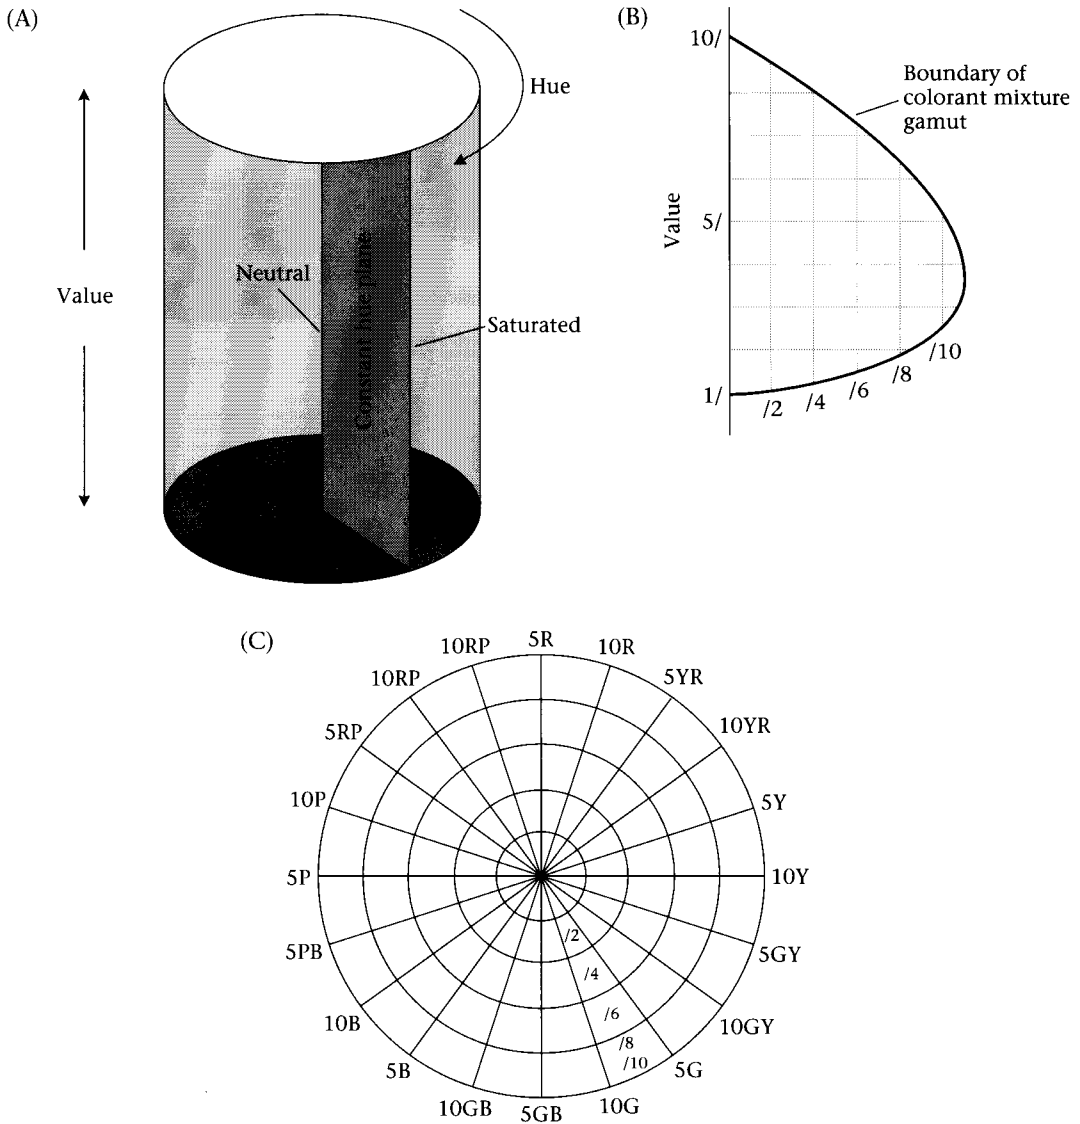
Color names are a coarse description of color experience. Moreover, names list but do not organize color experience. Thus, they are not helpful when we consider issues such as color similarity or color harmony. A more complete organization of color experience is based on the three perceptual attributes called **hue**, **saturation**, and **brightness**. Hue is the attribute that permits a color to be classified as red, yellow, green, and so forth. Saturation describes a color's similarity to a neutral gray or white. A gray object with a small reddish tint has little saturation, while a red object, with little white or gray, is very saturated. An object's brightness tells us about the relative ordering of the object on the dark-to-light scale.

Based on psychological studies of the similarity of colored patches with many different hues, saturations, and brightnesses, the artist Albert Munsell created a book of colored samples. The samples are organized with respect to hue, saturation, and brightness. Furthermore, the colored samples are spaced in equal perceptual steps. Munsell organized the samples within his book using a cylindrical organization as shown in Figure 9.13. The *Munsell Book of Colors* is sold by the Munsell Color Company, and is used as a reference in many design and engineering applications.

Perceptually, both saturation and brightness can be arranged using a linear ordering from small to large; hue, however, does not follow a linear ordering. Therefore, Munsell organized lightness along the main axis of the cylinder, and saturation as the distance from the center of the cylinder to the edge. The circular hue dimension was mapped around the circumference of the cylinder.

Munsell developed a special notation to refer to each of the samples in his book. To distinguish his notation from the colloquial usage, Munsell substituted the word **value** for lightness and the word **chroma** for saturation. He retained the word hue, apparently finding no adequate substitute. In the Munsell notation, the words hue, chroma, and value have specific and technical meanings. Each colored paper is described

¹³ White, black, red, green, yellow, blue, brown, purple, pink, orange, and gray.



9.13 THE MUNSELL BOOK OF COLORS is a collection of colored samples organized in terms of three perceptual attributes of color. (A) The samples are arranged using a cylindrical geometry. The main axis of the cylinder codes lightness; the distance from the center of the cylinder to the edge codes the Munsell property called value (saturation); the position around the circumference of the cylinder codes the Munsell property called chroma (hue). (B) A cross-section of the cylinder through the main axis represents a single hue, but different values and chroma. Not all combinations of hue, chroma and value can be reproduced; the general shape of the reproduction boundary is shown. (C) A cross-section perpendicular to the main axis includes samples at a single value but with different hues and chroma. Regions around the circle are labeled with basic color names (e.g., R = red, Y = yellow, GB = green-blue). These regions are divided into ten steps, though usually only the positions corresponding to steps 5 and 10 are shown. The hues are arranged so that similar hues are near to one another and complementary hues are on opposite sides of the circle.

using a three-part syntax of hue chroma/value. For example, 3YR 5/3 refers to a colored paper with the hue called 3YR, the chroma level 5, and the value level 3.

The *Munsell Book of Colors* was created before the CIE color standards described in Chapter 4. With the advent of the CIE measurement standard, based on the color-matching functions, there was a need for a method to convert the Munsell representation into the CIE standard representation. A committee of the Optical Society (Newhall et al., 1943) measured the CIE values of the Munsell samples in the published book, and the Munsell Corporation agreed to produce the colored samples to measurement standards defined by the Optical Society of America. The new standard for the *Munsell Book of Colors*, based on CIE values rather than pigment formulas, is called the **Munsell Renotation System**. Calibration tables that describe the color measurements of the Munsell samples are available, for example, in Wyszecki and Stiles (1982).

