In OSA Handbook of Optics, 2cd edition, M. Bass (ed.), Vol. 1, Chapter 26.
McGraw-Hill, New York, 1995.

CHAPTER 26
COLORIMETRY

David H. Brainard
Department of Psychology
University of California, Santa Barbara
Santa Barbara, California

26.1 GLOSSARY

This glossary describes the major symbol usage for the body of the chapter. Symbols used in the appendix are generic. The following notational conventions hold throughout: (1) scalars are denoted with plain symbols, (2) vectors are denoted with lowercase bold symbols, (3) matrices are denoted with uppercase bold symbols.

$\mathbf{a}$	linear model weights
$\mathbf{B}$	linear model basis vectors
$\mathbf{b}$	spectral power distribution; basis vector
$\mathbf{M}$	color space transformation matrix
N_b	linear model dimension
N_λ	number of wavelength samples
$\mathbf{P}$	linear model for primaries
$\mathbf{p}$	primary spectral power distribution
$\mathbf{R}$	cone (or sensor) sensitivities
$\mathbf{r}$	cone (or sensor) coordinates
$\mathbf{T}$	color matching functions
$\mathbf{t}$	tristimulus coordinates
v	luminance
$\mathbf{v}$	luminous efficiency function
λ	wavelength

26.2 INTRODUCTION

Scope

The goal of colorimetry is to incorporate properties of the human color vision system into the measurement and specification of visible light. This branch of color science has been quite successful. We now have efficient quantitative representations that predict when two lights will appear identical to a human observer. Although such colorimetric representations

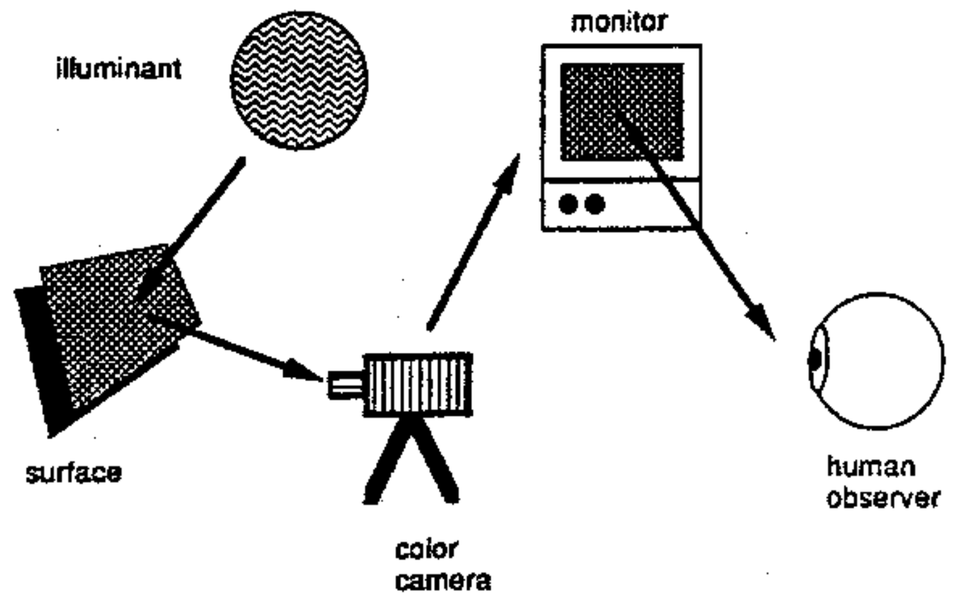

FIGURE 1 A typical image processing chain. Light reflects from a surface or collection of surfaces. This light is recorded by a color camera and stored in digital form. The digital image is processed by a computer and rendered on a color monitor. The reproduced image is viewed by a human observer.

do not directly predict the color sensation,[1-3] they do provide the foundation for the scientific study of color appearance. Moreover, colorimetry can be applied successfully in practical applications. Foremost among these is perhaps color reproduction.[4-7]

As an illustrative example, Fig. 1 shows an image processing chain. Light from an illuminant reflects from a collection of surfaces. This light is recorded by a color camera and stored in digital form. The digital image is processed by a computer and rendered on a color monitor. The reproduced image is viewed by a human observer. The goal of the image processing is to render an image with the same color appearance at each image location as the original. Although exact reproduction is not always possible with this type of system, the concepts and formulae of colorimetry do provide a reasonable solution.[4,8] To develop this solution, we will need to consider how to represent the spectral properties of light, the relation between these properties and color camera responses, the representation of the restricted set of lights that may be produced with a color monitor, and the way in which the human visual system encodes the spectral properties of light. We will treat each of these topics in this chapter, with particular emphasis on the role played by the human visual system.

Reference Sources

A number of excellent references are available that provide detailed treatments of colorimetry and its applications. Wyszecki and Stiles' comprehensive book[9] is an authoritative reference and provides numerous tables of standard colorimetric data. Pokorny and Smith[10] provide a handbook treatment complementary to the one developed here. Several publications of the Commission Internationale de l'Eclairage (International Commission on Illumination, commonly referred to as the CIE) describe current international technical standards for colorimetric measurements and calculations.[11-13] Other sources cover colorimetry's mathematical foundations,[14,15] its history,[16] its applications,[2,5,7,17] and its relation to neural mechanisms.[18,19] Chapters 27 and 28 of this volume and Vol. II, Chap. 24 are also relevant.

Chapter Overview

Colorimetry, Computers, and Linear Algebra. The personal computers that are now almost universally available in the laboratory can easily handle all standard colorimetric calculations. With this fact in mind, we have organized our treatment of colorimetry to allow direct translation between our formulation and its software implementation. In particular, we use vector and matrix representations throughout this chapter and use matrix algebra to express colorimetric formulae. Matrix algebra is being used increasingly in the colorimetric literature.[4.20–22] Appendix A reviews the elementary facts of matrix algebra required for this chapter. Numerous texts treat the subject in detail.[23–27] Various software packages provide extensive support for numerical matrix algebra.[28–31]

Chapter Organization. The rest of this chapter is organized into two main sections. The first section, "Fundamentals", reviews the empirical foundation of colorimetry and introduces basic colorimetric methods. The second section, "Topics", discusses a number of applications of colorimetry. It also includes a brief review of more advanced topics in color science.

26.3 FUNDAMENTALS

Stimulus Representation

Light at a Point. We describe the light reaching the eye from an image location by its spectral power distribution. The spectral power distribution generally specifies the radiant power density at each wavelength in the visible spectrum. For human vision, the visible spectrum extends roughly between 400 and 700 nm (but see "Sampling the Visible Spectrum" below). Depending on the viewing geometry, measures of radiation transfer other than radiant power may be used. These measures include radiance, irradiance, exitance, and intensity. The distinctions between these measures and their associated units are treated in Vol. II, Chap. 24 and are not considered here. That chapter also discusses measurement instrumentation and procedures.

Vector Representation of Spectral Functions. We will use discrete representations of spectral functions. Although a discrete representation samples the continuous functions of wavelength, the information loss caused by this sampling can be made arbitrarily small by increasing the number of sample wavelengths.

Suppose that spectral power density has been measured at N_λ discrete sample wavelengths $\lambda_1 \cdots \lambda_{N_\lambda}$, each separated by an equal wavelength step $\Delta\lambda$. As shown in Fig. 2, we can represent the measured spectral power distribution using an N_λ dimensional column vector **b**. The nth entry of **b** is simply the measured power density at the nth sample wavelength multiplied by $\Delta\lambda$. Note that the values of the sample wavelengths $\lambda_1 \cdots \lambda_{N_\lambda}$ and wavelength step $\Delta\lambda$ are not explicit in the vector representation. These values must be provided as side information when they are required for a particular calculation. In colorimetric applications, sample wavelengths are typically spaced evenly throughout the visible spectrum at steps of between 1 and 10 nm. We follow the convention that the entries of **b** incorporate $\Delta\lambda$, however, so that we need not represent $\Delta\lambda$ explicitly when we approximate integrals over wavelength.

Manipulation of Light. Intensity scaling is an operation that changes the overall power of a light at each wavelength without altering the relative power between any pair of wavelengths. One way to implement intensity scaling is to place a neutral density filter in

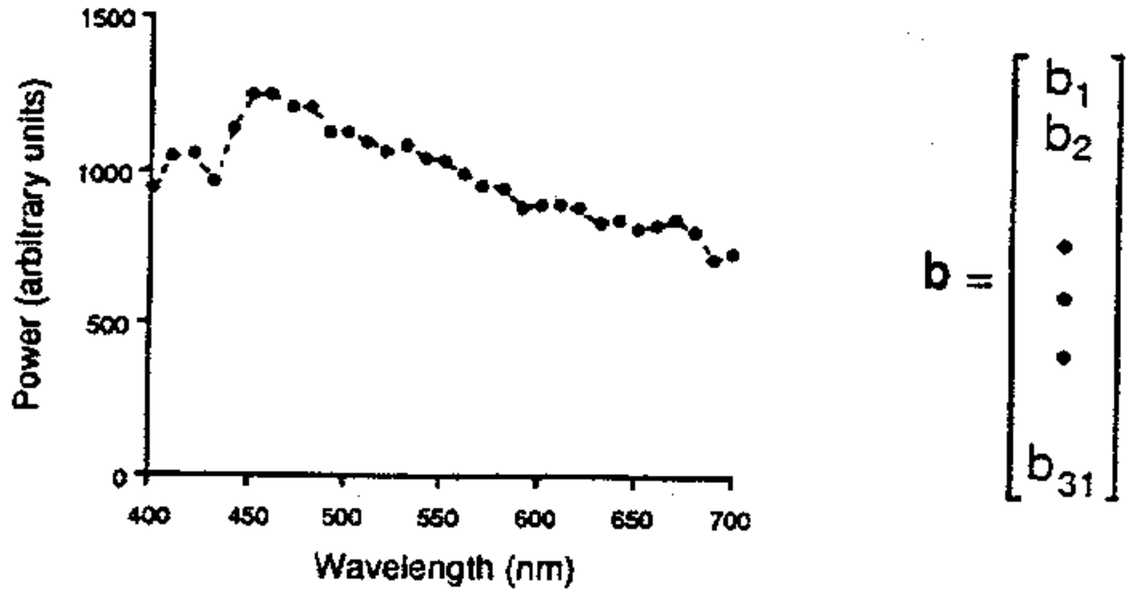

FIGURE 2 The vector representation of functions of wavelength. The plot shows a spectral power distribution measured at 10-nm intervals between 400 nm and 700 nm. Each point on the plot represents the power at a single sample wavelength. The vector **b** on the right depicts the vector representation of the same spectral power distribution. The **n**th entry of **b** is simply the measured power density at the **n**th sample wavelength times $\Delta\lambda$. Thus $\mathbf{b}_1$ is derived from the power density at 400 nm, $\mathbf{b}_2$ is derived from the power density at 410 nm, and $\mathbf{b}_{31}$ is derived from the power density at 700 nm.

the light path. The superposition of two lights is an operation that produces a new light whose power at each wavelength is the sum of the power in the original lights at the corresponding wavelength. One way to implement superposition is to use an optical beam splitter. The effects of both manipulations may be expressed using matrix algebra.

We use scalar multiplication to represent intensity scaling. If a light $\mathbf{b}_1$ is scaled by a factor a, then the result **b** is given by the equation $\mathbf{b} = \mathbf{b}_1 a$. The expression $\mathbf{b}_1 a$ represents a vector whose entries are obtained by multiplying the entries of the vector $\mathbf{b}_1 a$ by the scalar a. Similarly, we use vector addition to represent superposition. If we superimpose two lights $\mathbf{b}_1$ and $\mathbf{b}_2$, then the result **b** is given by the equation $\mathbf{b} = \mathbf{b}_1 + \mathbf{b}_2$. The expression $\mathbf{b}_1 + \mathbf{b}_2$ represents a vector whose entries are obtained by adding the entries of the vectors $\mathbf{b}_1$ and $\mathbf{b}_2$. Figures 3 and 4 depict both of these operations.

Linear Models for Spectral Functions. Intensity scaling and superposition may be used

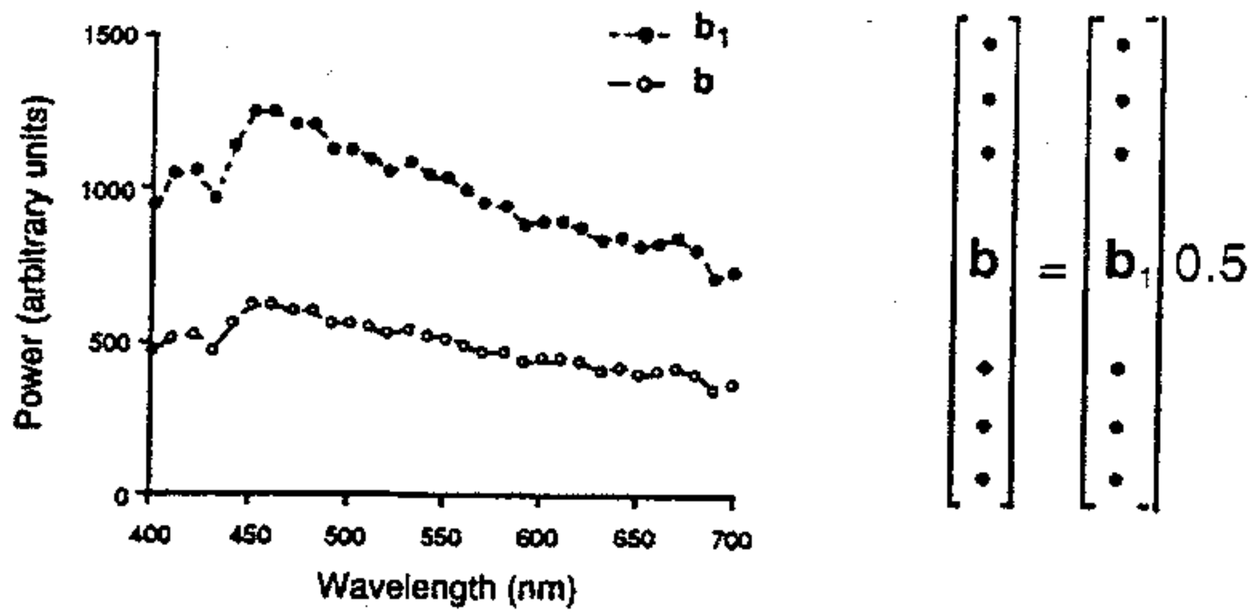

FIGURE 3 Representation of intensity scaling. Suppose that light **b** is created by reducing the power in light $\mathbf{b}_1$ by a factor of 0.5 at each wavelength. The result is shown graphically in the plot. The vector representation of the same relation is given by the equation $\mathbf{b} = \mathbf{b}_1\ 0.5$.

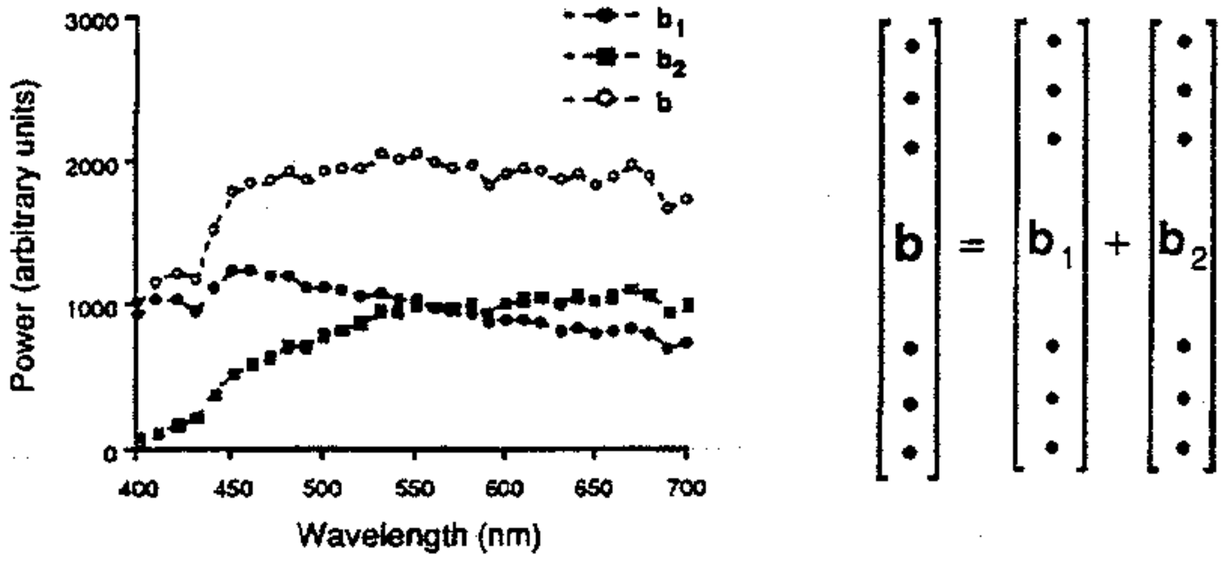

FIGURE 4 Representation of superposition. Suppose that light **b** is created by superimposing two lights **b** and $\mathbf{b}_2$. The result is shown graphically in the plot. The vector representation of the same relation is given by the equation $\mathbf{b} = \mathbf{b}_1 = \mathbf{b}_2$.

in combination to produce a wide range of spectral functions. Suppose that we have a set of N_b lights that we can individually scale and superimpose. Let the vectors $\mathbf{b}_1 \cdots \mathbf{b}_{N_b}$ represent the spectral power distributions of these lights. In this case, we can produce any spectral power distribution **b** that has the form

$$\mathbf{b} = \mathbf{b}_1 a_1 + \cdots + \mathbf{b}_{N_b} a_{N_b} \tag{1}$$

Suppose we know that a spectral function **b** is constrained to have the form of Eq. (1) where the vectors $\mathbf{b}_1 \cdots \mathbf{b}_{N_b}$ are known. Then we can specify **b** completely by providing the values of the scalars $a_1 \cdots a_{N_b}$. If the number of primaries N_b is less than the number of sample wavelengths N_λ, then this specification is more efficient (i.e., requires fewer numbers) than specifying the entries of **b** directly. We say that the spectral functions that satisfy Eq. (1) are described by (or lie within) a linear model. We call N_b the dimension of the linear model. We call the vectors $\mathbf{b}_1 \cdots \mathbf{b}_{N_b}$ the basis vectors for the model. We call the scalars $a_1 \cdots a_{N_b}$ required to construct any particular spectral function the model weights for that function.