Opponent Colors

One of the most remarkable and important insights about color appearance is the concept of **opponent colors**, first described by Hering (1878). Hering pointed out that there is a powerful psychological relationship between the different hues. While some pairs of hues can coexist in a single color sensation, others cannot. For example, orange is composed of red and yellow while cyan is composed of blue and green. But we never experience a hue that is simultaneously red and green. Nor do we experience a color sensation that is simultaneously blue and yellow. These two hue pairs, red–green and blue–yellow, are called **opponent colors**.

There is no physical reason why these two opponent-color pairs must exist. That we never perceive red and green, while we easily perceive red and yellow, must be due to the neural representation of colors. Hering argued that opponent colors exist because the sensations of red and green are encoded in a single visual pathway. The excitation of the pathway causes us to perceive one of the opponent colors; inhibition of the pathway causes us to perceive the other.

Hering made his point forcefully and extended his theory to explain various other aspects of color appearance, as well. But his insights were not followed by a set of quantitative studies. Perhaps for this reason, his ideas languished while the colorimetrists used color-matching to set standards for all of modern technology. This is not to say that Hering's work was forgotten. Colorimetrists who thought about color appearance invariably turned to Hering's insights. In a well-known review article (Judd, 1951), the eminent scientist D. B. Judd wrote (p. 836):

The Hering (1905) theory of opponent colors has come to be fairly well accepted as the most likely description of color processes in the optic nerve and cortex. Thus this theory reappears in the final stage in the stage theories of von Kries-Schrodinger (von Kries, 1905; Schrodinger,

1925), Adams (1923, 1942) and Muller (1924, 1930). By far the most completely worked out of these stage theories is that of Muller. . . . There is slight chance that all of the conjectures are correct, but, even if some of the solutions proposed by Muller prove to be unacceptable, he has nevertheless made a start toward the solution of important problems that will eventually have to be faced by other theorists.

Hue Cancellation


Several experimental observations, beginning in the mid-1950s, catapulted opponent-colors theory from a special-purpose model, known only to color specialists, to a central idea in vision science. The first was a behavioral experiment that defined a procedure for measuring opponent colors, the **hue-cancellation experiment**. The hue cancellation experiment was developed in a series of papers by Jameson and Hurvich (1955, 1957). By providing a method of quantifying the opponent-colors insight, Hurvich and Jameson made the idea accessible to other scientists, opening a major line of inquiry.

In the hue-cancellation experiment, the observer is asked to judge whether a test light appears to be, say, reddish or greenish. If the test light appears reddish, the subject adds green light in order to cancel precisely the red appearance of the test. If the light appears greenish, then the subject adds red light to cancel the green appearance. The added light is called the **canceling light**. Once the red or green hue of the test light is canceled, the test plus canceling light appears yellow, blue, or gray. The same experiment can be performed to measure the blue–yellow opponent-colors pairing. In this case the subject is asked whether the test light appears blue or yellow, and the canceling lights also appear blue or yellow.

Figure 9.14 shows a set of hue-cancellation measurements obtained by Jameson and Hurvich (1955, 1957). Subjects canceled the red–green or blue–yellow color appearance of a series of spectral lights. The vertical axis shows the relative intensity of the canceling lights, scaled so that when equal amounts of these lights are superimposed the result did not appear, say, red or green. The canceling lights always have positive intensity, but the intensity of the green and blue canceling lights are plotted as negative to permit you to distinguish which canceling light was used.

To what extent can we generalize from red–green measurements using monochromatic lights to other lights? To answer this question we must evaluate the linearity of the hue-cancellation experiment. If the experiment is linear, we can use the data in Figure 9.14 to predict whether any test light will appear red or green (or blue or yellow) since all lights are the sum of monochromatic lights. If the experiment is not linear, then the data represent only an interesting collection of observations.

To evaluate the linearity of the hue-cancellation experiment, one can perform the following experiment. Suppose the test light t_1 ap-



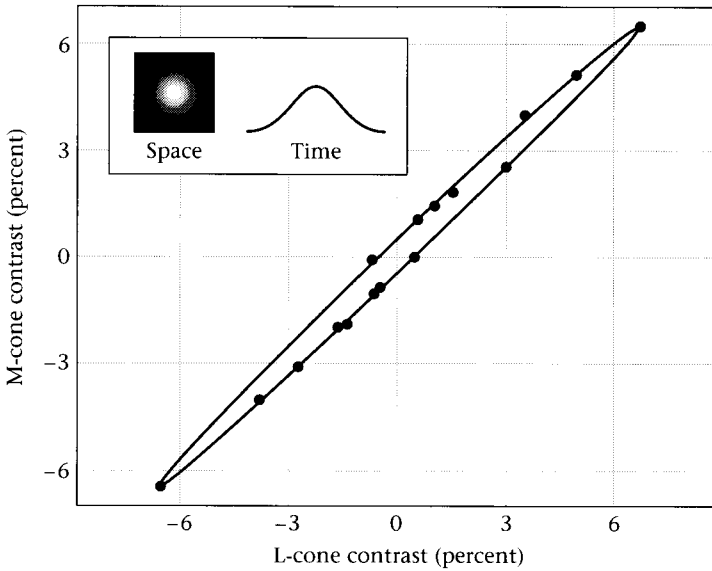
9.14 MEASUREMENTS FROM THE HUE-CANCELLATION EXPERIMENT. An observer is presented with a monochromatic test light. If the light appears red then some amount of a green canceling light is added to cancel the redness. If the light appears green, then a red canceling light is added to cancel the greenness. The horizontal axis of the graph measures the wavelength of the monochromatic test light, and the vertical axis measures the relative intensity of the canceling light. The entire curve (open triangles) represents the red–green appearance of all monochromatic lights. A similar procedure is used to measure blue–yellow hue cancellation (filled circles). The two graphs show results from two observers. Source: Hurvich and Jameson, 1957.

pears neither red nor green, and the test light t_2 appears neither red nor green. Does the superposition of these two test lights, $t_1 + t_2$, also appear neither red nor green? In general, the hue-cancellation experiment fails this test of linearity. If we superimpose two lights, neither of which appears red or green, the result can appear red; if we add two lights neither of which appears blue or yellow, the result can appear yellow. Hence, although the hue-cancellation studies provide a useful benchmark, we need a more complete (nonlinear) model before we can apply the hue-cancellation data in Figure 9.14 to predict the opponent-colors appearance of polychromatic test lights (Larimer et al. 1975; Burns et al., 1984; Ayama and Ikeda, 1989; Chichilnisky, 1995).

Opponent-Colors Measurements at Threshold

In addition to color-appearance judgments, one can also demonstrate the presence of essential opponent-colors signals behaviorally by **color test-mixture experiments**. These color experiments are direct analogues of the pattern-mixture experiments reviewed in Chapter 7 (see Figs. 7.7 and 7.9). A set of threshold measurements from a color test-mixture experiment are shown in Figure 9.15. The axes represent the absorption rates of the L- and M-cones. Each point represents a combination of absorption rates at detection threshold. The smooth curve fit through the data points is an ellipse.

The intersection of the ellipse with the horizontal axis, shows the relative L-cone absorption rate at detection threshold. The intersection



9.15 COLOR TEST-MIXTURE EXPERIMENTS demonstrate opponent-colors processes. The axes measure percentage change in cone absorption rates for the L- and M-cones. The points show the cone absorption rates at the detection threshold measured using different colored test lights. The smooth curve is an ellipse fit through the data points. The mixture experiment shows that the L- and M-cone signals cancel one another, so that lights that excite a mixture of L- and M-cones are harder to see than lights that stimulate just one of these two cone classes. The insets show the spatially blurred appearance of the test stimulus and the gradual timecourse as the stimulus contrast increases and then decreases. Source: Wandell, 1985.

of the ellipse with the vertical axis shows the relative M-cone absorption rate at detection threshold. The shape of the ellipse shows that signals that stimulate the L- and M-cones simultaneously are less visible than signals that stimulate only one or the other. At the most extreme points on the ellipse, the cone absorptions of the L- and M-cones are more than five times the rate required to detect a signal when each cone class is stimulated alone. The poor sensitivity to mixtures of signals from these two cone types shows that the signals must oppose one another. The cancellation of threshold-level signals from the L- and M-cones, as well as between the S-cones and the other two classes (not shown), has been observed in many different laboratories and under many different experimental conditions (e.g., Boynton et al., 1964; Pugh, 1976; Pugh and Mollon, 1979; Mollon and Polden, 1977; Sternheim, 1979; Wandell and Pugh, 1980a, b; Stromeyer et al., 1985).

In addition to demonstrating opponent colors, these threshold data reveal a second interesting and surprising feature of visual encoding. Two neural signals that are visible when they are seen singly become invisible when they are superimposed. It seems odd that the visual system should be organized so that plainly visible signals can be made invisible. From the figure we can see that this is a powerful effect, suppressing a signal that is more than five times threshold. This observation tells us

that in many operating conditions absolute sensitivity is not the dominant criterion. The visual pathways can sacrifice target visibility in order to achieve the goals of the opponent-colors encoding.

Opponent Signals in the Visual Pathways

In addition to these two types of behavioral evidence, there is also considerable physiological evidence that demonstrates the existence of opponent-colors signals in the visual pathway. In a report that gained widespread attention, Svaetichin (1956) measured the responses of three types of retinal neurons in a fish. He reported that the electrical responses were roughly consistent with Hering's notion of the opponent-colors representation. In two types of neurons, the electrical response increased to certain wavelengths of light and decreased in response to other wavelengths, paralleling the red-green and blue-yellow opponency in color perception.¹⁴ The electrical response of a third set of neurons increased to all wavelengths of light, as in a black-white representation. Shortly after Svaetichin's report, DeValois and his colleagues established the existence of opponent-colors neurons in the LGN of nonhuman primates. There is now a substantial literature documenting the presence of opponent-color signals in the visual pathways (e.g., DeValois et al., 1958a,b, DeValois, 1965a,b; DeValois et al., 1966; Wiesel and Hubel, 1966; Gouras, 1968; Derrington et al., 1984; Lennie et al., 1991).

The resemblance between the psychological organization of opponent colors measured in the hue-cancellation experiment and the neural opponent signals suggests a link from the neural responses to the perceptual organization. To make a convincing argument for the specific connection between opponent colors and a specific set of neural opponent signals, we must identify a linking hypothesis. The hypothesis should tell us how we can predict appearance from the activity of cells, and conversely how we can predict the activity of these cells from appearance.

A natural starting place is to suppose that there is a population of neurons whose members are excited when we perceive red and inhibited when we perceive green. From the linking hypothesis, we predict that neurons in this population will be unresponsive to lights that appear neither red nor green. There are two spectral regions that appear neither red nor green to human observers: one near 570 nm and a second near 470 nm. The question of whether there exists a population of neurons with the hypothetical response patterns was studied by DeValois and his collaborators in the LGN of the monkey. In their studies, DeValois and his colleagues studied the response of neurons to


¹⁴ At first, it was thought that these responses reflected the activity of the the cones. Subsequent investigations showed that the responses were from horizontal cells (Svaetichin and MacNichol, 1958).

monochromatic stimuli presented on a zero background. They found a weak correspondence between the neutral points of individual neurons and the perceptual neutral points (DeValois et al., 1966). More recently, Derrington et al. (1984) measured the responses of LGN neurons using contrast stimulus presented on a moderate neutral background. They estimated the input to these neurons from the different cone classes and confirmed the basic observations made by DeValois and his colleagues.

Derrington et al. (1984) reported that parvocellular neurons can be classified into two groups. One population of neurons receives opposing input from the L- and M-cones; the left panel of Figure 9.16 shows my estimate of the spectral sensitivity of this group of parvocellular neurons. While stimulus wavelengths near 570 nm are ineffective at exciting these cells, the precise wavelength of the zero-crossing varies widely across the neural population. Moreover, neurons in this population do not show a second zero-crossing near 470 nm that would parallel the human opponent-colors judgments in the hue-cancellation experiment. Hence, there is no strong evidence that this population of neurons represents red-green color appearance.

A second population of LGN neurons receives input from the S-cones and an opposing signal from a combination of the L- and M-cones. For these neurons, wavelengths near 500 nm are quite ineffective. The panel on the right of Figure 9.16 shows my estimate of the spectral sensitivity of this group of parvocellular neurons.

There was less order in the opponent-color signals of the magnocellular neurons. Many magnocellular units seemed to be driven by a difference between the L- and M-cones. A few parvocellular units and a few magnocellular units were driven by a positive sum of the signals from these two cone types.



9.16 OPPONENT SIGNALS MEASURED IN LGN NEURONS. These spectral response curves are inferred from the measured responses of LGN neurons to many different colored stimuli presented on a monitor. The vast majority of LGN neurons in the parvocellular layers can be divided into two groups based on their response to modulations of colored lights. (A) One group of neurons receives an opponent contribution from the L- and M-cones alone. (B) The second group of neurons receives a signal of like sign from the L- and M-cones, and an opposing signal from the S-cones. Source: Derrington et al., 1984.

The spectral responses of these neural populations suggest that there is only a loose connection between the signals coded by the neurons and the perceptual coding into opponent hues; it is unlikely that excitation and inhibition causes our perception of red–green and blue–yellow as opponent colors. One difficulty is the imperfect correspondence between the neural responses and the hue-cancellation measurements. The second difficulty is that there is no substantial population of neurons representing a white–black signal. This is a very important perceptual dimension which must be carried in the signals of the LGN. Yet, no clearly identified group of neurons can be assigned this role.¹⁵

Decorrelation of the Cone Absorptions

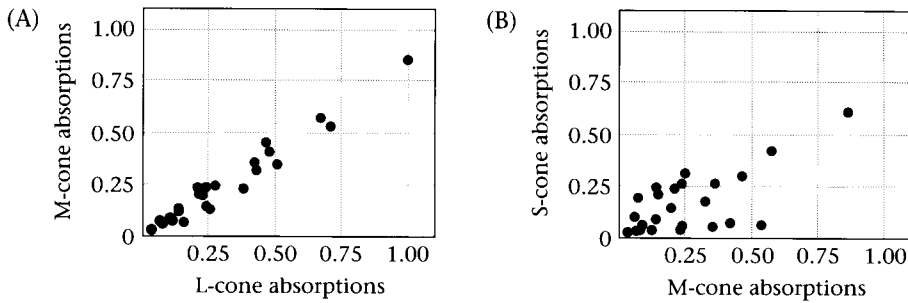
The opponent signals measured in the LGN probably represent a code used by the visual pathways in communicating information from the retina to the brain. The psychological opponent-colors coding may be a consequence of the coding strategy used to communicate information from the retina to the cortex. What reason might there be for using an opponent-signals coding?

One reason to use an opponent-signal representation has to do with the efficiency of the visual encoding. Because of the overlap of the L- and M-cone spectral sensitivities, the absorption rates of these two cone types are highly correlated. This correlation represents an inefficiency in the visual coding of spectral information. As I described in Chapter 8, decorrelating the signals can improve the efficiency of the neural representation.