Matrix Representation of Linear Models. Equation (1) can be written using vector and matrix notation. Let **B** be an N_λ by N_b dimensional matrix whose columns are the basis vectors $\mathbf{b}_1 \cdots \mathbf{b}_{N_b}$. We call **B** the basis matrix for the linear model. The composition of the basis matrix is shown pictorially on the left of Fig. 5. Let **a** be an N_b dimensional vector whose entries are the weights $a_1 \cdots a_{N_b}$. Figure 5 also depicts the vector **a**. Using **B** and **a** we can reexpress Eq. (1) as the matrix multiplication

$$\mathbf{b} = \mathbf{B}\mathbf{a} \tag{2}$$

The equivalence of Eqs. (1) and (2) may be established by direct expansion of the definition of matrix multiplication (see App. A). A useful working intuition for matrix multiplication is that the effect of multiplying a matrix times a vector (e.g., **Ba**) is to produce a new vector (e.g., **b**) that is a weighted superposition of the columns of the matrix (e.g., **B**), where the weights are given by the entries of the vector (e.g., **a**).

Use of Linear Models. When we know that a spectral function is described by a linear

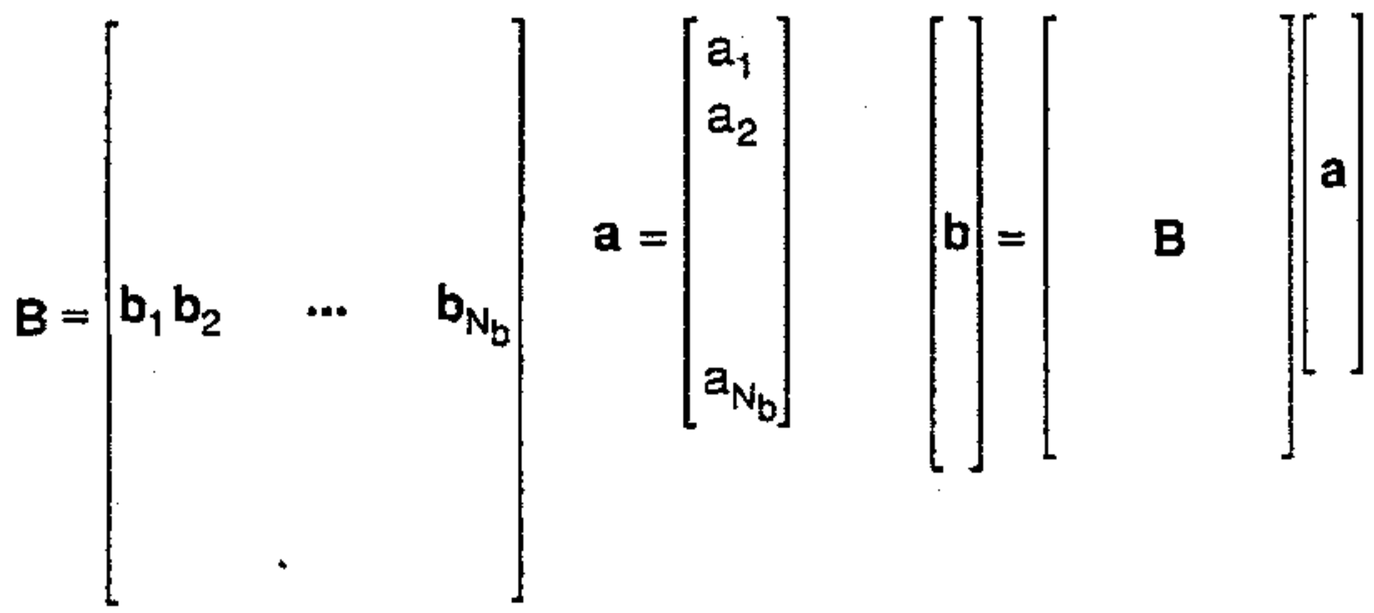

FIGURE 5 Vector representation of linear models. The matrix **B** represents the basis vectors of the linear model. The vector **a** represents the model weights required to form a particular spectral power distribution **b**. The relation between **b**, **a**, and **B** is given by Eq. (2) and is depicted on the right of the figure.

model, we can specify it by using the weight vector **a**. The matrix **B**, which is determined by the basis vectors, specifies the side information necessary to convert the vector **a** back to the discrete wavelength representation. When we represent spectral functions in this way, we say that we are representing the functions within the specified linear model.

Representing spectral functions within a small-dimensional linear model places strong constraints on the form of the functions. As the dimension of the model grows, linear models can represent progressively wider classes of functions. In many cases of interest, there is prior information that allows us to assume that spectra are indeed described by a linear model. A common example of this situation is the light emitted from a computer-controlled color monitor. Such a monitor produces different spectral power distributions by scaling the intensity of the light emitted by three different types of phosphor (see Vol. I, Chap. 27). Thus the emitted light lies within a three-dimensional linear model whose basis vectors are given by the emission spectra of the monitor's phosphors. Linear model constraints also turn out to be useful for describing naturally occurring surface and illuminant spectra (see later under "Surfaces and Illuminants").

Note that representing spectral functions within linear models is a generalization of, rather than an alternative to, the more traditional wavelength representation. To understand this, we need only note that we can choose the basis vectors of the linear model to be discrete delta functions centered at each sample wavelength. We refer to this special choice of basis vectors as the *identity basis* or *wavelength basis*. We refer to the corresponding linear model as the *identity model*. For the identity model, the basis matrix **B** is the N_λ by N_λ identity matrix, where N_λ is the number of sample wavelengths. The identity matrix contains 1's along its main diagonal and 0's elsewhere. Multiplying the identity matrix times any vector simply results in the same vector. From Eq. (2), we can see that when **B** is the identity matrix, the representation of any light **b** within the linear model is simply **a** = **b**.

Sampling the Visible Spectrum. To use a discrete representation for functions of wavelength, it is necessary to choose a sampling range and sampling increment. Standard practice varies considerably. The current recommendation of the CIE is that the visible spectrum be sampled at 5-nm increments between 380 and 780 nm.[11] Coarser sampling at 10 nm between 400 and 700 nm provides an adequate approximation for many applications. In this chapter, we provide tabulated spectral data at the latter wavelength sampling. Where possible, we provide references to where they may be obtained at finer sampling increments. In cases where a subset of the spectral data required for a calculation is not available, interpolation or extrapolation may be used to estimate the missing values.[11]

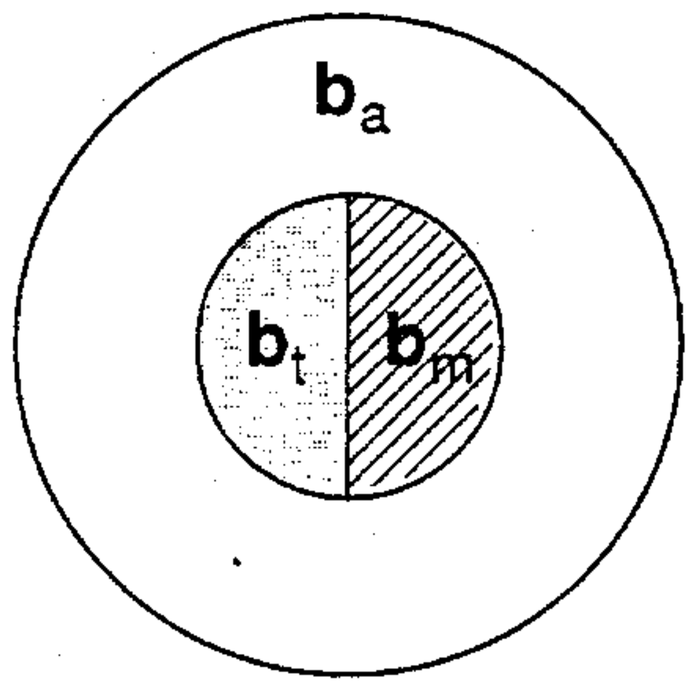

FIGURE 6 Schematic of the stimulus for the basic color-matching experiment. The observer views a bipartite field. A test is presented on one side of a bipartite field and a matching light is presented on the other side of the field. The test light's spectral power distribution $\mathbf{b}_t$ is under experimental control. The observer adjusts the spectral power distribution $\mathbf{b}_m$ of the matching light. The observer's task is to make the two halves of the bipartite field appear identical. The bipartite field is often surrounded by an annulus whose spectral power distribution $\mathbf{b}_a$ is held fixed during the matching process.

Empirical Foundations of Colorimetry

The Basic Color-Matching Experiment. In the absence of linear model constraints, the only complete representation of the spectral properties of light is the full spectral power distribution. The human visual system, however, does not encode all of the information available in the spectral power distribution. Understanding how the visual system encodes the spectral properties of light leads to an efficient representation of the information relevant to human vision. To develop this understanding, we consider the color-matching experiment.

The basic color-matching experiment is illustrated in Fig. 6. An observer views a bipartite field, as shown in the figure. Each side of the field is spatially uniform, but the two sides may differ in their spectral power distributions. The spectral power distribution $\mathbf{b}_t$ of one side of the bipartite field is under experimental control. We call this the *test light.* The spectral power distribution $\mathbf{b}_m$ of the light on the other side of the bipartite field is under the observer's control. We call this the *matching light.* The observer's task in the experiment is to adjust the matching light so that it appears identical in color to the test light. The bipartite field may be surrounded by an annulus whose spectral power distribution $\mathbf{b}_a$ is held fixed during the matching process.

Typically, the apparatus for the color-matching experiment is arranged so that the observer can adjust the matching light by controlling the intensity scaling of some number of superimposed primary lights. It is not a priori clear that it will be possible for the observer to make a match when the number of primaries is small. The results of a large number of color-matching studies, however, show that observers are able to set matches to any test light by adjusting the intensity of just three primaries. This result implies that there exist lights with different spectral power distributions that cannot be distinguished by a human observer. From the color-matching experiment, we may conclude that the human

visual system does not encode all of the information available in the full spectral power distribution.

We use the symbol ~ to indicate that two lights are a perceptual match. Perceptual matches are to be carefully distinguished from physical matches, which are denoted by the = symbol. Of course, when two lights are a physical match, they must be a perceptual match. Two lights that are a perceptual match but not a physical match are referred to as *metameric color stimuli* or *metamers.* The term *metamerism* is often used to refer to the fact that two physically different lights can appear identical.

Conditions for Trichromatic Color Matching. There are a number of qualifications to the empirical generalization that it is possible for observers to match any test light by adjusting the intensities of just three primaries. Some of these qualifications have to do with ancillary restrictions on the experimental conditions (e.g., the size of the bipartite field and the overall intensity of the test and matching lights). We discuss these later under "Limits of the Color-Matching Experiment." The other qualifications have to do with the choice of primaries and certain conventions about the matching procedure. First, the primaries must be chosen so that it is not possible to match any one of them with a weighted superposition of the other two. Second, the observer sometimes wishes to increase the intensity of one or more of the primaries above its maximum value. In this case, we must allow the observer to scale the intensity of the test light down. We follow the convention of saying that the match was possible, but scale up the reported primary weights by the same factor. Third, the observer sometimes wishes to decrease the intensity of one or more of the primaries below zero. In this case, we must allow the observer to superimpose each such primary on the test light rather than on the other primaries. We follow the convention of saying that the match was possible, but report with a negative sign the intensity of each transposed primary.

With these qualifications, matching with three primaries is always possible. This fact is often referred to as the *trichromacy* of normal human color vision.

Tristimulus Coordinates. The basic color-matching experiment is an empirical procedure that maps any light to three numbers: the weights on the color-matching apparatus primaries required to make a match. These three weights are often called the *tristimulus coordinates* of the test light. Given the qualifications and conventions discussed earlier, there are no restrictions on the magnitude of tristimulus coordinates. Thus, the matching light's spectral power distribution is described completely by a three-dimensional linear model whose basis vectors are the primary lights' spectral power distributions. The tristimulus coordinates of a test light are precisely the linear model weights required to form the matching light.

We denote the primary spectral power distributions by the vectors $\mathbf{p}_1 \cdots \mathbf{p}_3$. The associated linear model matrix $\mathbf{P}$ contains these vectors in its three columns. We denote the tristimulus coordinates of a light using the three dimensional vector $\mathbf{t}$. Thus we can use the tristimulus coordinates of any test light $\mathbf{b}$ to reconstruct a matching light $\mathbf{Pt}$ such that

$$\mathbf{b} \sim \mathbf{Pt} \tag{3}$$

We emphasize that, in general, $\mathbf{Pt}$ will not be equal to $\mathbf{b}$.

Critical Properties of Color Matching. Our discussion of the color-matching experiment suggests that we can represent the spectral properties of light using tristimulus coordinates. To ensure that this representation is general, however, we need to consider whether color matching exhibits a number of critical properties. We review these properties briefly

below. Given that they hold, it is possible to show that tristimulus coordinates provide a representation for the spectral properties of light. Krantz provides a detailed formal treatment.[14]

Reflexivity, Symmetry, and Transitivity. Reflexivity is the requirement that a light match itself. Symmetry is the requirement that if two lights match, they will continue to match when their roles in the color-matching experiment are reversed. Transitivity is the requirement if two lights each match a common third light, then they will match each other. Apart from small failures that arise from variability in observer's judgments, these three properties do generally hold for human color matching.[14]

Uniqueness of Color Matches. The tristimulus coordinates of a light should be unique. This is equivalent to the requirement that only one weighted combination of the apparatus primaries produces a match to any given test light. The uniqueness of color matches ensures that tristimulus coordinates are well defined. In conjunction with transitivity, uniqueness also guarantees that two lights that match each other will have identical tristimulus coordinates. It is generally accepted that, apart from variability, trichromatic color matches are unique for color-normal observers.

Persistance of Color Matches. The above properties concern color matching under a single set of viewing conditions. By viewing conditions, we refer to the properties of the image surrounding the bipartite field and the sequence of images viewed by the observer before the match was made. An important property of color matching is that lights that match under one set of viewing conditions continue to match when the viewing conditions are changed. This property is referred to as the *persistence* or *stability* of color matches.[9,18] It holds to good approximation for a wide range of viewing conditions. We discuss conditions where the persistence law fails later in the chapter. The importance of the persistence law is that it allows a single set of tristimulus values to be used across viewing conditions.

Consistency Across Observers. Finally, for the use of tristimulus coordinates to have general validity, it is important that there be good agreement about matches across observers. For the majority of the population, there is good agreement about which lights match. We discuss individual differences in color matching in a later section.

Computing Tristimulus Coordinates. *Grassmann's Laws.* Tristimulus coordinates provide an efficient representation for the spectral properties of light. For this representation to be useful, we require a model of the color-matching experiment that allows us to compute tristimulus coordinates from spectral power distributions. To develop such a model, we rely on two regularities of color matching that were first characterized by Grassmann.[32] These (sometimes in conjunction with the properties discussed above) are generally referred to as Grassmann's laws. They have been tested extensively and hold well.[9,18]

The two regularities may be expressed as follows: (1) (proportionality law) if two lights match, they will continue to match if they are both scaled by the same factor, (2) (additivity law) if two pairs of lights match each other, then the superposition of the two will also match. We can express the two laws formally as follows. The proportionality law states:

$$\text{if } \mathbf{b}_1 \sim \mathbf{b}_2, \text{ then } \mathbf{b}_1 a \sim \mathbf{b}_2 a \tag{4}$$

where a is a scalar that represents any intensity scaling. The additivity law states:

$$\text{if } \mathbf{b}_1 \sim \mathbf{b}_2 \text{ and } \mathbf{b}_3 \sim \mathbf{b}_4, \text{ then } \mathbf{b}_1 + \mathbf{b}_3 \sim \mathbf{b}_2 + \mathbf{b}_4 \tag{5}$$

The proportionality law allows us to determine the relation between the tristimulus coordinates of a light and the tristimulus values of a scaled version of that light. Suppose that $\mathbf{b} \sim \mathbf{Pt}$. Applying the proportionality law, we conclude that for any scalar a, we have $\mathbf{b}a \sim (\mathbf{Pt})a$. Because matrix multiplication is associative, we can conclude that:

$$\text{if } \mathbf{b} \sim \mathbf{Pt}, \text{ then } \mathbf{b}a \sim \mathbf{P}(\mathbf{t}a) \tag{6}$$

This means that the tristimulus coordinates of a light $\mathbf{b}a$ may be obtained by scaling the tristimulus coordinates $\mathbf{t}$ of the light $\mathbf{b}$. A similar argument shows that the additivity law determines the relation between the tristimulus coordinates of two lights and the tristimulus coordinates of their superposition:

$$\text{if } \mathbf{b}_1 \sim \mathbf{Pt}_1 \text{ and } \mathbf{b}_2 \sim \mathbf{Pt}_2, \text{ then } \mathbf{b}_1 + \mathbf{b}_2 \sim \mathbf{P}(\mathbf{t}_1 + \mathbf{t}_2) \tag{7}$$

Implication of Grassmann's Laws. If the tristimulus coordinates of the basis vectors for a linear model are known, then Grassmann's laws allow us to determine the tristimulus coordinates of any light within the linear model. Let $\mathbf{t}_1 \cdots \mathbf{t}_{N_b}$ be the tristimulus coordinates corresponding to the model basis vectors and let $\mathbf{T}_B$ be the 3 by N_b matrix whose columns are $\mathbf{t}_1 \cdots \mathbf{t}_{N_b}$. For any light $\mathbf{b}$ within the linear model, we can write that $\mathbf{b} = \mathbf{Ba}$. By expanding this matrix product and applying Eqs. (6) and (7), it is possible to show that the tristimulus coordinates of $\mathbf{b}$ are given by the matrix product:

$$\mathbf{t} = \mathbf{T}_B\mathbf{a} \tag{8}$$

Equation (8) is very important. It tells how to compute the tristimulus coordinates for any light within a linear model from the tristimulus coordinates for each of the basis vectors. Thus a small number of color matches (one for each of the basis vectors) allows us to predict color matches for a large number of lights. We call the rows of the matrix $\mathbf{T}_B$ the color matching functions with respect to the linear model defined by the columns of $\mathbf{B}$. This is a generalization of the standard usage of the term *color-matching functions,* which is introduced below.

Color-Matching Functions. Let $\mathbf{T}$ be the corresponding matrix of tristimulus values for the basis vectors of the identity model. In this case, $\mathbf{T}$ has dimensions 3 by N_λ, where N_λ is the number of sample wavelengths. Each column of $\mathbf{T}$ is the tristimulus coordinates for a monochromatic light. Within the identity model, the representation of any light $\mathbf{b}$ is simply $\mathbf{b}$ itself. From Eq. (8) we conclude directly that the tristimulus values for any light are given by

$$\mathbf{t} = \mathbf{Tb} \tag{9}$$

Once we know the tristimulus coordinates for a set of monochromatic lights centered at each of the sample wavelengths, we use Eq. (9) to compute the tristimulus coordinates of any light.

We can regard each of the rows of $\mathbf{T}$ as a function of wavelength. We refer to these

functions as a set of color matching functions. Color matching functions are often plotted as a function of wavelength, as illustrated in Fig. 7. It is important to note, however, they do not represent spectral power distributions.

Mechanisms of Color Matching

Cone Mechanisms. The basic color-matching experiment can be understood in terms of the action of retinal photoreceptors. Three distinct classes of cone photoreceptors are generally believed to participate in normal color vision. These are often referred to as the long (L), middle (M), and short (S) wavelength sensitive cones. The information a cone provides about a light is mediated by a single number: the rate of photopigment isomerizations caused by the absorption of light quanta.[33] If two lights produce identical absorption rates in all three classes of cones, the visual system will not be able to distinguish them. The conventional mechanistic explanation for the results of the color-matching experiment is thus that two lights match if and only if they result in the same number of photopigment absorptions in all three classes of cones.[18]

To compute a cone's quantal absorption rate from a light's spectral power distribution, we use the cone's spectral sensitivity function. For each class of cones, the function specifies the number of quanta that will be absorbed for a monochromatic light of unit power centered on each of the sample wavelengths. When the application is to predict an observer's color matches, the sensitivity function should incorporate wavelength-dependent filtering by the ocular media. To compute the quantal absorption rate, we multiply the light power by the cone's sensitivity at each wavelength and then sum the results over wavelength.

Suppose that we represent the spectral sensitivity functions of the three classes of cones by the rows of a 3 by N_λ matrix **R**. Let **r** be a three-dimensional vector whose entries are the cone quantal absorption rates. For a light with spectral power distribution **b**, we can compute the absorption rates through the matrix equation

$$\mathbf{r} = \mathbf{Rb} \tag{10}$$

This computation accomplishes the wavelength-by-wavelength multiplication and summation for each cone class.

Figure 8 shows estimates of human cone spectral sensitivity functions. The three curves are independently normalized to a maximum of one. We discuss these estimates more thoroughly later.

Cone Coordinates. We use the term *cone coordinates* to refer to the vector **r**. We can relate cone coordinates to tristimulus coordinates in a straightforward manner. Suppose that in a color-matching experiment performed with primaries **P** we find that a light **b** has tristimulus coordinates **t**. From our mechanistic explanation and Eq. (10) we have

$$\mathbf{r} = \mathbf{Rb} = \mathbf{RPt} \tag{11}$$

If we define the matrix $\mathbf{M}_{T,R} = (\mathbf{RP})$, we see that the tristimulus coordinates of a light are related to its cone coordinates by a linear transformation

$$\mathbf{r} = \mathbf{M}_{T,R}\mathbf{t} \tag{12}$$

By comparing Eq. (9) with Eq. (11) and noting that these equations hold for any light **b**, we derive

$$\mathbf{R} = \mathbf{M}_{R,T}\mathbf{T} \tag{13}$$

Equation (13) has the interesting implication that the color-matching experiment determines the cone sensitivities up to a free linear transformation.

Common Color Coordinate Systems

Color Coordinate Systems. For the range of conditions where the color-matching experiment obeys the properties described in the previous sections, tristimulus coordinates (or cone coordinates) provide a complete and efficient representation for human color vision. When two lights have identical tristimulus coordinates, they are indistinguishable to the visual system and may be substituted for one another. When two lights have tristimulus coordinates that differ substantially, they can be distinguished by an observer with normal color vision.

The relation between spectral power distributions and tristimulus coordinates depends on the choice of primaries used in the color-matching experiment. In this sense, the choice of primaries in colorimetry is analogous to the choice of unit (e.g., foot versus meter) in the measurement of length. We use the terms color coordinate system and color space to refer to a representation derived with respect to a particular choice of primaries. We also use the term color coordinates as a synonym for tristimulus coordinates.

Although the choice of primaries determines a color space, specifying primaries alone is not sufficient to compute tristimulus coordinates. Rather, it is the color-matching functions that characterize the properties of the human observer with respect to a particular set of primaries. Knowledge of the color-matching functions allows us to compute tristimulus coordinates through Eq. (9). As we show below, however, knowledge of a single set of color-matching functions also allows us to derive color-matching functions with respect to other sets of primaries. Thus in practice we can specify a color space either by its primaries or by its color-matching functions.

A large number of different color spaces are in common use. The choice of which color space to use in a given application is governed by a number of considerations. If all that is of interest is to use a three-dimensional representation that accurately predicts the results of the color-matching experiment, the choice revolves around the question of finding a set of color-matching functions that accurately capture color-matching performance for the set of observers and viewing conditions under consideration. From this point of view, color spaces that differ only by an invertible linear transformation are equivalent. But there are other possible uses for color representation. For example, one might wish to choose a space that makes explicit the responses of the physiological mechanisms that mediate color vision. We discuss a number of commonly used color spaces below.

Stimulus Spaces. A stimulus space is the color space determined by the primaries of a particular apparatus. For example, stimuli are often specified in terms of the excitation of three monitor phosphors. Stimulus color spaces have the advantage that they provide a direct description of the physical stimulus. On the other hand, they are nonstandard and their use hampers comparison of data collected in different laboratories. A useful compromise is to transform the data to a standard color space, but to provide enough side information to allow exact reconstruction of the stimulus. Often this side information can take the form of a linear model whose basis functions are the apparatus primaries.

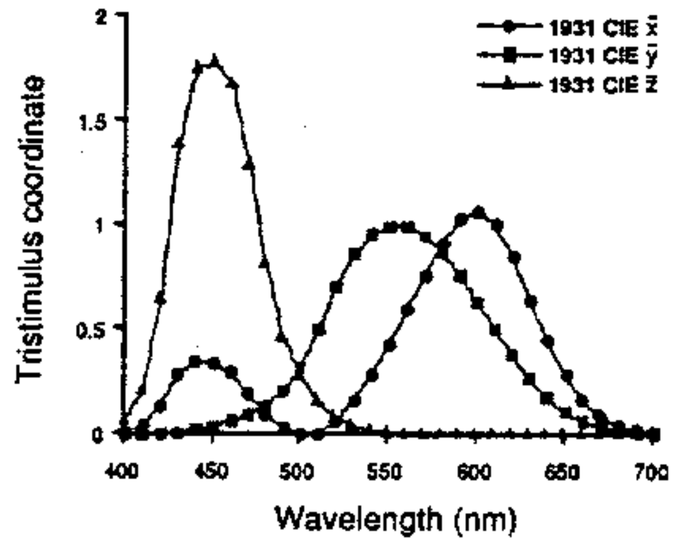

FIGURE 7 Color-matching functions. As described in the text, the rows of the matrix T may be viewed as a function of wavelength. Color-matching functions are often plotted as a function of wavelength. The plot shows a set of color-matching functions for normal human vision standardized by the CIE in 1931. The three individual functions are usually referred to as the CIE $\bar{x}$, $\bar{y}$, and $\bar{z}$ color-matching functions.

CIE 1931 Color-Matching Functions. In 1931, the CIE chose a standard set of primaries and integrated a large body of empirical data to determine a standard set of color-matching functions. The notion was that these functions would describe the results of a color-matching experiment performed on an "average" color-normal human observer. The result of the standardization process was a set of color-matching functions which are often called the CIE 1931 color-matching functions.[12] These functions are plotted as a function of wavelength in Fig. 7 and tabulated in Table 1. Tristimulus coordinates computed with respect to these color-matching functions are called CIE 1931 *XYZ* tristimulus coordinates. The hypothetical observer whom the CIE 1931 color-matching functions describe is often referred to as the CIE 1931 standard observer.

There is now some evidence that the color-matching functions standardized by the CIE in 1931 are slightly different from those of the average human observer.[9,34] A large body of extant data is available only in terms of the CIE 1931 system, however, and many colorimetric instruments are designed around it. Therefore, it seems likely that the CIE 1931 system will continue to be of practical importance for some time.

Judd-Vos Modified Color-Matching Functions. In 1951 Judd reconsidered the 1931 color-matching functions and came to the conclusion that they could be improved.[35] Vos[36] refined Judd's analysis. The Judd-Vos modifications lead to a set of color-matching functions that are probably more typical of the average human observer than the original CIE 1931 color-matching functions. These functions were never officially standardized. They are the basis of a number of estimates of the human cone spectral sensitivities and are thus widely used in practice, especially in vision science. We provide the Judd-Vos modified *XYZ* color-matching functions in Table 2.

1964 10° Color-Matching Functions. In 1964, the CIE standardized a second set of color-matching functions appropriate for larger field sizes. These color-matching functions take into account the fact that human color matches depend on the size of the matching fields. The CIE 1964 10° color-matching functions are an attempt to provide a standard

TABLE 1 CIE 1931 *XYZ* Color-Matching Functions

Wavelengths are in nm. The data are available at 1-nm intervals between 360 nm and 830 nm.[12]

Wavelength	X	Y	Z
400	1.43E-02	3.96E-04	6.79E-02
410	4.35E-02	1.21E-03	2.07E-01
420	1.34E-01	4.00E-03	6.46E-01
430	2.84E-01	1.16E-02	1.39E+00
440	3.48E-01	2.30E-02	1.75E+00
450	3.36E-01	3.80E-02	1.77E+00
460	2.91E-01	6.00E-02	1.67E+00
470	1.95E-01	9.10E-02	1.29E+00
480	9.56E-02	1.39E-01	8.13E-01
490	3.20E-02	2.08E-01	4.65E-01
500	4.90E-03	3.23E-01	2.72E-01
510	9.30E-03	5.03E-01	1.58E-01
520	6.33E-02	7.10E-01	7.82E-02
530	1.66E-01	8.62E-01	4.22E-02
540	2.90E-01	9.54E-01	2.03E-02
550	4.33E-01	9.95E-01	8.70E-03
560	5.95E-01	9.95E-01	3.90E-03
570	7.62E-01	9.52E-01	2.10E-03
580	9.16E-01	8.70E-01	1.65E-03
590	1.03E+00	7.57E-01	1.10E-03
600	1.06E+00	6.31E-01	8.00E-04
610	1.00E+00	5.03E-01	3.40E-04
620	8.54E-01	3.81E-01	2.00E-04
630	6.42E-01	2.65E-01	0.00E+00
640	4.48E-01	1.75E-01	0.00E+00
650	2.84E-01	1.07E-01	0.00E+00
660	1.65E-01	6.10E-02	0.00E+00
670	8.74E-02	3.20E-02	0.00E+00
680	4.68E-02	1.70E-02	0.00E+00
690	2.27E-02	8.20E-03	0.00E+00
700	1.14E-02	4.10E-03	0.00E+00

observer for these larger fields. The use of 10° color-matching functions is recommended by the CIE when the sizes of the regions under consideration are larger than 4°.[11] We provide the CIE 10° color-matching functions in Table 3.

Cone Sensitivities. For some applications it is desirable to relate the color stimulus to the responses of the cone mechanisms. A number of psychophysical and physiological techniques have been developed for determining the foveal cone spectral sensitivities.[9,34,37,38] At present, the most widely used estimates are those of Smith and Pokorny,[37,39] which are based on the Judd-Vos modified *XYZ* color-matching functions. Stockman, MacLeod, and Johnson,[34] however, argue for a set of small-field cone sensitivities derived from the CIE 1964 10° color-matching functions. We provide both the Smith-Pokorny and Stockman-MacLeod-Johnson estimates in Tables 4 and 5. Other estimates exist.[9,38]

Note that there is no universally accepted standard for scaling the sensitivities of each cone class relative to one another. One common method normalizes the sensitivities to a maximum of one for each cone class. Care should be taken when drawing conclusions that depend on the scaling chosen.

Opponent and Modulation Spaces. Cone coordinates are useful because they make explicit the responses of the initial physiological mechanisms thought to mediate color vision. A number of investigators have begun to use representations that attempt to

TABLE 2 Judd-Vos Modified *XYZ* Color-Matching Functions

Wavelengths are in nm. The data are available at 5-nm increments between 380 nm and 825 nm.[36]

Wavelength	X	Y	Z
400	3.80E-02	2.80E-03	1.74E-01
410	9.99E-02	7.40E-03	4.61E-01
420	2.29E-01	1.75E-02	1.07E+00
430	3.11E-01	2.73E-02	1.47E+00
440	3.33E-01	3.79E-02	1.62E+00
450	2.89E-01	4.68E-02	1.47E+00
460	2.33E-01	6.00E-02	1.29E+00
470	1.75E-01	9.10E-02	1.11E+00
480	9.19E-02	1.39E-01	7.56E-01
490	3.17E-02	2.08E-01	4.47E-01
500	4.85E-03	3.23E-01	2.64E-01
510	9.29E-03	5.03E-01	1.54E-01
520	6.38E-02	7.10E-01	7.66E-02
530	1.67E-01	8.62E-01	4.14E-02
540	2.93E-01	9.54E-01	2.00E-02
550	4.36E-01	9.95E-01	8.78E-03
560	5.97E-01	9.95E-01	4.05E-03
570	7.64E-01	9.52E-01	2.28E-03
580	9.16E-01	8.70E-01	1.81E-03
590	1.02E+00	7.57E-01	1.23E-03
600	1.06E+00	6.31E-01	9.06E-04
610	9.92E-01	5.03E-01	4.29E-04
620	8.43E-01	3.81E-01	2.56E-04
630	6.33E-01	2.65E-01	9.77E-05
640	4.41E-01	1.75E-01	5.12E-05
650	2.79E-01	1.07E-01	2.42E-05
660	1.62E-01	6.10E-02	1.19E-05
670	8.58E-02	3.20E-02	5.60E-06
680	4.58E-02	1.70E-02	2.79E-06
690	2.22E-02	8.21E-03	1.31E-06
700	1.11E-02	4.10E-03	6.48E-07

Source: From "Colorimetric and photometric properties of a two degree fundamental observer." J. J. Vos, *Color Research and Application*, Copyright © 1978. Reprinted by permission of John Wiley & Sons, Inc.

represent the responses of subsequent mechanisms. Two basic ideas underlie these representations. The first is the general opponent processing model described later under "Opponent Process Model." We call representations based on this idea *opponent color spaces*. The second idea is that stimulus contrast is more relevant than stimulus magnitude.[40] We call spaces that are based on this second idea *modulation color spaces*. Some color spaces are both opponent and modulation color spaces.