We can illustrate this principle by working an example, parallel to the one in Chapter 8. Consider the cone absorptions for a set of surfaces. Because of the overlap in spectral sensitivities, the cone absorptions between, say, the L- and M-cones will be correlated. To remove the correlation, we create a new representation of the signals consisting of the L- cone absorptions alone and a weighted combination of the L-, M-, and S-cone absorptions. We will choose the weighted combination of signals so that the new signal is independent of the L-cone absorptions. As we reviewed in the earlier chapter, by decorrelating the cone absorptions before they travel to the brain, we make effective use of the dynamic range of the neurons transmitting the information (Buchsbau and Gottschalk, 1984).

The graphs in Figure 9.17 show examples of the correlation of the cone absorptions for a particular set of surfaces and illuminant. These

¹⁵ Some authors have suggested that a single group of LGN neurons codes a white–black sensation for high-spatial-frequency patterns and a red–green sensation for low-spatial-frequency patterns. This is an interesting hypothesis that abandons the idea that there is a specific color sensation associated with the response of LGN neurons. Instead, the perceived hue depends on the pattern of neural activation (Derrington, et al., 1984; Ingling and Martinez, 1984).



9.17 ABSORPTIONS IN THE THREE CONE CLASSES ARE HIGHLY CORRELATED.

The correlation between cone absorptions can be measured using correlograms. In this figure, correlograms are shown for the cone absorptions from the surfaces in the Macbeth ColorChecker illuminated by average daylight. (A) A correlogram of the L- and M-cone absorptions. (B) A correlogram of the M- and S-cone absorptions.

plots represent the cone absorptions from light reflected by a Macbeth ColorChecker viewed under mean daylight illumination. The correlations shown in these two plots are typical of natural images: the L- and M-cone absorptions are highly correlated (panel A); the M- and S-cone absorptions are also correlated (panel B).

As described in Chapter 8, we decorrelate the signals derived from the cone absorptions by forming new signals that are weighted combinations of the cone absorptions. There are many linear transforms of the cone absorptions that could serve to decorrelate these absorptions. One such transformation is represented by the three linear equations,¹⁶

$$\begin{aligned}
 O_1(\lambda) &= 1.0L(\lambda) + 0.0M(\lambda) + 0.0S(\lambda) \\
 O_2(\lambda) &= -0.59L(\lambda) + 0.80M(\lambda) + -0.12S(\lambda) \\
 O_3(\lambda) &= -0.34L(\lambda) + -0.11M(\lambda) + 0.93S(\lambda)
 \end{aligned} \tag{9.16}$$

or, written in matrix form,

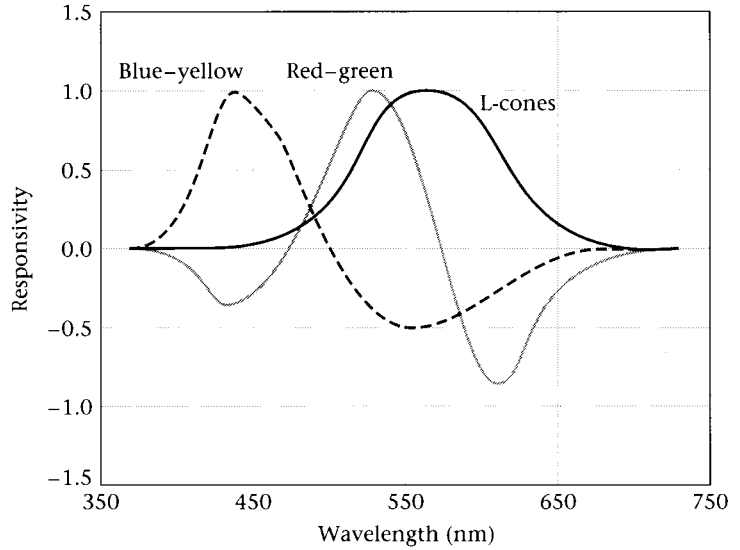
$$\begin{pmatrix} O_1(\lambda) \\ O_2(\lambda) \\ O_3(\lambda) \end{pmatrix} = \begin{pmatrix} 1.00 & 0.00 & 0.00 \\ -0.59 & 0.80 & -0.12 \\ -0.34 & -0.11 & 0.93 \end{pmatrix} \begin{pmatrix} L(\lambda) \\ M(\lambda) \\ S(\lambda) \end{pmatrix}$$

The new signals, $O_i(\lambda)$, are related to the cone absorptions by a linear transformation. These three signals are decorrelated with respect to this particular collection of surfaces and illuminant.

The spectral sensitivity of the three decorrelated signals are shown in Figure 9.18. The two opponent spectral sensitivities are reminiscent of the hue-cancellation measurements and the opponent signals measured in the LGN. One of the sensors has two zero-crossings, near

¹⁶The decorrelation is based on the singular value decomposition of the cone absorptions.

9.18 THE SPECTRAL RESPONSIVITY OF A SET OF COLOR SENSORS whose responses to the Macbeth ColorChecker under mean daylight are decorrelated. The spectral sensitivities of these sensors resemble the spectral sensitivities of LGN neurons and the color-appearance judgments measured in the hue-cancellation experiment.



570 nm and 470 nm. A second sensor has one zero-crossing at about 490 nm. The third sensor has no zero-crossings, as required for a white-black pathway. The similarity among the decorrelated signals, the opponent signals in the lateral geniculate nucleus, and the hue-cancellation experiment suggests a purpose for opponent-colors organization. Opponent colors may exist to decorrelate the cone absorptions and provide an efficient neural representation of color (Buchsbaum and Gottschalk, 1984; Derrico and Buchsbaum, 1991).

The opponent-colors representation is a universal property of color appearance for humans, just as the need for efficient coding is a simple and universal idea. We should expect to find a precise connection between opponent-colors appearance and neural organization in the central visual pathways. The hue-cancellation experiment provides us with a behavioral method of quantifying opponent-colors organization. Hue-cancellation measurements establish a standard for neurophysiologists to use when evaluating opponent signals in the visual pathways as candidates for the opponent-colors representation. Opponent-colors organization is a simple and important idea; pursuing its neural basis will lead us to new ideas about the encoding of visual information.

Spatial Pattern and Color

Color Plate 2 and Figure 9.1 show that the color appearance at a location depends on the local image contrast, that is, the relationship between the local cone absorptions and the mean image absorptions. The targets we used to demonstrate this dependence are very simple spatial patterns, squares or lines, with no internal spatial structure of their own. In this section, we will review how color appearance can also depend on the spatial structure, such as the texture or spatial frequency, of the target itself.

Color Plate 6 shows two square waves composed of alternating blue and yellow bars. One square wave is at a higher spatial frequency than the other. The average signal reflected from the regions containing the square waves is the same; that is, these are pure contrast modulations about a common mean field. If you examine the square waves from a close distance, you will see that bars in the square-wave patterns are drawn with the same ink. If you place this book a few meters away from you, say, across the room, the color of the bars in the high-spatial-frequency pattern will appear different from the color of the bars in the low-spatial-frequency pattern. The bars in the high-spatial-frequency patterns will appear to be light and dark modulations about the green average. The bars in the low-spatial-frequency pattern will continue to look like distinct blue and yellow bars.¹⁷