Cone Modulation Space. To derive coordinates in the cone modulation color space, the stimulus is first expressed in terms of its cone coordinates. The cone coordinates of a white point are then chosen. Usually these are the cone coordinates of a uniform adapting field or the spatio-temporal average of the cone coordinates of the entire image sequence. The cone coordinates of the white point are subtracted from the cone coordinates of the stimulus and the resulting differences are normalized by the corresponding cone coordinate of the white point.

The DKL Color Space. Derrington, Krauskopf, and Lennie[41] introduced an opponent modulation space that is now widely used. This space is closely related to the chromaticity diagram suggested by MacLeod and Boynton.[42] To derive coordinates in the DKL color space, the stimulus is first expressed in cone coordinates. As with cone modulation space,

TABLE 3 CIE 1964 10° *XYZ* Color-Matching Functions

Wavelengths are in nm. The data are available at 1-nm intervals between 360 nm and 830 nm.[12]

Wavelength	X	Y	Z
400	1.91E-02	2.00E-03	8.60E-02
410	8.47E-02	8.80E-03	3.89E-01
420	2.05E-01	2.14E-02	9.73E-01
430	3.15E-01	3.87E-02	1.55E+00
440	3.84E-01	6.21E-02	1.97E+00
450	3.71E-01	8.95E-02	1.99E+00
460	3.02E-01	1.28E-01	1.75E+00
470	1.96E-01	1.85E-01	1.32E+00
480	8.05E-02	2.54E-01	7.72E-01
490	1.62E-02	3.39E-01	4.15E-01
500	3.80E-03	4.61E-01	2.19E-01
510	3.75E-02	6.07E-01	1.12E-01
520	1.18E-01	7.62E-01	6.07E-02
530	2.37E-01	8.75E-01	3.05E-02
540	3.77E-01	9.62E-01	1.37E-02
550	5.30E-01	9.92E-01	4.00E-03
560	7.05E-01	9.97E-01	0.00E+00
570	8.79E-01	9.56E-01	0.00E+00
580	1.01E+00	8.69E-01	0.00E+00
590	1.12E+00	7.77E-01	0.00E+00
600	1.12E+00	6.58E-01	0.00E+00
610	1.03E+00	5.28E-01	0.00E+00
620	8.56E-01	3.98E-01	0.00E+00
630	6.48E-01	2.84E-01	0.00E+00
640	4.32E-01	1.80E-01	0.00E+00
650	2.68E-01	1.08E-01	0.00E+00
660	1.53E-01	6.03E-02	0.00E+00
670	8.13E-02	3.18E-02	0.00E+00
680	4.09E-02	1.59E-02	0.00E+00
690	1.99E-02	7.70E-03	0.00E+00
700	9.60E-03	3.70E-03	0.00E+00

the cone coordinates of a white point are then subtracted from the cone coordinates of the stimulus of interest. The next step is to reexpress the resulting difference as tristimulus coordinates with respect to a new choice of primaries that are thought to isolate the responses of postreceptoral mechanisms.[43,44] The three primaries are chosen so that modulating two of them does not change the response of the photopic luminance mechanism (discussed later). The color coordinates corresponding to these two primaries are often called the constant B and constant R and G coordinates. Modulating the constant R and G coordinate of a stimulus modulates only the S cones. Modulating the constant B coordinate modulates both the L and M cones but keeps the S cone response constant. Because the constant R and G coordinate is not allowed to change the response of the photopic luminance mechanism, the DKL color space is well defined only if the S cones do not contribute to luminance. The third primary of the space is chosen so that it has the same relative cone coordinates as the white point. The coordinate corresponding to this third primary is called the luminance coordinate. Flitcroft[45] provides a detailed treatment of the DKL color space.

Caveats. The basic ideas underlying the use of opponent and modulation color spaces seem to be valid. On the other hand, there is not general agreement about how signals from cones are combined into opponent channels, about how this combination depends on adaption, or about how adaptation affects signals originating in the cones. Since a specific

TABLE 4 Smith-Pokorny Estimates of the Cone Sensitivities

Wavelengths are in nm. The data are available at 1-nm increments between 400 nm and 700 nm.[37]

Wavelength	L	M	S
400	2.66E-03	2.82E-03	1.08E-01
410	6.89E-03	7.67E-03	2.85E-01
420	1.58E-02	1.89E-02	6.59E-01
430	2.33E-02	3.17E-02	9.08E-01
440	3.01E-02	4.77E-02	1.00E+00
450	3.43E-02	6.35E-02	9.10E-01
460	4.12E-02	8.60E-02	7.99E-01
470	6.27E-02	1.30E-01	6.89E-01
480	1.02E-01	1.89E-01	4.68E-01
490	1.62E-01	2.67E-01	2.76E-01
500	2.63E-01	3.96E-01	1.64E-01
510	4.23E-01	5.95E-01	9.56E-02
520	6.17E-01	8.08E-01	4.74E-02
530	7.73E-01	9.41E-01	2.56E-02
540	8.83E-01	9.97E-01	1.24E-02
550	9.54E-01	9.87E-01	5.45E-03
560	9.93E-01	9.22E-01	2.53E-03
570	9.97E-01	8.06E-01	1.44E-03
580	9.65E-01	6.51E-01	1.16E-03
590	8.94E-01	4.77E-01	8.12E-04
600	7.95E-01	3.18E-01	6.10E-04
610	6.70E-01	1.93E-01	3.12E-04
620	5.30E-01	1.10E-01	1.98E-04
630	3.80E-01	5.83E-02	9.03E-05
640	2.56E-01	2.96E-02	5.25E-05
650	1.59E-01	1.44E-02	2.51E-05
660	9.14E-02	6.99E-03	1.44E-05
670	4.82E-02	3.33E-03	7.56E-06
680	2.57E-02	1.64E-03	4.02E-06
690	1.24E-02	7.50E-04	1.94E-06
700	6.21E-03	3.68E-04	9.72E-07

model of these processes is implicit in any opponent or modulation color space, coordinates in these spaces must be treated carefully. This is particularly true of modulation spaces, where the relation between the physical stimulus and coordinates in the space depends on the choice of white point. As a consequence, radically different stimuli can have identical coordinates in a modulation space. For example, 100 percent contrast monochromatic intensity gratings are all represented by the same coordinates in modulation color spaces, independent of their wavelength. Nonetheless, such stimuli appear very different to human observers. Identity of coordinates in a modulation color space does not imply identity of appearance across different choices of white points.

Transformations Between Color Spaces

Because of the large number of color spaces currently in use, the ability to transform data between various color spaces is of considerable practical importance. The derivation of such transformations depends on what is known about the source and destination color spaces. Below we discuss cases where both the source and destination color space are derived from the same underlying observer (i.e., when the source and destination color

TABLE 5 Stockman-MacLeod-Johnson Estimates of Cone Sensitivities

Wavelengths are in nm. The data are available at 5-nm increments between 390 nm and 730 nm.[34]

Wavelength	L	M	S
400	2.23E-03	1.93E-03	5.74E-02
410	8.76E-03	8.46E-03	2.39E-01
420	1.75E-02	2.05E-02	5.28E-01
430	2.69E-02	3.88E-02	8.03E-01
440	3.84E-02	6.36E-02	9.87E-01
450	4.88E-02	8.79E-02	9.50E-01
460	6.52E-02	1.20E-01	8.12E-01
470	9.69E-02	1.76E-01	6.51E-01
480	1.38E-01	2.38E-01	3.94E-01
490	1.88E-01	3.04E-01	2.09E-01
500	2.97E-01	4.48E-01	1.19E-01
510	4.59E-01	6.47E-01	6.21E-02
520	6.34E-01	8.33E-01	2.94E-02
530	7.76E-01	9.46E-01	1.28E-02
540	8.83E-01	1.00E+00	5.61E-03
550	9.42E-01	9.77E-01	2.50E-03
560	9.88E-01	9.19E-01	1.15E-03
570	9.99E-01	8.03E-01	5.44E-04
580	9.68E-01	6.45E-01	2.63E-04
590	9.25E-01	4.88E-01	1.31E-04
600	8.36E-01	3.35E-01	6.65E-05
610	7.12E-01	2.10E-01	3.46E-05
620	5.64E-01	1.22E-01	1.83E-05
630	4.15E-01	6.93E-02	9.94E-06
640	2.71E-01	3.41E-02	5.49E-06
650	1.66E-01	1.57E-02	3.08E-06
660	9.40E-02	7.70E-03	1.76E-06
670	4.99E-02	3.69E-03	1.03E-06
680	2.51E-02	1.75E-03	6.07E-07
690	1.22E-02	8.33E-04	3.64E-07
700	5.87E-03	3.95E-04	2.22E-07

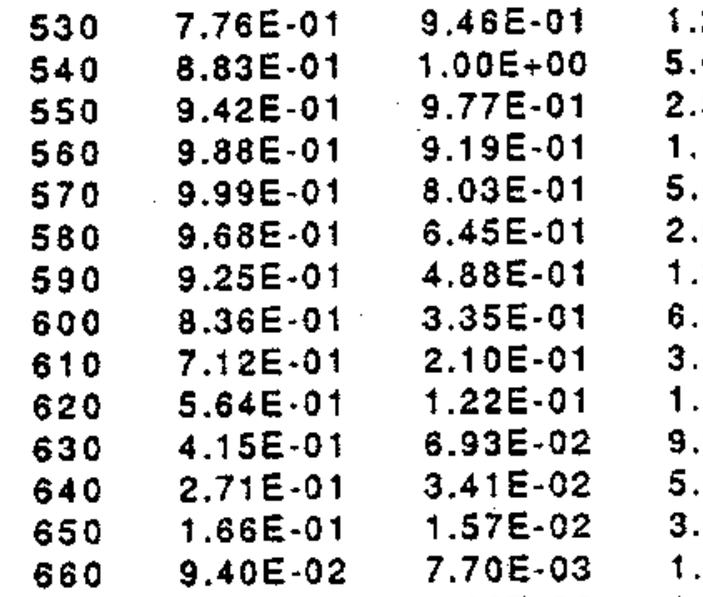

FIGURE 8 Human cone spectral responsivity functions normalized to a maximum of one. These estimates are due to Smith and Pokorny.[37,39]

TABLE 6 Color Space Transformations

The table summarizes how to form the matrix **M** that transforms color coordinates between two spaces.

Spectral functions known

Source	Destination	M	Notes
Primaries P_1	CMFs T_2	$M = T_2P_1$	
CMFs T_1	Primaries P_2	$M = (T_1P_2)^{-1}$	
Primaries P_1	Primaries P_2	$M = (T\ P_2)^{-1}(T\ P_1)$	T is any set of CMFs.
CMFs T_1	CMFs T_2	$T_2 = M\ T_1$	Use regression to find M.

One space specified in terms of other

Known tristimulus coordinates	How to construct M
Source primaries known in destination space.	Put them in columns of M.
Source CMFs known in destination space.	Put them in rows of M^{-1}.
Destination primaries known in source space.	Put them in columns of M^{-1}.
Destination CMFs known in source space.	Put them in rows of M.

* CMFs stands for Color Matching Functions

spaces both predict identical color matches). Table 6 summarizes these transformations. When the source and destination color spaces are characterized by a different underlying observer, the transformation is more difficult and often cannot be done exactly. We discuss possible approaches in a later section.

Source Primaries and Destination Color-Matching Functions Known. Let $\mathbf{P}_1$ be the matrix representing a set of known primaries. Let $\mathbf{T}_2$ be the matrix representing a known set of color-matching functions (e.g., the CIE 1931 *XYZ* color-matching functions). We would like to determine a transformation between the color coordinate system specified by $\mathbf{P}_1$ and that specified by $\mathbf{T}_2$. For example, linearized frame buffer values input to a computer-controlled color monitor may be thought of as tristimulus coordinates in a color space defined by the monitor's phosphor emission spectra. This transformation thus allows computation of the CIE 1931 *XYZ* tristimulus coordinates from the linearized frame buffer values.

We start by using Eq. (9) to compute the tristimulus coordinates, with respect to $\mathbf{T}_2$, for all three primary lights specified by $\mathbf{P}_1$. Each of these primaries is contained in a column of $\mathbf{P}_1$, so that we may perform this calculation directly through the matrix multiplication

$$\mathbf{M}_{P,T} = \mathbf{T}_2\mathbf{P}_1 \tag{14}$$

Let the matrix $\mathbf{P}_2$ represent the destination primaries. We do not need to know these explicitly, only that they exist. The meaning of Eq. (14) is that

$$\mathbf{P}_1 \sim \mathbf{P}_2\mathbf{M}_{P,T} \tag{15}$$

where we have generalized the symbol "~" to denote a column-by-column visual match for the matrices on both sides of the relation. This relation follows because the columns of $\mathbf{M}_{P,T}$ specify how the destination primaries should be mixed to match the source primaries. Equation (15) tells us that we can substitute the three lights represented by the columns of $\mathbf{P}_2\mathbf{M}_{P,T}$ for the three lights represented by the columns of $\mathbf{P}_1$ in any color-matching

experiment. In particular, we may make this substitution for any light **b** with tristimulus coordinates $\mathbf{t}_1$ in the source color coordinate system. We have

$$\mathbf{b} \sim \mathbf{P}_1\mathbf{t}_1 \sim \mathbf{P}_2\mathbf{M}_{P,T}\mathbf{t}_1 \tag{16}$$

By inspection, this tells us that the three-dimensional vector

$$\mathbf{t}_2 = \mathbf{M}_{P,T}\mathbf{t}_1 \tag{17}$$

is the tristimulus coordinates of **b** in the destination color coordinate system. Equation (17) provides us with the means to transform tristimulus coordinates from a coordinate system where the primaries are known to one where the color-matching functions are known. The transformation matrix $\mathbf{M}_{P,T}$ required to perform the transformation depends only on the known primaries $\mathbf{P}_1$ and the known color-matching functions $\mathbf{T}_2$. Given these, $\mathbf{M}_{P,T}$ may be computed directly from Eq. (14).

Source Color-Matching Functions and Destination Primaries Known. A second transformation applies when the color-matching functions in the source color space and the primaries in the destination color space are known. This will be the case, for example, when we wish to render a stimulus specified in terms of CIE 1931 tristimulus coordinates on a calibrated color monitor.

Let $\mathbf{T}_1$ represent the known color-matching functions and $\mathbf{P}_2$ represent the known primaries. By applying Eq. (17) we have that the relation between source tristimulus coordinates and the destination tristimulus coordinates is given by $\mathbf{t}_1 = \mathbf{M}_{P,T}\mathbf{t}_2$. This is a system of linear equations that we may solve to find an expression for $\mathbf{t}_2$ in terms of $\mathbf{t}_1$. In particular, as long as the matrix $\mathbf{M}_{P,T}$ is nonsingular, we can convert tristimulus coordinates using the relation

$$\mathbf{t}_2 = \mathbf{M}_{T,P}\mathbf{t}_1 \tag{18}$$

where we define

$$\mathbf{M}_{T,P} = (\mathbf{M}_{P,T})^{-1} = (\mathbf{T}_1\mathbf{P}_2)^{-1} \tag{19}$$

Source and Destination Primaries Known. A third transformation applies when the primaries of both the source and destination color spaces are known. One application of this transformation is to generate matching stimuli on two different calibrated monitors.

Let $\mathbf{P}_1$ and $\mathbf{P}_2$ represent the two sets of primaries. Let **T** represent a set of color-matching functions for any human color coordinates system. (There is no requirement that the color-matching functions be related to either the source or the destination primaries. For example, the CIE 1931 *XYZ* color-matching functions might be used.) To do the conversion, we simply use Eq. (17) to transform from the color coordinate system described by $\mathbf{P}_1$ to the coordinate system described by **T**. Then we use Eq. (18) to transform from the coordinates system described by **T** to the coordinate system described by $\mathbf{P}_2$. The overall transformation is given by

$$\mathbf{t}_2 = \mathbf{M}_{P,P}\mathbf{t}_1 = (\mathbf{M}_{T,P_2})(\mathbf{M}_{P_1,T})\mathbf{t}_1 = (\mathbf{T}\mathbf{P}_2)^{-1}(\mathbf{T}\mathbf{P}_1)\mathbf{t}_1 \tag{20}$$

It should not be surprising that this transformation requires the specification of a set of color-matching functions. These color-matching functions are the only source of information about the human observer in the transformation equation.

Source and Destination Color-Matching Functions Known. Finally, it is sometimes of

interest to transform between two color spaces that are specified in terms of their color-matching functions. An example is transforming between the space defined by the Judd-Vos modified *XYZ* color-matching functions and the space defined by the Smith-Pokorny cone fundamentals.

Let $\mathbf{T}_1$ and $\mathbf{T}_2$ represent the source and destination color-matching functions. Our development above assures us that there is some three-by-three transformation matrix, call it $\mathbf{M}_{T,T}$, that transforms color coordinates between the two systems. Recall that the columns of $\mathbf{T}_1$ and $\mathbf{T}_2$ are themselves tristimulus coordinates for corresponding monochromatic lights. Thus $\mathbf{M}_{T,T}$ must satisfy

$$\mathbf{T}_2 = \mathbf{M}_{T,T}\mathbf{T}_1 \tag{21}$$

This is a system of linear equations where the entries of $\mathbf{M}_{T,T}$ are the unknown variables. This system may be solved using standard regression methods. Once we have solved for $\mathbf{M}_{T,T}$, we can transform tristimulus coordinates using the equation

$$\mathbf{t}_2 = \mathbf{M}_{T,T}\mathbf{t}_1 \tag{22}$$

The transformation specified by Eq. (22) will be exact as long as the two sets of color-matching functions $\mathbf{T}_1$ and $\mathbf{T}_2$ characterize the performance of the same observer. One sometimes wishes, however, to transform between two color spaces that are defined with respect to different observers. For example, one might want to convert CIE 1931 *XYZ* tristimulus values to Judd-Vos modified tristimulus values. Although the regression procedure described here will still produce a transformation matrix in this case, the result of the transformation is not guaranteed to be correct.[4] We return to this topic later.

Interpreting the Transformation Matrix. It is useful to interpret the rows and columns of the matrices derived above. Let $\mathbf{M}$ be a matrix that maps the color coordinates from a source color space to a destination color space. Both source and destination color spaces are associated with a set of primaries and a set of color-matching functions. From our derivations above, we can conclude that the columns of $\mathbf{M}$ are the coordinates of the source primaries in the destination color space [see Eq. (14)] and the rows of $\mathbf{M}$ provide the destination color-matching functions with respect to the linear model whose basis functions are the primaries of source color space. Similarly, the columns of $\mathbf{M}^{-1}$ are the coordinates of the destination primaries in the source color-matching space and the rows of $\mathbf{M}^{-1}$ are the source color-matching functions with respect to the linear model whose basis functions are the primaries of the destination color space. Thus in many cases it is possible to construct the matrix $\mathbf{M}$ without full knowledge of the spectral functions. This can be of practical importance. For example, monitor manufacturers often specify the CIE 1931 *XYZ* tristimulus coordinates of their monitors' phosphors. In additional, colorimeters that measure tristimulus coordinates directly are often more readily available than spectral radiometers.

Transforming Primaries and Color-Matching Functions. We have shown that color coordinates in any two color spaces may be related by applying a linear transformation $\mathbf{M}$. The converse is also true. If we pick any nonsingular linear transformation $\mathbf{M}$ and apply it to a set of color coordinates we have defined a new color space that will successfully predict color matches. The color-matching functions for this new space will be given by $\mathbf{T}_2 = \mathbf{M}\mathbf{T}_1$. A set of primaries for the new space will be given by $\mathbf{P}_2 = \mathbf{P}_1\mathbf{M}^{-1}$. These derived primaries are not unique. Any set of primaries that match the constructed primaries will also work.

The fact that new color spaces can be constructed by applying linear transformations has an important implication for the study of color. If we restrict attention to what we may conclude from the color-matching experiment, we can determine the psychological representation of color only up to a free linear transformation. There are two attitudes

one can take toward this fact. The conservative attitude is to refrain from making any statements about the nature of color vision that depend on a particular choice of color space. The other is to appeal to experiments other than the color-matching experiment to choose a privileged representation. At present, there is not universal agreement about how to choose such a representation and we therefore advocate the conservative approach.

Visualizing Color Data

A challenge facing today's color scientist is to produce and interpret graphical representations of color data. Because the visual representation of light is three-dimensional, it is difficult to plot this representation on a two-dimensional page. Even more difficult is to represent a dependent measure of visual performance as a function of color coordinates. We discuss several approaches.

Three-Dimensional Approaches. One strategy is to plot the three-dimensional data in perspective, as shown on the top of Fig. 9. In many cases, the projection viewpoint may be chosen to provide a clear view of the regularities of interest in the data. The three-dimensional structure of the data may be emphasized by the addition of various monocular depth cues. A number of computer-graphics packages now provide facilities to aid in the preparation of three-dimensional perspective plots. Often these programs allow variation of the viewpoint and automatic inclusion of monocular depth cues.

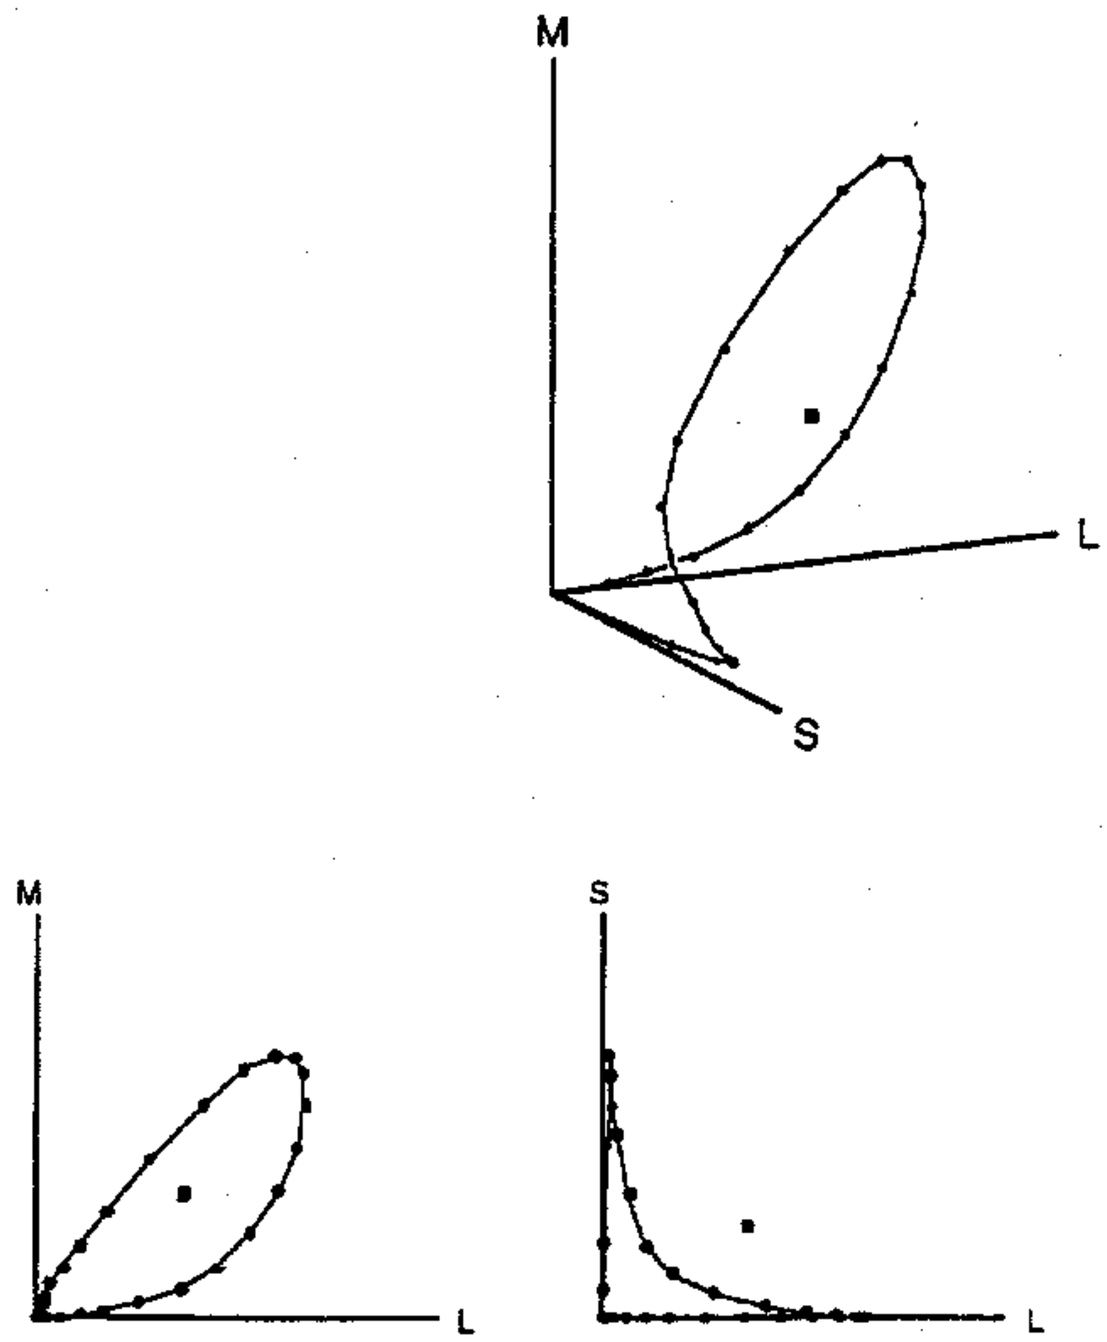

FIGURE 9 Three dimensional views of color data. The figure shows the color coordinates of an equal energy spectrum in color space defined by the human cone sensitivities (closed circles) and the color coordinates of CIE daylight D65 (closed squares). The top panel shows the data in perspective. The bottom panels show three two-dimensional views of the same data.

A second approach to showing the three-dimensional structure of color data is to provide multiple two-dimensional views, as in a drafter's sketch. This approach is shown on the bottom of Fig. 9.

Computer display technology provides promise for improved methods of viewing three-dimensional data. For example, it is now possible to produce computer animations that show plots that vary over time. Such plots have the potential for representing multidimensional data in a manner that is more comprehensible to a human viewer than a static plot. Other interesting possibilities include the use of stereo depth cues and color displays. At present, the usefulness of these techniques is largely confined to exploratory data analysis because there are no widely accepted standards for publication.

Chromaticity Diagrams. A second strategy for plotting color data is to reduce the dimensionality of the data representation. One common approach is through the use of chromaticity coordinates. Chromaticity coordinates are defined so that any two lights with the same relative color coordinates have identical chromaticity coordinates. That is, the chromaticity coordinates of a light are invariant with respect to intensity scaling. Because chromaticity coordinates have one less degree of freedom than color coordinates, they can be described by just two numbers and plotted in a plane. We call a plot of chromaticity coordinates a *chromaticity diagram.* A chromaticity diagram eliminates all information about the intensity of a stimulus.

There are many ways to normalize color coordinates to produce a set of chromaticity coordinates. When the underlying color coordinates are CIE 1931 *XYZ* tristimulus coordinates, it is conventional to use CIE 1931 *xy* chromaticity coordinates. These are given by

$$x = X/(X + Y + Z)$$
$$y = Y/(X + Y + Z) \tag{23}$$

A graphical way to understand these chromaticity coordinates is illustrated in Fig. 10.

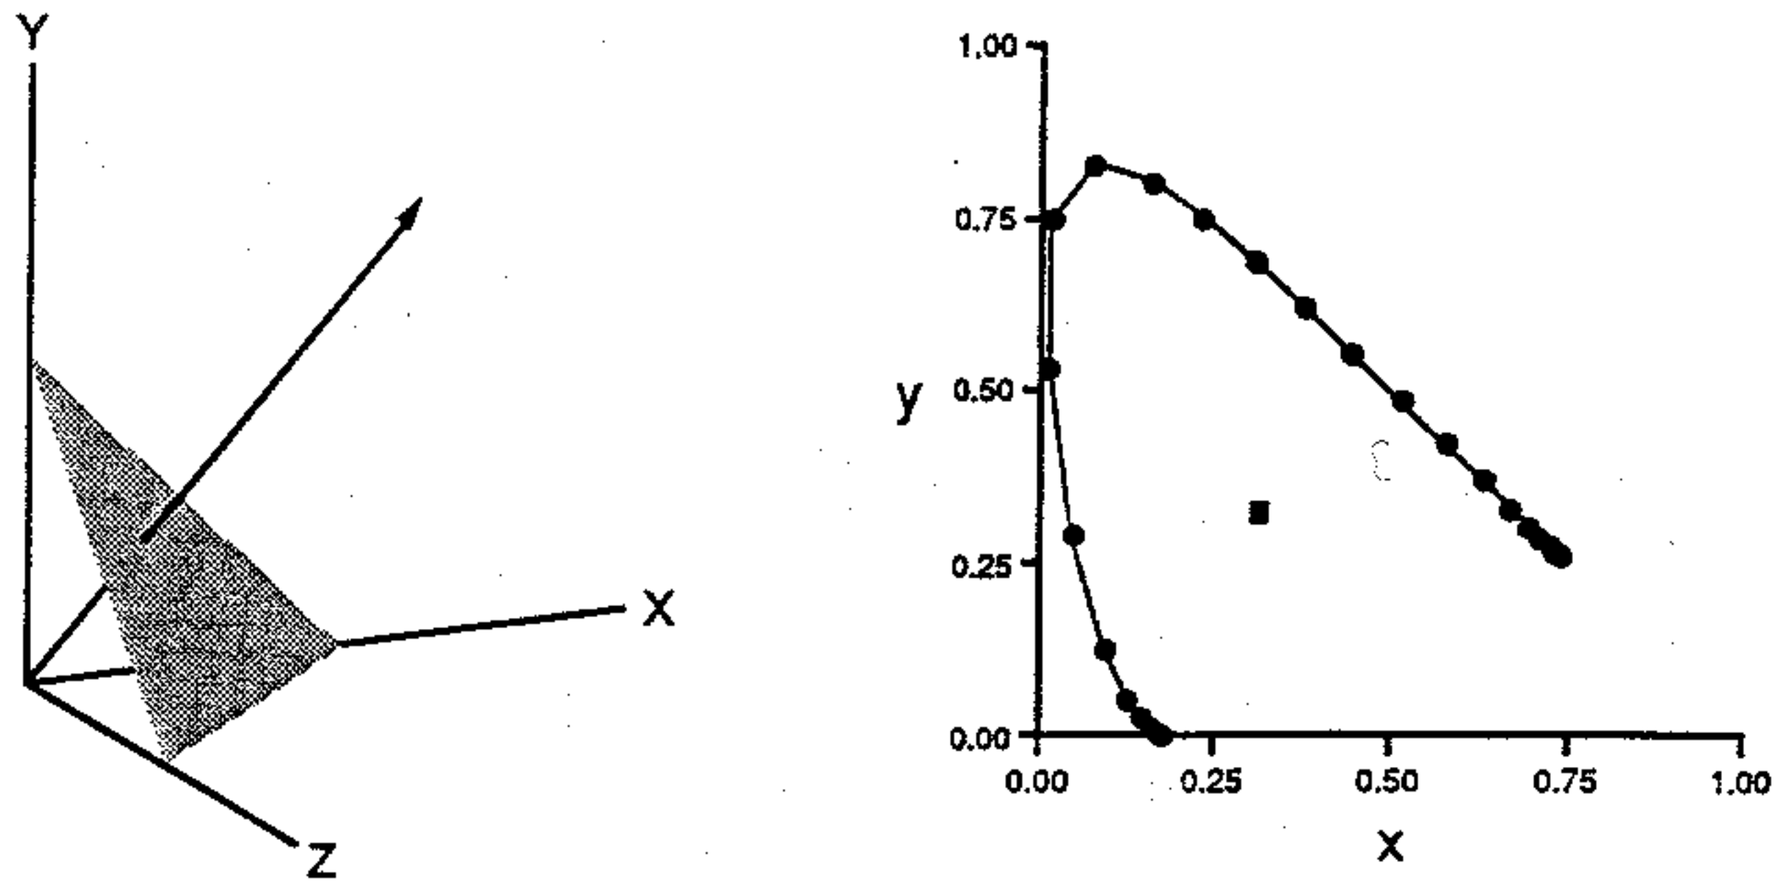

FIGURE 10 CIE 1931 *xy* chromaticity diagram. The left side of the figure shows a perspective view of the CIE 1931 *XYZ* tristimulus space. The ray shows a locus of points with constant chromaticity coordinates. The actual chromaticity coordinates for each ray are determined by where the ray intersects the plane described by the equation $X + Y + Z = 1$. This plane is indicated. The X and Y tristimulus values at the point of intersection are the x and y chromaticity coordinates for the ray. The right side of the figure shows the chromaticity coordinates for the ray. The right side of the figure shows the chromaticity coordinated of an equal energy spectrum (closed circles) and of CIE daylight D65 (closed square).