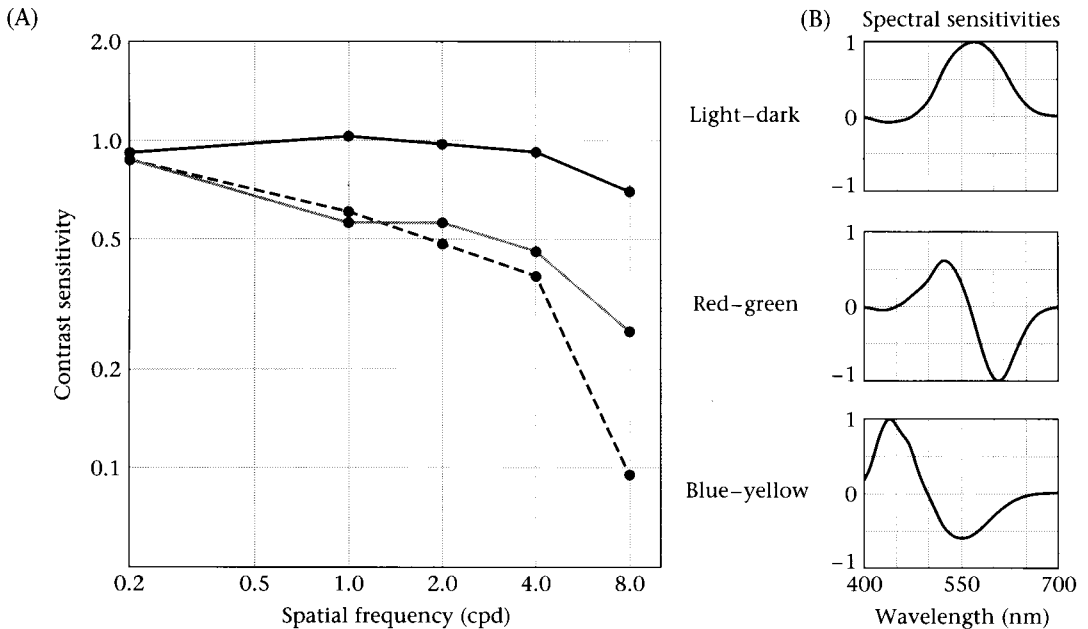
Poirson and Wandell (1993) used an asymmetric color-matching task to study how color appearance changes with spatial frequency of the square-wave pattern. Subjects viewed square-wave patterns whose bars were colored modulations about a neutral gray background; that is, the average of the two bars comprising the pattern was equal to the mean background level. Subjects adjusted the appearance of a 2-degree square patch to have the same color appearance as each of the bars in the pattern.

Two qualitative observations stood out in this study. First, spatial patterns of moderate- and high-spatial-frequency patterns (above 8 cpd) appear mainly light–dark, with little saturation. Thus, no matter what the relative cone absorptions of a high-spatial-frequency target, the target appeared to be a light–dark variation about the mean level. Second, the spatially asymmetric color appearance matches are not photopigment matches. This can be deduced from the first observation: because of axial chromatic aberration, moderate-frequency square-wave contrast patterns (4 and 8 cpd) cannot stimulate the S-cones significantly. Yet, subjects match the bars in these high-frequency patterns using a 2-degree patch with considerable S-cone contrast. The asymmetric color matches are established at neural sites central to the photoreceptors.

Poirson and I explained the asymmetric spatial color matches using a pattern–color separable model. In this model, we supposed that the color appearance of the target was determined by the response of three color mechanisms, and that the response of each mechanism was separable with respect to pattern and color. We derived the spatial and spectral responsivities of these pathways from the observers' color matches; the estimated sensitivities are shown in Figure 9.19.

Interestingly, the three color pathways that we derived from the asymmetric matching experiment correspond well to the opponent-color mechanisms derived from the hue-cancellation experiment. One pathway is sensitive mainly to light–dark variation; this pathway has

¹⁷ You can also alter the relative color appearance of the patterns by moving the book rapidly up and down. You will see that the low-frequency square wave retains its appearance, while the high-frequency square wave becomes a green blur.



9.19 ESTIMATES OF PATTERN-COLOR SEPARABLE SENSITIVITY of pathways mediating color appearance. By measuring spatially asymmetric color matches, it is possible to deduce the pattern and color sensitivity of three visual mechanisms that mediate color appearance judgments. (A) The pattern sensitivities and (B) the wavelength sensitivities derived from experimental measurements are shown for light-dark, red-green, and blue-yellow mechanisms. Source: Poirson and Wandell, 1993.

the best spatial resolution. The other two pathways are sensitive to red-green and blue-yellow variation. The blue-yellow pathway has the worst spatial resolution. Granger and Heurtley (1973), Mullen (1985), and Sekiguchi et al. (1993a, b) made measurements that presupposed the existence of opponent-color pathways and estimated similar pattern sensitivities of the three mechanisms. Notice that the derivation of the opponent-colors representation in this experiment did not involve asking the observers any questions about the hue or saturation of the targets. The observers simply set color-appearance matches; the opponent-colors mechanisms were needed to predict the color matches.

One of the more striking aspects of opponent-colors representations is that the apparent spatial sharpness, or focus, of a color image depends mainly on the sharpness of the light-dark component of the image; apparent sharpness depends very little on the spatial structure of the opponent-color image components. This is illustrated in the three images shown in Color Plate 7.

We can take advantage of the low spatial acuity of the opponent-colors representations when we code color images for image storage and transmission. We can allocate much less information about the opponent-colors components of color images without changing the apparent spatial sharpness of the image. This property of human per-

ception was important in shaping broadcast television standards and digital image-compression algorithms. As a quantitative prediction, we should expect to find that neurons in the central visual pathways that represent light–dark information should be able to represent spatial information at a much higher resolution than neurons that code opponent-colors information. Consequently, we should expect that the largest fraction of central neurons encode light–dark, rather than the other two opponent-colors signals.

The differences between the light–dark encoding and the opponent-colors encoding are of great perceptual significance. Consequently, several authors have studied hypotheses based on the idea that opponent-colors signals and light–dark signals are found in separate areas of the brain. In the final section of this chapter, we will consider some of the evidence concerning the representation of color information in the visual cortex.

The Cortical Basis of Color Appearance

Clinical Studies

In 1974, Meadows reviewed case studies of 14 patients who had lost their ability to see colors due to brain injuries. For some patients, the colors of objects appeared wrong. Other patients saw the world entirely in shades of gray. However, these patients still had good visual acuity.

The syndrome Meadows reviewed, which I will call **cerebral dyschromatopsia**,¹⁸ had been described in reports spanning a century (Zeki, 1990a,b). But the cases were rare, poor methods were used to study the patients, and the color loss was not well dissociated from other visual deficits. Consequently, at the time Meadows wrote his review, several well-known investigators had expressed doubt about even the existence of cerebral dyschromatopsia (e.g., Teuber et al., 1960).¹⁹ By bringing together a number of new cases and studying them with much better methods, Critchley (1965), Meadows (1974), Zeki (1990b), and others (e.g., Green and Lesell, 1977; Domasio et al., 1980; Mollon, 1980;

¹⁸Some of the terms used to describe color loss vary between authors. The terms trichromacy, dichromacy, and monochromacy are precise, referring to the number of primary lights necessary to complete the color-matching experiment. Some authors use the phrase *cerebral achromatopsia*, meaning “without color vision,” to describe a loss of color vision while others use *cerebral dyschromatopsia*. I prefer the second term because in these cases insensitivity to hue is often not complete and because these patients still distinguish the colors white and black. When the behavioral evidence warrants it, one might append a modifier, such as *monochromatic dyschromatopsia*, to describe the the color loss more precisely.

¹⁹In his book and in a long article, Zeki (1990b, 1993) argued that the skepticism concerning cerebral dyschromatopsia was due to the acceptance of a profoundly misguided theory concerning the significance of visual area V1. I agree with Meadows’s gentler assessment; the early evidence in support of cerebral dyschromatopsia was spotty and poorly argued. There was room for some skepticism.

Heywood et al., 1987, 1992; Victor et al., 1987) have removed any doubt about the existence and significance of the syndrome.