Wyszecki and Stiles[9] review a number of other standard chromaticity diagrams. MacLeod and Boynton[42] propose a chromaticity diagram that is derived from cone coordinates.

A useful property of most chromaticity diagrams is that the chromaticity coordinates of the mixture of two lights is always a weighted combination of chromaticity coordinates of the individual lights. This is easily verified for the CIE 1931 *xy* chromaticity diagram by algebraic manipulation. Thus the chromaticity of a mixture of lights will plot somewhere on the chord connecting the chromaticities of the individual lights.

Implicit in the use of chromaticity coordinates is the assumption that scalar multiplication of the stimuli does not affect the visual performance being plotted. If the overall intensity of the stimuli matter, then the use of chromaticity coordinates can obscure important regularities. For example, the shape of color discrimination contours (see "Color Discrimination" later in chapter) depends on how the overall intensity of the stimuli covaries with their chromaticities. Yet these contours are often plotted on a chromaticity diagram. This practice can lead to misinterpretation of the discrimination data. We recommend that plots of chromaticity coordinates be treated with some caution.

Functions of Wavelength. Color data are often represented as functions of wavelength. The wavelength spectrum parameterizes a particular path through the three-dimensional color space. The exact path depends on how overall intensity covaries with wavelength. For an equal energy spectrum, the path is illustrated in Fig. 9.

Wavelength representations are particularly useful in situations where knowing the value of a function for the set of monochromatic stimuli provides a complete characterization of performance. Color-matching functions, for example, are usefully plotted as functions of wavelength because these functions may be used to predict the tristimulus coordinates of any light. Plots of detection threshold versus wavelength, on the other hand, cannot be used to predict the detection threshold for arbitrary lights.[46] Just as the chromaticity diagram tends to obscure the potential importance of manipulating the overall intensity of light, wavelength representations tend to obscure the potential importance of considering mixtures of monochromatic lights.

Colorimetric measurements

To apply the formulae described in this chapter, it is often necessary to measure the colorimetric properties of stimuli. The most general approach is to measure the full spectral power distribution of the stimuli. Standard instrumentation and methods are discussed in Vol. II, Chap. 24. Often, however, it is not necessary to know the full spectral power distribution; knowledge of the tristimulus coordinates (in some standard color space) is sufficient. For example, the color space transformations summarized earlier in Table 6 depend on the full spectral power distributions of the primaries only through their tristimulus coordinates.

Specialized instruments, called *colorimeters,* can measure tristimulus coordinates directly. These instruments typically operate using the same principles as photometers (see Chap. 24 of Vol. II of this Handbook) with the exception that they have three calibrated filters rather than just one. Each filter mimics the spectral shape of one of the color-matching functions. Wyszecki and Stiles discuss colorimeter design.[9] Colorimeters are generally less expensive than radiometers and are thus an attractive option when full spectral data is not required.

Two caveats are worth note. First, it is technically difficult to design filters that exactly match a desired set of color-matching functions. Generally, commercial colorimeters are calibrated so that they give accurate readings for stimuli with broad spectral power

distributions. For narrowband stimuli (e.g., the light emitted by the red phosphor of many color monitors) the reported readings may be quite inaccurate. Second, most colorimeters are designed to the CIE 1931 standard. This may not be an optimal choice for the purpose of predicting the matches of an average human observer.

26.4 TOPICS

Surfaces and Illuminants

As shown in Fig. 1, the light reaching the eye is often formed when light from an illuminant reflects from a surface. Illuminants and surfaces are likely to be of interest in color reproduction applications involving inks, paints, and dyes, and in lighting design applications.

Reflection Model. Illuminants are specified by their spectral power distributions. We will use the vector **e** to represent the illuminant spectral power distributions. In general, the interaction of light with matter is quite complex (see Vol. I, Chaps. 7 and 9). For many applications, however, a rather simple model is acceptable. Using this model, we describe a surface by its surface reflectance function. The surface reflectance function specifies, for each sample wavelength, the fraction of illuminant power that is reflected to the observer. We will use the vector **s** to represent surface reflectance spectra. Each entry of **s** gives the reflectance measured at a single sample wavelength. Thus the spectral power distribution **b** of the reflected light is given by the wavelength-by-wavelength product of the illuminant spectral power distribution and the surface reflectance function.

The most important consideration neglected in this formulation is viewing geometry. The relation between the radiant power emitted by a source of illumination, the material properties of a surface, and the radiant power reaching an observer can depend strongly on the viewing geometry. In our formulation, these geometrical factors must be incorporated implicitly into the specification of the illuminant and surface properties, so that any actual calculation is specific to a particular viewing geometry. Moreover, the surface reflectance must be understood as being associated with a particular image location, rather than with a particular object. A great deal of current research in photorealistic computer graphics is concerned with accurate and efficient ways to specify illuminants and surfaces for spatially complex scenes.[7,47] A second complexity that we neglect is fluorescence.

Computing the Reflected Light. The relation between the surface reflectance function and the reflected light spectral power distribution is linear if the illuminant spectral power distribution is held fixed. We form the N_λ by N_λ diagonal illuminant matrix **E** whose diagonal entries are the entries of **e**. This leads to the relation $\mathbf{b} = \mathbf{Es}$. By substituting into Eq. (9), we arrive at an expression for the tristimulus coordinates of the light reflected from a surface

$$\mathbf{t} = (\mathbf{TE})\mathbf{s} \tag{24}$$

The matrix (**TE**) in this equation plays exactly the same role as the color-matching functions do in Eq. (9). Any result that holds for spectral power distributions may thus be directly extended to a result for surface reflectance functions when the illuminant is known and held fixed.

Linear Model Representations for Surfaces and Illuminants. Judd, MacAdam, and Wyszecki[48] measured the spectral power distributions of a large number of naturally occurring daylights. They determined that a four-dimensional linear model provided a good description of their spectral measurements. Consideration of their results and other

TABLE 7 CIE Basis Vectors for Daylights

Wavelengths are in nm. The data are available at 5-nm increments from 300 nm to 830 nm.[11]

Wavelength	Vector 1	Vector 2	Vector 3
400	9.48E+01	4.34E+01	-1.10E+00
410	1.05E+02	4.63E+01	-5.00E-01
420	1.06E+02	4.39E+01	-7.00E-01
430	9.68E+01	3.71E+01	-1.20E+00
440	1.14E+02	3.67E+01	-2.60E+00
450	1.26E+02	3.59E+01	-2.90E+00
460	1.26E+02	3.26E+01	-2.80E+00
470	1.21E+02	2.79E+01	-2.60E+00
480	1.21E+02	2.43E+01	-2.60E+00
490	1.14E+02	2.01E+01	-1.80E+00
500	1.13E+02	1.62E+01	-1.50E-01
510	1.11E+02	1.32E+01	-1.30E+00
520	1.07E+02	8.60E+00	-1.20E+00
530	1.09E+02	6.10E+00	-1.00E+00
540	1.05E+02	4.20E+00	-5.00E-01
550	1.04E+02	1.90E+00	-3.00E-01
560	1.00E+02	0.00E+00	0.00E+00
570	9.60E+01	-1.60E+00	2.00E-01
580	9.51E+01	-3.50E+00	5.00E-01
590	8.91E+01	-3.50E+00	2.10E+00
600	9.05E+01	-5.80E+00	3.20E+00
610	9.03E+01	-7.20E+00	4.10E+00
620	8.84E+01	-8.60E+00	4.70E+00
630	8.40E+01	-9.50E+00	5.10E+00
640	8.51E+01	-1.09E+01	6.70E+00
650	8.19E+01	-1.07E+01	7.30E+00
660	8.26E+01	-1.20E+01	8.60E+00
670	8.49E+01	-1.40E+01	9.80E+00
680	8.13E+01	-1.36E+01	1.02E+01
690	7.19E+01	-1.20E+01	8.30E+00
700	7.43E+01	-1.33E+01	9.60E+00

daylight measurements led the CIE to standardize a three-dimensional linear model for natural daylights. The basis vectors for this model are provided in Table 7. Figure 11 depicts a daylight spectral power distribution (measured at the author's laboratory in Santa Barbara, Calif.) and its approximation using the first two basis vectors of the CIE linear model for daylight.

Cohen[49] analyzed the reflectance spectra of a large set of Munsell papers[50,51] and concluded that a four-dimensional linear model provided a good approximation to the entire data set. The basis vectors for Cohen's linear model are provided in Table 8. Maloney[52] reanalyzed these data, plus a set of natural spectra measured by Krinov[53] and confirmed Cohen's conclusion. More recently, reflectance measurements of the spectra of additional Munsell papers and of natural objects[54,55] have been described by small-dimensional linear models. Figure 12 shows a measured surface reflectance spectrum (red cloth, measured in the author's laboratory) and its approximation using Cohen's four-dimensional linear model.

It is not yet clear why natural illuminant and surface spectra are well approximated by small-dimensional linear models or how general this conclusion is. Maloney[52] provides some speculations. Nonetheless, the assumption that natural spectra do lie within small-dimensional linear models seems reasonable in light of the currently available evidence. This assumption makes possible a number of interesting practical calculations, as we illustrate in some of the following sections.

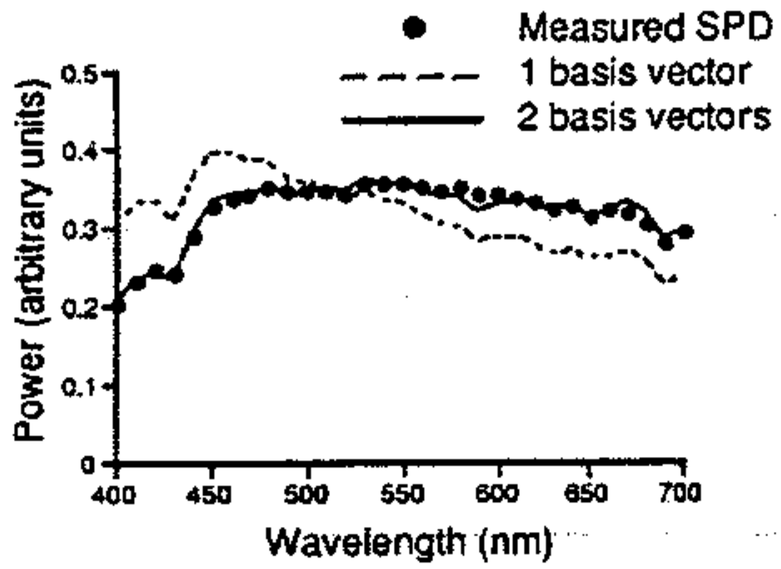

FIGURE 11 The figure shows a daylight spectral power distribution and its approximation using the CIE linear model for daylight. For this particular illuminant, only two basis functions were required to provide a very good fit.

TABLE 8 Cohen's Basis Vectors for Reflectance Functions

Wavelengths are in nm. The data are available at 10-nm increments from 380 nm to 770 nm.[49]

Wavelength	Vector 1	Vector 2	Vector 3	Vector 4
400	3.50E+00	-9.14E-01	-7.83E-01	-2.28E-01
410	3.51E+00	-9.54E-01	-7.78E-01	-2.50E-01
420	3.52E+00	-9.94E-01	-7.65E-01	-2.66E-01
430	3.54E+00	-1.05E+00	-7.39E-01	-2.64E-01
440	3.55E+00	-1.12E+00	-7.09E-01	-2.44E-01
450	3.57E+00	-1.20E+00	-6.63E-01	-2.07E-01
460	3.57E+00	-1.28E+00	-5.81E-01	-1.63E-01
470	3.56E+00	-1.35E+00	-4.62E-01	-1.13E-01
480	3.58E+00	-1.38E+00	-2.47E-01	-4.24E-02
490	3.62E+00	-1.38E+00	8.40E-03	4.18E-02
500	3.67E+00	-1.31E+00	2.79E-01	1.20E-01
510	3.80E+00	-1.13E+00	6.10E-01	2.62E-01
520	3.91E+00	-9.28E-01	8.89E-01	4.00E-01
530	3.91E+00	-7.84E-01	1.03E+00	4.07E-01
540	3.91E+00	-6.33E-01	1.10E+00	3.53E-01
550	3.94E+00	-4.66E-01	1.13E+00	2.80E-01
560	3.94E+00	-2.78E-01	1.12E+00	2.15E-01
570	4.04E+00	-7.45E-02	1.13E+00	8.13E-02
580	4.17E+00	1.71E-01	1.07E+00	-8.77E-02
590	4.42E+00	4.20E-01	1.01E+00	-1.94E-01
600	4.55E+00	6.63E-01	7.18E-01	-3.45E-01
610	4.73E+00	8.70E-01	4.75E-01	-3.74E-01
620	4.84E+00	1.00E+00	2.76E-01	-3.80E-01
630	4.90E+00	1.08E+00	1.47E-01	-3.87E-01
640	4.93E+00	1.12E+00	5.61E-02	-3.92E-01
650	4.97E+00	1.15E+00	-1.21E-02	-3.87E-01
660	5.01E+00	1.16E+00	-6.80E-02	-3.67E-01
670	5.08E+00	1.16E+00	-1.20E-01	-3.19E-01
680	5.17E+00	1.16E+00	-1.61E-01	-2.40E-01
690	5.27E+00	1.16E+00	-1.99E-01	-1.29E-01
700	5.39E+00	1.14E+00	-2.40E-01	1.28E-02

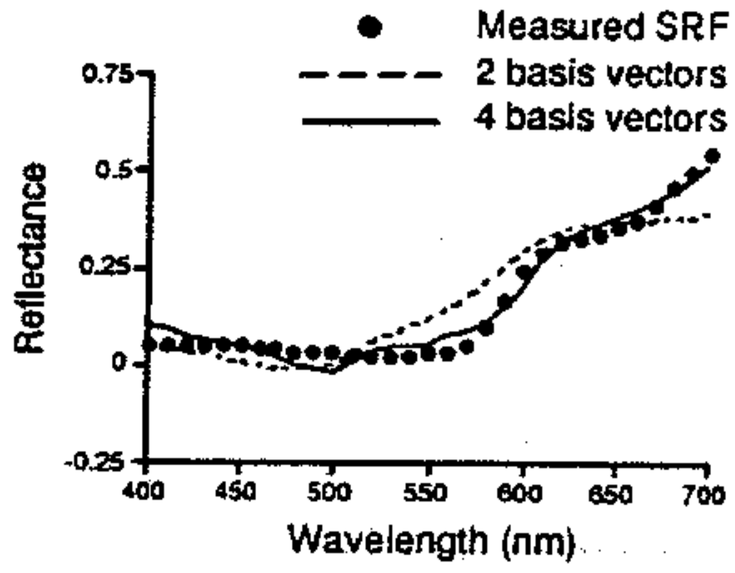

FIGURE 12 The figure shows a measured surface reflectance function and a fit to it using Cohen's 4-dimensional linear model.[49]

Determining a Linear Model from Raw Spectral Data. Given a set of spectral measurements, it is possible, for any integer N_b, to find the N_b dimensional linear model that best approximates the spectral data set (in a least-squares sense). Suppose that the data set consists of N_{meas} spectra, each of which is represented at N_λ sample wavelengths. Let $\mathbf{X}$ be an N_λ by N_{meas} data matrix whose columns represent the individual spectral measurements. The goal of the calculation is to determine an N_λ by N_b matrix $\mathbf{B}$ and an N_b by N_{meas} matrix of coefficients $\mathbf{A}$ such that the linear model approximation $\hat{\mathbf{X}} = \mathbf{BA}$ is the best least-squares approximation to the data matrix $\mathbf{X}$ over all possible choices of $\mathbf{B}$ and $\mathbf{A}$.