Congenital Monochromacy

Usually, observers are dichromats or monochromats because they are missing one of the cone photopigments (see e.g., Alpern, 1974; Smith and Pokorny, 1972). There are also reports of congenital cone monochromacy of central origin. In a thorough and fascinating study, Weale (1953) searched England for individuals who (a) could not tell color photographs from black and white, (b) were not photophobic, and (c) had good visual acuity. [Requirements (b) and (c) eliminated rod monochromats]. Weale found three cone monochromats, that is, individuals who could adjust the intensity of a single primary light to match the appearance of any other test light. Yet, based on direct measurements of the photopigment in the eye of one of the observers, as well as behavioral measurements, some of these cone monochromats were shown to have more than one cone photopigment (Weale, 1959; Gibson, 1962; see also Alpern, 1974). Hence, Weale's subjects had a congenital dyschromatopsia caused by deficiencies central to the photopigments. At present, we know little more about them.

Regularities of the Cerebral Dyschromatopsia Syndrome

When color loss arises from damage to the brain, the distortion of color appearance can take several forms. In some cases, patients report that colors have completely lost their saturation and hue and the world becomes gray. In other cases, color appearance may become desaturated. Some observers can perform some simple color-discrimination tasks, but they report that the colors of familiar objects do not appear right. In many cases the loss is permanent, but there are also reports of transient dyschromatopsia. For example, Lawden and Cleland (1993) recently reported on the case of a woman who suffers from migraines. During the migraine attacks, her world becomes transiently colorless.

The variability in the case studies suggest that a variety of mechanisms may disturb color appearance. Across this variability, however, there are also some regularities. First, Meadows (1974) observed that every patient with dyschromatopsia was blind in some portion of the upper visual field.

Meadows also examined the reverse correlation: do patients with purely upper visual field losses tend to have cerebral dyschromatopsia? In the literature, he found 12 patients with a purely upper visual field loss, seven of whom had dyschromatopsia; of 16 patients with a purely lower visual field loss, none had dyschromatopsia. In humans, the upper visual field is represented along the lower part of the calcarine sulcus (Chapter 6). The correlation between field loss and dyschromatopsia suggests a second regularity: damage that leads to dyschromatopsia is either near the lower portion of the calcarine or somewhere along the

path traced out by the nerve fibers whose signal enters the lower portion of the calcarine cortex.

Third, many of the patients suffer from a syndrome called **prosopagnosia**, the inability to recognize familiar faces. Twelve of the 14 patients described by Meadows had this syndrome. The patient with migraines also has transient prosopagnosia (Lawden and Cleland, 1993). The co-occurrence of dyschromatopsia and prosopagnosia suggests that the neural mechanisms necessary for recall of familiar faces and color are located close to one another or that they rely on the same visual signal. Based on his review of the literature, Meadows (1974, p. 622) concluded that, "The evidence on localization in cases of cerebral achromatopsia points to the importance of bilateral, inferiorly placed, posterior lesions of both cerebral hemispheres."²⁰

Behavioral Studies of Patients with Cerebral Dyschromatopsia

Patients with cerebral dyschromatopsia often fail to identify any of the test patterns on the Ishihara plates (see Chapter 4).²¹ Mollon et al. (1980) reported on a patient who failed to identify the targets on the Ishihara plates at reading distance, but who could distinguish the targets when the plates were viewed from a distance of 2 meters. At the 2-meter viewing distance, the neutral areas separating the target and background are barely visible, and the target and background appear contiguous. Twelve years after the original study, Heywood et al. (1992) replicated the finding on the same patient. They also showed that the patient could discriminate contiguous colors, but not colors separated by a gray stripe. Hence, in this patient, cerebral dyschromatopsia involves color and pattern together (see also Victor et al., 1987).

Cerebral dyschromatopsics score quite poorly on the Farnsworth–Munsell hue test (see Chapter 4). The pattern of errors does not correspond to the errors made by any class of dichromat. The results of the test of one such patient is shown in Figure 9.20. The errors are large in all directions, though there is some hint that the errors may be somewhat larger in the blue and yellow portions of the hue circle.


How Many Cone Types Are Functional?

The patients' errors on the Ishihara color plates and the Farnsworth–Munsell hue test are not consistent with visual-pigment loss. Nonetheless, we cannot tell from their performance on these tests whether the

²⁰ As Zeki (1990b) points out, Meadows's conclusion echoes a disputed suggestion made a century earlier. While studying a patient who reported a loss of color vision, the French physician, Verrey (1888) concluded, "Le centre du sense chromatique se trouverait dans la partie la plus inferieure du lobe occipital, probablement dans la partie posterieure des plis lingual et fusiforme." [Translation: The center of the chromatic sense will be found in the inferior part of the occipital lobe, probably in the posterior part of the lingual and fusiform gyrus.]

²¹ Meadows (1974) and Victor et al. (1987) describe patients with cerebral dyschromatopsia who nonetheless could read all of the plates.

9.20 RESULTS OF THE FARNSWORTH-MUNSELL HUE TEST FROM A PATIENT SUFFERING CEREBRAL DYSCHROMATOPSIA. The patient's error scores are high in all hue directions. This pattern of scores is not consistent with any of the usual pattern of errors observed by cone dichromats who are missing one of their cone photopigments (see Figure 4.21). Source: Meadows, 1974.



separate cone classes are functioning or whether the loss of color perception is due, in part, to cone dysfunction.

Gibson (1962; see also Alpern, 1974; Mollon et al., 1980) developed a behavioral test to infer whether patients with cerebral dyschromatopsia had more than a single class of functioning cones. The logic of this behavioral test is based on the fact the cone signals are scaled to correct for changes in the ambient lighting conditions. For example, in the presence of a long-wavelength background, the sensitivity of the L-cones is suppressed, while the sensitivity of the S-cones remains unchanged.

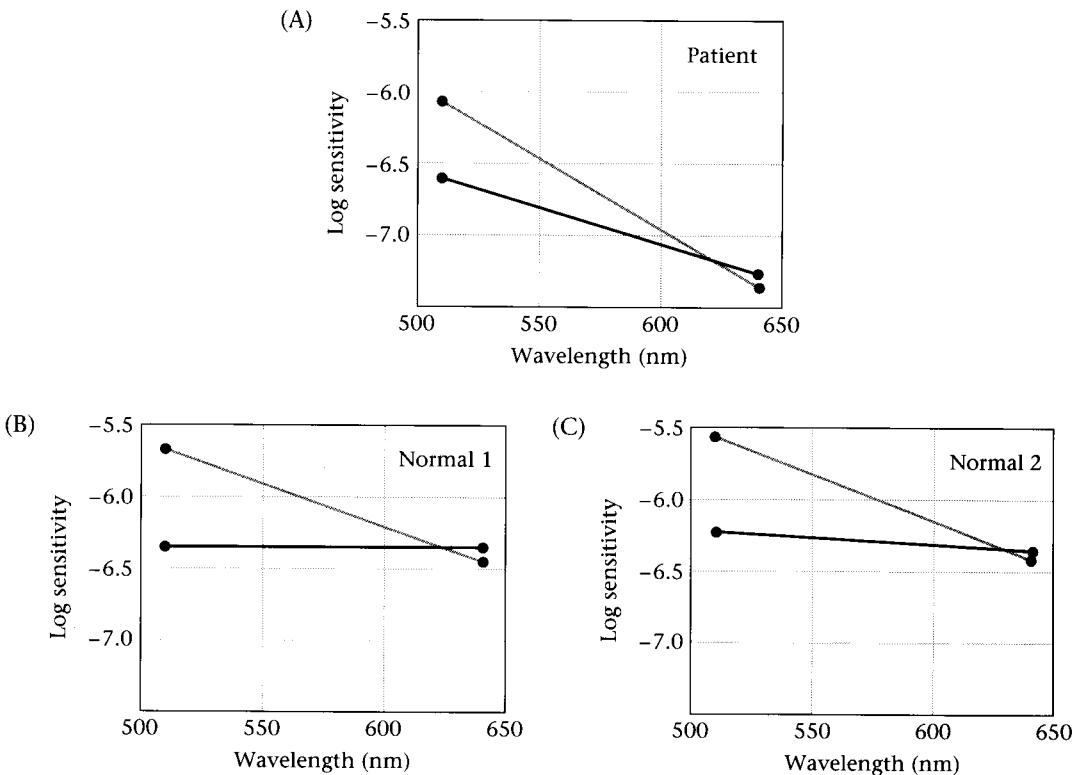
Suppose a subject has only a single type of cone. For this observer, wavelength sensitivity is determined by the spectral sensitivity of a single cone photopigment. Changes of the background illumination will not change the observer's relative wavelength sensitivity. This is the situation for normal observers under scotopic viewing conditions, when we see only through the rods. Under scotopic conditions, wavelength sensitivity is determined by the rhodopsin photopigment; changing the background does not change in the relative sensitivity to different test wavelengths.²²

²²In a beautiful series of experiments, Stiles (1939, 1959, 1978) studied how sensitivity varies as one changes the wavelength and intensity of a test and background lights. He developed a penetrating analysis of this experimental paradigm and identified candidate processes which he believed might describe photoreceptor adaptation. He referred to these processes as π -mechanisms.

If an individual has two functional cone classes, however, changes in the sensitivity of one cone class relative to the other will change the behavioral wavelength sensitivity. Hence, we can detect the presence of two cone classes by measuring wavelength sensitivity on two different backgrounds and noting a change in the observer's relative wavelength sensitivity.

Mollon et al. (1980) measured a cerebral dyschromatopsic's relative wavelength sensitivity to test wavelengths (510 nm and 640 nm) on two different backgrounds (510 nm and 650 nm). I have replotted their data in Figure 9.21. When the background changes, the relative test-wavelength sensitivity changes, showing that the subject has at least two functional cone classes, as in the case of congenital monochromats.

Clinical studies of cerebral dyschromatopsia show that central lesions can disturb color vision severely, while sparing many other aspects



9.21 EXPERIMENTAL DEMONSTRATION THAT A PATIENT WITH CEREBRAL DYSCHROMATOPSIA HAS MORE THAN A SINGLE FUNCTIONING CONE CLASS.

(A) The patient's threshold sensitivity was measured to two monochromatic test lights on two different backgrounds. The change in background illumination changed the patient's relative wavelength sensitivity. (B, C) The results of performing the same experiment on normal observers. The results from the normal observers and the patient are quite similar. Source: Mollon et al., 1980.

of visual performance. This clinical syndrome suggests that some of the neural mechanisms essential to the sensation of color appearance may be anatomically separate from the mechanisms required for other visual tasks, such as acuity, motion perception, and depth perception. Clinical lesions are not neat and orderly, however, and the syndrome of cerebral dyschromatopsia is quite varied. Alternative hypotheses, for example, that neurons carrying color information are more susceptible to stroke damage than other neurons, are also consistent with the clinical observations (Mollon et al., 1980). To pursue the question of the neural representation of color information, we need to consider other forms of evidence concerning the localization of color appearance.

Physiological Studies of Color Appearance

Much of the agenda for modern research on the cortical representation of color appearance has been set by Zeki via a hypothesis he calls **functional segregation** (Zeki, 1974, 1993; see Chapter 6). Zeki argues that there is a direct correlation between the neural responses in cortical areas beyond V1 and various perceptual features, such as color, motion, and form. This is not the only hypothesis we might entertain for the relationship between brain structures and perceptual function. An alternative view has been expressed by Livingstone and Hubel (1984, 1987) who argue that perceptual function can be localized to groups of neurons residing within single visual areas. Specifically, they have argued that differences in the density of the enzyme cytochrome oxidase within cell bodies serves as a clue to the localization of perceptual processing (see Chapter 6). This criterion for identifying neural segregation of function seems relevant in areas V1 and V2, since the anatomical interconnections between these areas appear to respect the differences in cytochrome oxidase density (Burkhalter et al., 1989).

Livingstone and Hubel's hypothesis need not conflict with Zeki's, since information may be intertwined within peripheral visual areas only to be segregated later. But the presence of subdivisions within areas V1 and V2 raises the question of whether more detailed study might not reveal functional subdivisions within the other cortical areas as well (see e.g., Born and Tootell, 1992; Deyoe et al., 1994).

The principal line of evidence used to support Zeki's hypothesis of functional segregation is Barlow's neuron doctrine, namely, that the receptive field of a neuron corresponds to the perceptual experience the animal will have when the neuron is excited (Chapter 6). Based on this doctrine, neurophysiologists frequently assume that neurons with spatially oriented receptive fields are responsible for the perception of form; neurons that are inhibited by some wavelengths and excited by others are responsible for opponent-color percepts; neurons with motion-selective receptive fields are responsible for motion perception.

Zeki's suggestion that monkey area V4 is a color center and area MT is a motion center is based on differences in the receptive-field properties of neurons in these two areas. The overwhelming majority of neurons in area MT show motion direction selectivity. Zeki reported that many neurons in area V4 exhibited an unusual wavelength selectivity (Zeki, 1973, 1980, 1990a,b).

As we have already seen, qualitative observations concerning receptive neural wavelength selectivity is not a firm basis to establish these neurons as being devoted mainly to color. For example, the vast majority of neurons in the LGN respond with opponent signals, and these neurons have no orientation selectivity. Yet, we know that these neurons surely represent color, form, and motion information.

Moreover, the quality of the receptive-field measurements in area V4 has not achieved the same level of precision as measurements in the retina or area V1. Because these cells appear to be highly nonlinear, there are no widely agreed upon methods for fully characterizing their responses. Also, there have been disputes concerning even the qualitative properties of area-V4 receptive fields. For example, Desimone et al. (1985) report that many cells are selective for orientation, direction of motion, and spatial frequency. Like Zeki, these authors accept the basic logic of the neuron doctrine. They conclude from the variation of receptive field properties that "like V1, area V4 is not specialized to analyze one particular stimulus quality, but rather processes several in parallel" (p. 447). They then develop an alternative notion of the role of area V4 and later visual areas (Desimone and Schein, 1987).

Reasoning about the Cortex and Perception

While hypotheses about the roles of different cortical areas in perception are being debated, and although experiments to test these hypotheses have begun, we are at quite an early stage in our understanding of cortical function. This should not be too surprising; after all, the scientific investigation of the relationship between cortical responses and perception is a relatively new scientific endeavor, less than 100 years old. At this point in time we should expect some controversy and uncertainty regarding the status of early hypotheses. Much of the controversy stems from uncertainty concerning which experimental measurements will prove to be reliable sources of information and which will not.

In thinking about what we have learned about cortical function, I find it helpful to consider these two questions:

1. What do we want to know about cortical function?
2. What are the logical underpinnings of the experimental methods we have available to determine the relationship between cortical responses and perception?

When one discovers a new structure in the brain, it is almost impossible to refrain from asking: what does this part of the brain do? Once one poses this question, the answer is naturally formulated in terms of the *localization* of perceptual function. Our mindset becomes one of asking what happens where, rather than simply asking what happens. Hypotheses concerning the localization of function are the usual way to begin a research program on brain function. Moreover, I think any fair reading of the historical literature will show that hypotheses about what functions are localized within the brain region serve the useful purpose of organizing early experiments and theory.