The process of finding the matrix $\mathbf{B}$ is called *one-mode components analysis.*[56] It is very closely related to the principle components analysis technique discussed in most multivariate statistics texts.[27,57] One-mode components analysis may be accomplished numerically through the use of the singular value decomposition.[25,29] We define the singular value decomposition in App. A. To see how the singular value decomposition is used to determine an N_b dimensional linear model for $\mathbf{X}$, consider Fig. 13. The top part of the figure depicts the singular value decomposition of an N_λ by N_{meas} matrix $\mathbf{X}$ for the case $N_{\text{meas}} > N_\lambda$, where the two matrices $\mathbf{D}$ and $\mathbf{V}^T$ have been collapsed. This form makes it clear that each column of $\mathbf{X}$ is given by a linear combination of the columns of $\mathbf{U}$. Furthermore, for each column of $\mathbf{X}$, the weights needed to combine the columns of $\mathbf{U}$ are given by the corresponding column of the matrix $\mathbf{DV}^T$. Suppose we choose an N_b

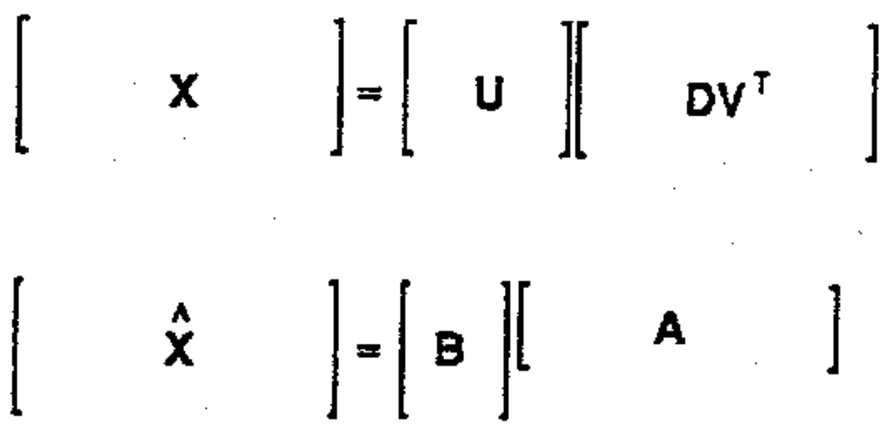

FIGURE 13 The top part of the figure depicts the singular value decomposition (SVD) of an N_λ by N_{meas} matrix $\mathbf{X}$ for the case $N_{meas} > N_\lambda$. In this view we have collapsed the two matrices $\mathbf{D}$ and $\mathbf{V}^T$. To determine an N_b dimensional linear model $\mathbf{B}$ for the data in $\mathbf{X}$ we let $\mathbf{B}$ consist of the first N_b columns of $\mathbf{U}$. As shown in bottom of the figure, the linear model approximation of the data is given by $\hat{\mathbf{X}} = \mathbf{BA}$ where $\mathbf{A}$ consists of the first N_b rows of $\mathbf{DV}^T$.

dimensional linear model **B** for the data in **X** by extracting the first N_b columns of **U**. In this case, it should be clear that we can form an approximation $\hat{\mathbf{X}}$ to the data **X** as shown in the bottom part of the figure. Because the columns of **U** are orthogonal, the matrix **A** consists of the first N_b rows of $\mathbf{DV}^T$. The accuracy of the approximation $\hat{\mathbf{X}}$ depends on how important the columns of **U** excluded from **B** were to the original expression for **X**. It can be shown that choosing **B** as above produces a linear model that minimizes the approximation error of the linear model, for any choice of N_b.[56] Thus, computing the singular value decomposition of **X** allows us to find a good linear model of any desired dimension for $N_b < N_\lambda$. Computing linear models from data is quite feasible on modern personal computers.

Although the above procedure produces the linear model that provides the best least-squares fit to a data set, there are a number of additional considerations that should go into choosing a linear model. First, we note that the choice of linear model is not unique. Any nonsingular linear combination of the columns of **B** will produce a linear model that provides an equally good account of the data. Second, the least-squares error measure gives more weight to spectra with large amplitudes. In the case of surface spectra, this means that the more reflective surfaces will tend to drive the choice of basis vectors. In the case of illuminants, the more intense illuminants will tend to drive the choice. To avoid this weighting, the measured spectra are sometimes normalized to unit length before performing the singular value decomposition. The normalization equalizes the effect of the relative shape of each spectrum in the data set.[52,58] Third, it is sometimes desired to find a linear model that best describes the variation of a data set around its mean. To do this, the mean of the data set should be subtracted before performing the singular value decomposition. When the mean of the data is subtracted, one-mode components analysis is identical to principle components analysis. Finally, there are circumstances where the linear model will be used not to approximate spectra but rather to approximate some other quantity (e.g., color coordinates) that depend on the spectra. In this case, more general techniques, closely related to those discussed here, may be used.[59]

Approximating a Spectrum with Respect to a Linear Model. Given an N_b dimensional model **B**, it is straightforward to find the representation of any spectrum with respect to the linear model. Let **X** be a matrix representing the spectra of functions to be approximated. These spectra do not need to be members of the data set that was used to determine the linear model. To find the matrix of coefficients **A** such that $\hat{\mathbf{X}} = \mathbf{BA}$ best approximates **X** we use simple linear regression. Regression routines to solve this problem are provided as part of any standard matrix algebra software package.

Digital Image Representations. If, in a given application, illuminants and surfaces may be represented with respect to small-dimensional linear models, then it becomes feasible to use point-by-point representations of these quantities in digital image processing. In typical color image processing, the image data at each point are represented by three numbers at each location. These numbers are generally tristimulus coordinates in some color space. In calibrated systems, side information about the color-matching functions or primary spectral power distributions that define the color space is available to interpret the tristimulus coordinates. It is straightforward to generalize this notion of color images by allowing the images to contain N_b numbers at each point and allowing these numbers to represent quantities other than tristimulus coordinates.[8] For example, in representing the image produced by a printer, it might be advantageous to represent the surface reflectance at each location.[60] If the gamut of printed reflectances can be represented within a small-dimensional linear model, then representing the surface reflectance functions with respect to this model would not require much more storage than a traditional color image.[8] The basis functions for the linear model need be represented only once, not at each location. But by representing reflectances rather than tristimulus values, it becomes possible to compute what the tristimulus coordinates reflected from the printed image

would be under any illumination. We illustrate the calculation below. Because of the problem of metamerism, this calculation is not possible if only the tristimulus coordinates are represented in the digital image.

Simulation of Illuminated Surfaces. Consider the problem of producing a signal on a monitor that has the same tristimulus coordinates as a surface under a variety of different illuminants. The solution to this problem is straightforward and is useful in a number of applications. These include rendering digitally archived paintings,[8,61] generating stimuli for use in psychophysics,[62] and producing photorealistic computer-generated imagery.[7] We show the calculation for the data at a single image location. Let **a** be a representation of the surface reflectance with respect to an N_b dimensional linear model **B**. Let **E** represent the illuminant spectral power distribution in diagonal matrix form. Let **T** represent the color-matching functions for a human observer, and **P** represent the primary phosphor spectral power distributions for the monitor on which the surface will be rendered. We wish to determine tristimulus coordinates **t** with respect to the monitor primaries so that the light emitted from the monitor will appear identical to the light reflected from the simulated surface under the simulated illuminant. From Eqs. (2) (cast as $\mathbf{s} = \mathbf{Ba}$), (24), (18), and (19) we can write directly the desired rendering equation

$$\mathbf{t} = ((\mathbf{TP})^{-1}(\mathbf{TE})\mathbf{B})\mathbf{a} \tag{25}$$

The rendering matrix $((\mathbf{TP})^{-1}(\mathbf{TE})\mathbf{B})$ has dimensions 3 by N_b and maps the surface weights directly to monitor tristimulus coordinates. It is quite general, in that we may use it for any calibrated monitor and any choice of linear models. It does not depend on the particular surface being rendered and may be computed once for an entire image. Because the rendering matrix is of small dimension, rendering of this sort is feasible, even for very large images. As discussed earlier in the chapter, it may be possible to determine the matrix $\mathbf{M}_{T,P} = (\mathbf{TP})^{-1}$ directly. A similar shortcut is possible for the matrix $(\mathbf{TE})\mathbf{B}$. Each column of this matrix is the tristimulus coordinates of one linear model basis vector under the illuminant specified by the matrix **E**.

Color Coordinates of Surfaces. Our discussion thus far has emphasized describing the color coordinates of lights. In many applications of colorimetry, it is desirable to describe the color properties of reflective objects. One efficient way to do this, as described above, is to use linear models to describe the full surface reflectance functions. Another possibility is to specify the color coordinates of the light reflected from the surface under standard illumination. This method allows the assignment of tristimulus values to surfaces in an orderly fashion. The CIE has standardized several illuminant spectral power distributions that may be used for this purpose (see following section). Using the procedures defined above, one can begin with the spectral power distribution of the illuminant and the surface reflectance function and from there calculate the desired color-matching coordinates. The relative size of the tristimulus values assigned to a surface depends on its spectral reflectance function and on the illuminant chosen for specification. To factor the intensity of the illuminant out of the surface representation, the CIE specified a normalization of the color coordinates for use with 1931 *XYZ* tristimulus coordinates. This normalization consists of multiplying the computed tristimulus coordinates by the quantity $100/Y_0$, where Y_0 is the Y tristimulus coordinate for the illuminant.

The tristimulus coordinates of a surface provide enough information to match the surface when it is viewed under the illuminant used to compute those coordinates. It is important to bear in mind that two surfaces that have the same tristimulus coordinates under one illuminant do not necessarily share the same tristimulus coordinates under another illuminant. A more complete description can be generated using the linear model approach described above.

Standard Sources of Illumination. The CIE has standardized a number of illuminant

TABLE 9 CIE Illuminants A and D65

Wavelengths are in nm. The data are available at 1-nm increments from 300 nm to 830 nm.[13]

Wavelength	A	D_{65}
400	1.47E+01	8.28E+01
410	1.77E+01	9.15E+01
420	2.10E+01	9.34E+01
430	2.47E+01	8.67E+01
440	2.87E+01	1.05E+02
450	3.31E+01	1.17E+02
460	3.78E+01	1.18E+02
470	4.29E+01	1.15E+02
480	4.82E+01	1.16E+02
490	5.39E+01	1.09E+02
500	5.99E+01	1.09E+02
510	6.61E+01	1.08E+02
520	7.25E+01	1.05E+02
530	7.91E+01	1.08E+02
540	8.60E+01	1.04E+02
550	9.29E+01	1.04E+02
560	1.00E+02	1.00E+02
570	1.07E+02	9.63E+01
580	1.14E+02	9.58E+01
590	1.22E+02	8.87E+01
600	1.29E+02	9.00E+01
610	1.36E+02	8.96E+01
620	1.44E+02	8.77E+01
630	1.51E+02	8.33E+01
640	1.58E+02	8.37E+01
650	1.65E+02	8.00E+01
660	1.72E+02	8.02E+01
670	1.79E+02	8.23E+01
680	1.85E+02	7.83E+01
690	1.92E+02	6.97E+01
700	1.98E+02	7.16E+01

spectral power distributions.[13] These were designed to be typical of various common viewing conditions and are useful as specific choices of illumination when the illuminant cannot be measured directly. CIE illuminant A is designed to be representative of tungsten-filament illumination. CIE illuminant D65 is designed to be representative of average daylight. The relative spectral power distributions of these two illuminants are provided in Table 9. Other CIE standard daylight illuminants may be computed using the basis vectors in Table 7 and formulae specified by the CIE.[11] Spectra representative of fluorescent lamps and other artificial sources are also available.[9,11]

Metamerism

Recovering Spectral Power Distributions from Tristimulus Coordinates. It is not possible in general to recover a spectral power distribution from its tristimulus coordinates. If some a priori information about the spectral power distribution of the color signal is available, however, then recovery may be possible. Such recovery is of most interest in applications where direct spectral measurements are not possible and where knowing the full spectrum is important. For example, the effect of lens chromatic aberrations on cone quantal absorption rates depends on the full spectral power distribution.[45,63]

Suppose the spectral power distribution of interest is known to lie within a three-dimensional linear model. We may write $\mathbf{b} = \mathbf{Ba}$, where the basis matrix $\mathbf{B}$ has dimensions N_λ by 3. Let $\mathbf{t}$ be the tristimulus coordinates of the light with respect to a set of color-matching functions $\mathbf{T}$. Following the development earlier in the chapter, we can conclude that $\mathbf{a} = (\mathbf{TB})^{-1}\mathbf{t}$, which imples

$$\mathbf{b} = \mathbf{B}(\mathbf{TB})^{-1}\mathbf{t} \tag{26}$$

When we do not have a prior constraint that the signal belongs to a three-dimensional linear model, we may still be able to place some linear model constraint, of dimension higher than 3, on the spectral power distribution. For example, when we know that the signal was produced by the reflection of daylight from a natural object, it is reasonable to assume that the color signal lies within a linear model of dimension that may be as low as 9.[64] In this case, we can still write $\mathbf{b} = \mathbf{Ba}$, but we cannot apply Eq. (26) directly because the matrix $(\mathbf{TB})$ will be singular. To deal with this problem, we can choose a reduced linear model $\hat{\mathbf{B}}$ with only three dimensions. We then proceed as outlined above, but substitute the reduced model for the true model. This will lead to an estimate $\hat{\mathbf{b}}$ for the actual spectral power distribution $\mathbf{b}$. If the reduced linear model $\hat{\mathbf{B}}$ provides a reasonable approximation to $\mathbf{b}$, the estimation error may be quite small. The estimate $\hat{\mathbf{b}}$ will have the property that it is a metamer of $\mathbf{b}$. The techniques described above for finding linear model approximations may be used to choose an appropriate reduced model.

Finding Metamers of a Light. It is often of interest to find metamers of a light. We discuss two approaches here. Wyszecki and Stiles[9] treat the problem in considerable detail.

Using a Linear Model. If we choose any three-dimensional linear model $\hat{\mathbf{B}}$ we can combine Eq. (26) with the fact the fact that $\mathbf{t} = \mathbf{Tb}$ [Eq. (9)] to compute a pair of metameric spectral power distributions $\mathbf{b}$ and $\hat{\mathbf{b}}$:

$$\hat{\mathbf{b}} = \hat{\mathbf{B}}(\mathbf{T}\hat{\mathbf{B}})^{-1}\mathbf{Tb} \tag{27}$$

Each choice of $\hat{\mathbf{B}}$ will lead to a different metamer $\hat{\mathbf{b}}$. Figure 14 shows a number of metameric spectral power distributions generated in this fashion.

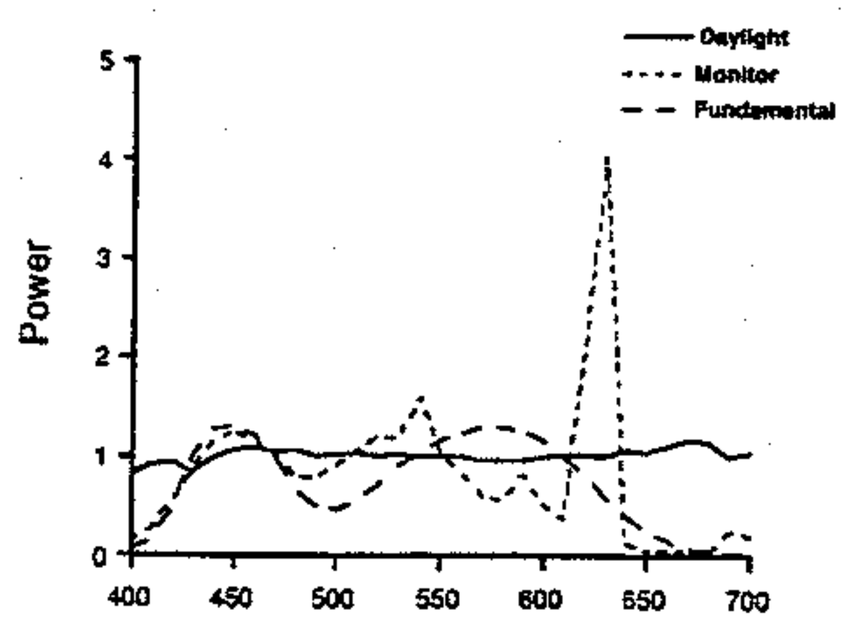

FIGURE 14 The figure shows three metameric color signals with respect to the CIE 1931 standard observer. The three metamers were computed using Eq. (27). The initial spectral power distribution $\mathbf{b}$ (not shown) was an equal energy spectrum. Three separate linear models were used: one that describes natural daylights, one typical of monitor phosphor spectral power distributions, and one that provides Cohen's "fundamental metamer."

Metameric Blacks. Another approach to generating metamers is to note that there will be some spectral power distributions $\mathbf{b}_0$ that have the property $\mathbf{Tb}_0 = \mathbf{0}$. Wyszecki referred to such distributions as metameric blacks, since they have the same tristimulus coordinates as no light at all.[9,65] Grassmann's laws imply that adding a metameric black $\mathbf{b}_0$ to any light $\mathbf{b}$ yields a metamer of $\mathbf{b}$. Given a linear model $\mathbf{B}$ with dimension greater than 3 it is possible to find a second linear model $\mathbf{B}_0$ such that (1) all lights that lie in $\mathbf{B}_0$ also lie in $\mathbf{B}$ and (2) all lights in $\mathbf{B}_0$ are metameric blacks. We determine $\mathbf{B}_0$ by finding a linear model for the null space of the matrix $\mathbf{TB}$. The null space of a matrix consists of all vectors that are mapped to $\mathbf{0}$ by the matrix. Finding a basis for the null space of a matrix is a standard operation in numerical matrix algebra. If we have a set of basis vectors $\mathbf{N}_0$ for the null space of $\mathbf{TB}$, we can form $\mathbf{B}_0 = \mathbf{BN}_0$. This technique provides a way to generate a large list of metamers for any given light $\mathbf{b}$. We choose a set of weights $\mathbf{a}$ at random and construct $\mathbf{b}_0 = \mathbf{B}_0\mathbf{a}$. We then add $\mathbf{b}_0$ to $\mathbf{b}$ to form a metamer. To generate more metamers, we simply repeat with new choices of weight vector $\mathbf{a}$.