On the other hand, in those portions of the visual pathways where our understanding is relatively mature, localization is rarely the central issue. We know that the retina is involved in visual function, and we know that some features of the retinal encoding are important for acuity, adaptation, wavelength encoding, and so forth. Our grasp of retinal function is sufficiently powerful so that we no longer frame questions about retinal function as a problem of localization. Instead, we pose problems in terms of the flow of information; we try to understand how information is represented and transformed within the retina.

For example, we know that information about the stimulus wavelength is represented by the relative absorption rates of the three cone photopigments. The information is not localized in any simple anatomical sense: no single neuron contains all the necessary information, nor are the neurons that represent wavelength information grouped together. Perhaps one might argue that acuity is localized, since acuity is greatest in the fovea. Even so, acuity depends on image formation, proper spacing of the photoreceptors, and appropriate representation of the photoreceptors in the optic tract. Without all of these other components in place, the observer will not have good visual acuity. The important questions about visual acuity are questions about the nature of the information and how the information is encoded and transmitted. That the fovea is the region of highest acuity is important, but not a solution to the question of how we resolve fine detail.

The most important questions about vision are those that Helmholtz posed: What are the principles that govern how the visual pathways make inferences from the visual image? How do we use image information to compute these perceptual inferences? We seek to understand these principles of behavior and neural representations with the same precision as we understand color-matching and the cone photopigments. We begin with spatial localization of brain function so that we can decide where to begin our work, not how to end it.

Thus, as our understanding becomes more refined we no longer formulate hypotheses based on localization of function alone. Instead, we use quantitative methods to compare neural responses and behavioral measurements. Mature areas of vision science relate perception and neural response by demonstrating correlations between the information

in the neural signals and the computations applied to those signals. The information contained in the neural response and the transformations applied to that information are the essence of perception.

Summary and Conclusions

Color appearance, like so much of vision, is an inference. Mainly, color is a perceptual representation of the surface reflectance of an object. There are two powerful obstacles that make it difficult to infer surface reflectance from the light incident at the eye. First, the reflected light confounds information about the surface and the illuminant. Second, the human eye has only three types of cones to encode a spectral signal consisting of many different wavelengths.

We began this chapter by asking what aspects of color imaging might make it feasible to perform this visual inference. Specifically, we studied how surface reflectance might be estimated from the light incident at the eye. We concluded that it is possible to draw accurate inferences about surface-reflectance functions when the surface and illuminant spectral curves are regular functions that can be well approximated by low-dimensional linear models. When the input signals are constrained, it is possible to design simple algorithms that use the cone absorptions to estimate surface reflectance accurately.

Next, we considered whether human judgments of color appearance share some of the properties used by algorithms that estimate surface reflectance. As a test of the correspondence between these abstract algorithms and human behavior, we reviewed how judgments of color appearance vary with changes in the illumination. Experimental results using the asymmetric color-matching method show that color-appearance judgments of targets seen under different illuminants can be predicted by matches between scaled responses of the human cones. The scale factor depends on the difference in illumination. To a large degree, these results are consistent with the general principle we have observed many times: judgments of color appearance are described mainly by the local contrast of the cone signals, not their absolute level. By basing color-appearance judgments on the scaled signal, which approximates the local cone contrast, color appearance is more closely correlated with surface reflectance than with the light incident at the eye.

Then we turned to a more general review of the organizational principles of color appearance. There are two important means of organizing color experience. Many color representations, like the Munsell representation, emphasize the properties of hue, saturation, and lightness. A second organizational theme is based on Hering's observation that red-green and blue-yellow are opponent-colors pairs and that we never experience these hues together in a single color. The opponent-colors organization has drawn considerable attention with the discovery that

many neurons carry opponent signals, increasing their response to some wavelengths of light and decreasing in response to others.

In recent years, there have been many creative and interesting attempts to study the representation of color information in the visual cortex. Most prominent among the hypotheses generated by this work is the notion that opponent-colors signals are spatially localized in the cortex. The evidence in support of this view comes from two types of experiments. First, clinical observations show that certain individuals lose their ability to perceive color although they retain high visual acuity. Second, studies of the receptive fields of individual neurons suggest that opponent-colors signals are represented in spatially localized brain areas. These hypotheses are new and unproven. But whether they are ultimately right or wrong, these hypotheses are the important opening steps in the modern scientific quest to understand the neural basis of conscious experience.

Exercises

1. Apart from surface reflectance, what other aspects of the material world might be represented by color vision?
2. Color vision has evolved separately in birds, fish, and primates. What does this imply about the significance of wavelength information in the environment?
3. We have used finite-dimensional linear models of surfaces and illuminants to understand how the visual system might infer surface reflectance from the responses of only three cone types.
 - (a) Offer a definition of a linear model that relates them to the discrete Fourier series.
 - (b) Describe the relationship between linear models used in this chapter and linear models used in spatial pattern sensitivity and multiresolution image representations.
 - (c) List some reasons why light-bulb manufacturers would be interested in linear models for illuminants.
 - (d) Using a linear model for daylight, we can design a device with two color sensors to estimate the relative spectral power distribution of daylight.
 - i. Suppose a device has sensor sensitivities S . Write an equation showing how the sensor responses depend on the illuminant model basis functions and the illuminant coefficients.
 - ii. Describe why you can solve this equation for the linear-model weights.
 - iii. Once you have calculated the linear-model weights, how do you calculate the estimated spectral power distribution?
4. (a) What was von Kries's hypothesis concerning the biological mechanisms of color constancy?
 - (b) How did Wassef test this theory experimentally?

- (c) What was Land's view of color science and its role in understanding color appearance?
- (d) In his recent book, Zeki (1993, p. 12) writes:

Why did colour constancy play such a subsidiary role in enquiries on colour vision? Almost certainly because, until only very recently it has been treated as a departure from a general rule, although it is in fact the central problem of colour vision. That general rule supposes that there is a precise and simple relationship between the wavelength composition of the light reaching the eye from every point on a surface and the colour of that point.

Zeki also quotes Newton's famous passage in support of his view:

Every Body reflects the rays of its own colour more copiously than the rest, and from their excess and predominance in the reflected light has its colour.

Does Newton's passage support Zeki's assertion that most authors thought it was the light at the eye that governs color appearance? Or, does Newton's passage support Helmholtz's view that surface reflectance determines appearance? What do *you* think, based on your reading of this chapter? Go to the historical literature and see what others have written on the topic.

5. (a) Design a neural wiring diagram such that the system will see yellow and blue at the same time, but not yellow and red.
- (b) Given the strong chromatic aberration of the eye, what spatial sensitivity properties would you assign to the blue-yellow neural pathway?
6. Stiles developed a sophisticated quantitative behavioral analysis of light adaptation to uniform fields. His papers were collected in a 1979 volume entitled *Mechanisms of Colour Vision*. Stiles's work makes substantial use of linear methods coupled with static nonlinearities. Read the experimental literature in this area and consider your answers to the following questions.
- (a) What role did homogeneity play in Stiles's analysis of the color mechanisms?
- (b) Which aspect of his measurements led him to believe that the data might be revealing the action of independent cone pathways?
- (c) As you read about Stiles's experiments, consider how you would perform experimental tests of superposition. Then, read the papers on the Stiles mechanisms by Pugh and Mollon (1979) and Wandell and Pugh (1980). How did they test superposition?
7. Speculate about why prosopagnosia (the inability to recognize familiar faces) and cerebral dyschromatopsia are frequently associated in patients. Make sure that you list some reasons that have to do with functional issues concerning visual computation and some that have to do with the organization of the neural and vascular structures related to vision. Design experimental tests of your hypotheses.