Surface and Illuminant Metamerism. The formal similarity between Eq. (9) (which gives the relation between spectral power distributions and tristimulus coordinates) and Eq. (24) (which gives the relation between surface reflectance functions and tristimulus coordinates when the illuminant is known) makes it clear that our discussion of metamerism can be applied to surface reflectance spectra. Two physically different surfaces will appear identical if the tristimulus coordinates of the light reflected from them is identical. This fact can be used to good purpose in some color reproduction applications. Suppose that we have a sample surface or textile whose color we wish to reproduce. It may be that we are not able to reproduce the sample's surface reflectance function exactly because of various limitations in the available color reproduction technology. If we know the illuminant under which the reproduction will be viewed, we may be able to determine a reproducible reflectance function that is metameric to that of the desired sample. This will give us a sample whose color appearance is as desired. Applications of this sort make heavy use of the methods described above to determine metamers.

But what if the illuminant is not known or if it is known to vary? In this case there is an additional richness to the topic of determining metamers. We can pose the problem of finding surface reflectance functions that will be metameric to a desired reflectance under multiple specified illuminants or under all of the illuminants within some linear model. The general methods developed here have been extended to analyze this case.[64,66] Similar issues arise in lighting design, where we desire to produce an artificial light whose color-rendering properties match those of a specified light (such as natural daylight). When wavelength-by-wavelength matching of the spectra is not feasible, it may still be possible to find a spectrum so that the light reflected from surfaces within a linear model is identical for the two light sources. Because of the symmetric role of illuminants and surfaces in reflection, this problem may be treated by the same methods used for surface reproduction.

Cohen's "Matrix R." Cohen and Kappauf[21,67,68] have proposed that a useful way to associate a spectral power distribution with a set of tristimulus coordinates is to choose as basis functions the color-matching functions themselves. That is, we choose the matrix $\hat{\mathbf{B}} = \mathbf{T}^T$. When $\hat{\mathbf{B}}$ is chosen in this way, the estimated color signal from Eq. (26) is given by $\hat{\mathbf{b}} = \mathbf{T}^T(\mathbf{TT}^T)^{-1}\mathbf{t}$. From Eq. (27) we also have $\hat{\mathbf{b}} = \mathbf{T}^T(\mathbf{TT}^T)^{-1}\mathbf{Tb}$. Cohen and Kappauf refer to this $\hat{\mathbf{b}}$ as the "fundamental metamer" of $\mathbf{b}$. The matrix $\mathbf{T}^T(\mathbf{TT}^T)^{-1}\mathbf{T}$ is often referred to as "matrix R." It is easy to show that "matrix R" is invariant when a linear transformation is applied to the color-matching functions. There is no reason to believe the "fundamental metamer" will be a good estimate (in a least-squares sense) of the original spectral power distribution (see Fig. 14).

Color Cameras and Other Visual Systems

We have treated colorimetry from the point of view of specifying the spectral information available to a human observer. We have developed our treatment, however, in such a way that it may be applied to handle other visual systems. Suppose that we wish to define color coordinates with respect to some arbitrary visual system with N_{device} photosensors. This visual system might be an artificial system based on a color camera or scanner, a nonhuman biological visual system, or the visual system of a color-anomalous human observer. We assume that the sensitivities of the visual system's photosensors are known up to a linear transformation. Let $\mathbf{T}_{device}$ be an $\mathbf{N}_{device}$ by $\mathbf{N}_\lambda$ matrix whose entries are the sensitivities of each sensor at each sample wavelength. We can compute the responses of these sensors to any light $\mathbf{b}$. Let $\mathbf{t}_{device}$ be a vector containing the responses of each sensor type to the light. Then we have $\mathbf{t}_{device} = \mathbf{T}_{device}\mathbf{b}$. We may use $\mathbf{t}_{device}$ as the device color coordinates of $\mathbf{b}$.

Transformation Between Color Coordinates of Different Visual Systems. Suppose that we have two different visual systems and we wish to transform between the color coordinates of each. A typical example might be trying to compute the CIE 1931 *XYZ* tristimulus coordinates of a light from the responses of a color camera. Let N_s be the number of source sensors, with sensitivities specified by $\mathbf{T}_s$. Similarly, let N_d be the number of destination sensors with sensitivities specified by $\mathbf{T}_d$. For any light $\mathbf{b}$ we know that the source device color coordinates are given by $\mathbf{t}_s = \mathbf{T}_s\mathbf{b}$ and the destination device color coordinates $\mathbf{t}_d = \mathbf{T}_d\mathbf{b}$. We would like to transform between $\mathbf{t}_s$ and $\mathbf{t}_d$ without direct knowledge of $\mathbf{b}$.

If we can find an N_d by N_s matrix $\mathbf{M}$ such that $\mathbf{T}_d = \mathbf{M}\mathbf{T}_s$, then it is easy to show that the matrix $\mathbf{M}$ may be used to compute the destination device color coordinates from the source device color coordinates through $\mathbf{t}_d = \mathbf{M}\mathbf{t}_s$. We have already considered this case (in a less general form). The extension here is that we allow the possibility that the dimensions of the two color coordinate systems differ. When a linear transformation between $\mathbf{T}_s$ and $\mathbf{T}_d$ exists, it can be found by standard regression methods.

Horn demonstrated that when no exact linear transformation between $\mathbf{T}_s$ and $\mathbf{T}_d$ exists, it is not, in general, possible to transform between the two sets of color coordinates.[4] The reason for this is that there will always exist a pair of lights that have the same color coordinates for the source device but different color coordinates for the destination device. The transformation will therefore be incorrect for at least one member of this pair. When no exact linear transformation exists, it is still possible to make an approximate transformation. One approach is to use linear regression to find the best linear transformation $\mathbf{M}$ between the two sets of color-matching functions in a least-squares sense. This transformation is then applied to the source color coordinates as if it were exact.[4] Although this is an approximation, in many cases the results will be acceptable. In the absence of prior information about the spectral power distribution of the original light $\mathbf{b}$, it is a sensible approach.

A second possibility is to use prior constraints on the spectral power distribution of the light to guide the transformation.[20,69] Suppose that we know that the light is constrained to lie within an N_b dimensional linear model $\mathbf{B}$. Then we can find the best linear transformation $\mathbf{M}$ between the two matrices $\mathbf{T}_s\mathbf{B}$ and $\mathbf{T}_d\mathbf{B}$. This transformation may then be used to transform the source color coordinates to the destination color coordinates. It is easy to show that the transformation will be exact if $\mathbf{T}_d\mathbf{B} = \mathbf{M}\mathbf{T}_s\mathbf{B}$. Otherwise, it is a reasonable approximation that takes the linear model constraint into account.

Computational Color Constancy. An interesting application is the problem of estimating surface reflectance functions from color coordinates. This problem is of interest for two reasons. First, it appears that human color vision makes some attempt to perform this estimation, so that our percept of color is more closely associated with object surface properties than with the proximal properties of the light reaching the eye. Second, an

artificial system that could estimate surface properties would have an important cue to aid object recognition. In the case where the illuminant is known, the problem of estimating surface reflectance properties is the same as the problem of estimating the color signal, because the illuminant spectral power distribution can simply be incorporated into the sensor sensitivities. In this case the methods outlined above for estimating color signal spectral properties can be used.

The more interesting case is where both the illuminant and the surface reflectance are unknown. In this case, the problem is more difficult. Considerable insight has been gained by applying linear model constraints to both the surface and illuminant spectral power distributions. In the past decade, a large number of approaches have been developed for recovering surface reflectance functions.[20,64,70–78] Each approach differs (1) in the additional assumptions that are made about the properties of the image and (2) in the sophistication of the model of illuminant surface interaction and scene geometry used. A thorough review of all of these methods is beyond the scope of this chapter. It is instructive, however, to review one of the simpler methods, that of Buchsbaum.[72]

Buchsbaum assumed that in any given scene, the average reflectance function of the surfaces in the scene is known. This is commonly called the "gray world" assumption. He also assumed that the illuminant was diffuse and constant across the scene and that the illuminants and surfaces in the scene are described by linear models with the same dimensionality as the number of sensors. Let $\mathbf{S}_{avg}$ be the spectral power distribution of the known average surface, represented in diagonal matrix form. Then it is possible to write the relation between the space average of the sensor responses and the illuminant as

$$\mathbf{t}_{avg} = \mathbf{T}\mathbf{S}_{avg}\mathbf{B}_e\mathbf{a}_e \tag{28}$$

where $\mathbf{a}_e$ is a vector containing the weights of the illuminant within the linear model representation $\mathbf{B}_e$. Because we assume that the dimension $N_e = N_t$, the matrix $\mathbf{T}\mathbf{S}_{avg}\mathbf{B}_e$ will be square and typically may be inverted. From this we recover the illuminant as $\mathbf{e} = \mathbf{B}_e(\mathbf{T}\mathbf{S}_{avg}\mathbf{B}_e)^{-1}\mathbf{t}_{avg}$. If we let $\mathbf{E}$ represent the recovered illuminant in matrix form, then at each image location we can write

$$\mathbf{t} = \mathbf{T}\mathbf{E}\mathbf{B}_s\mathbf{a}_s \tag{29}$$

where $\mathbf{a}_s$ is a vector containing the weights of the surface within the linear model representation $\mathbf{B}_s$. Proceeding exactly as we did for the illuminant, we may recover the surface reflectance from this equation.

Although Buchsbaum's method depends on rather strong assumptions about the nature of the scene, subsequent algorithms have shown that these assumptions can be weakened. Maloney and Wandell demonstrated that the gray world assumption can, under certain circumstances, be relaxed.[20,73] Several recent reviews emphasize the relation between computational color constancy and the study of human vision.[79–82]

Color Discrimination

Measurement of Small Color Differences. Our treatment so far has not included any discussion of the precision to which observers can judge identity of color appearance. To specify tolerances for color reproduction, it would be helpful to know how different the color coordinates of two lights must be for an observer to reliably distinguish between them. A number of techniques are available for measuring human ability to discriminate between colored lights.

One method, employed in seminal work by MacAdam,[83,84] is to examine the variability in individual color matches. That is, if we have observers set matches to the same test stimulus, we will discover that they do not always set exactly the same values. Rather, there will be some trial-to-trial variability in the settings. MacAdam and others[85,86] used the sample covariance of the individual match tristimulus coordinates as a measure of observers' color discrimination.

A second approach is to use psychophysical methods (see Vol. I, Chap. 39) to measure

observers' thresholds for discriminating between pairs of colored lights. Examples include increment threshold measurements for monochromatic lights[87] and thresholds measured systematically in a three-dimensional color space.[88,89]

Measurements of small color differences are often summarized with isodiscrimination contours. An isodiscrimination contour specifies the color coordinates of lights that are equally discriminable from a common standard light. Figure 15 shows an illustrative isodiscrimination contour. Isodiscrimination contours are often modeled as ellipsoids[88,89] and the figure is drawn to the typical ellipsoidal shape. The well-known MacAdam ellipses[83] are an example of representing discrimination data using the chromaticity coordinates of a cross section of a full three-dimensional isodiscrimination contour (see the legend of Fig. 15).

CIE Uniform Color Spaces. Figure 16 shows chromaticity plots of representative isodiscrimination contours. A striking feature of the plots is that the size and shape of the contours depends on the standard stimulus. For this reason, it is not possible to predict whether two lights will be discriminable solely on the basis of the Euclidean distance between their color coordinates. The heterogeneity of the isodiscrimination contours must also be taken into account.

The CIE provides two sets of formulae that may be used to predict the discriminability of colored lights. Each set of formulae specifies a nonlinear transformation from CIE 1931 *XYZ* color coordinates to a new set of color coordinates. Specifically, the *XYZ* coordinates may be transformed either to CIE 1976 L*u*v* (CIELUV) color coordinates or to CIE

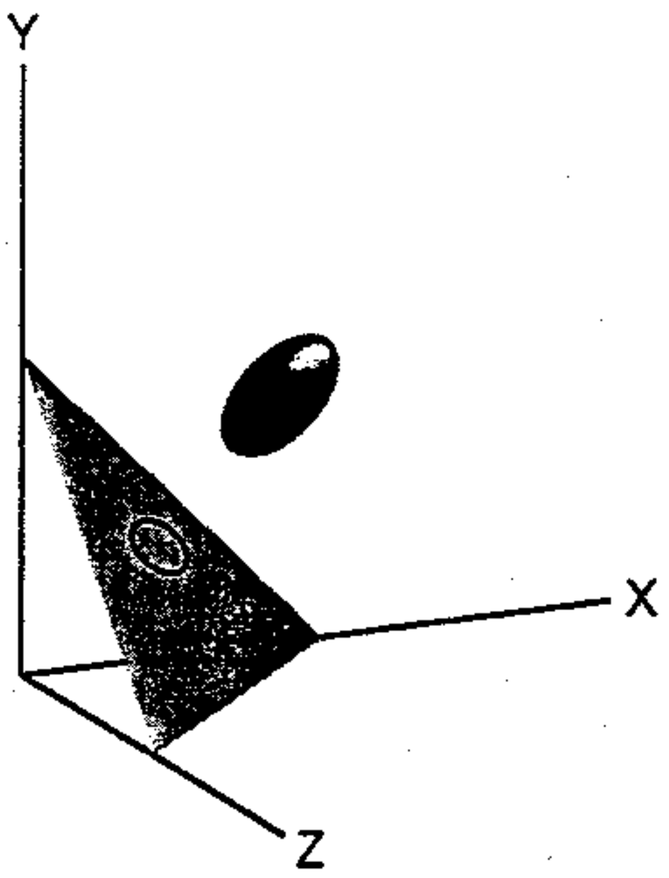

FIGURE 15 Isodiscrimination contour. The plotted ellipsoid shows a hypothetical isodiscrimination contour in the CIE *XYZ* color space. This contour represents color discrimination performance for the standard light whose color coordinates are located at the ellipsoid's center. Isodiscrimination contours such as the one shown are often summarized by a two-dimensional contour plotted on a chromaticity diagram (see Fig. 16). The two-dimensional contour is obtained form a cross section of the full contour, and its shape can depend on which cross section is used. This information is not available directly from the two-dimensional plot. A common criterion for choice of cross section is isoluminance. The ellipsoid shown in the figure is schematic and does not represent actual human performance.

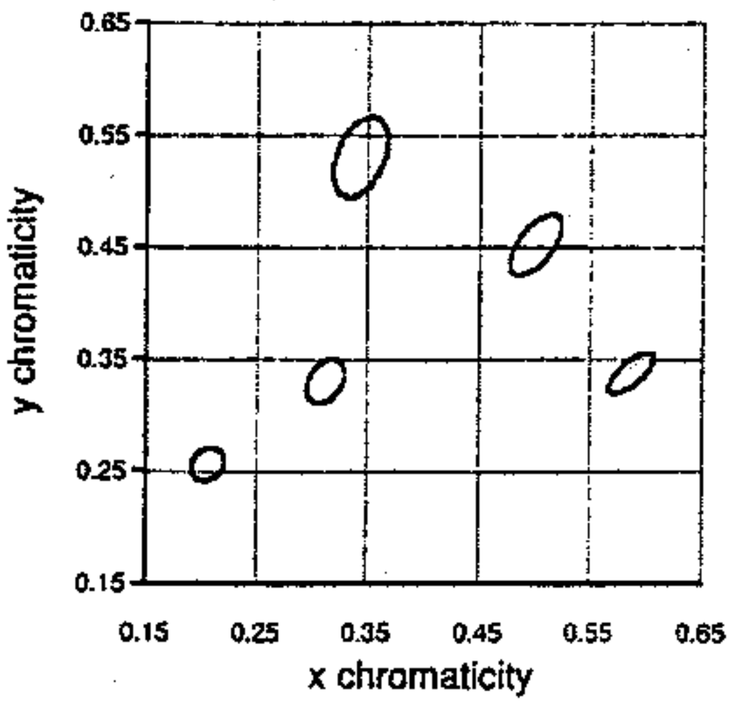

FIGURE 16 Isodiscrimination contours plotted in the chromaticity diagram. These were computed using the CIE L*u*v* uniform color space and provide an approximate representation of human performance. For each standard stimulus, the plotted contour represents the color coordinated of lights that differ from the standard by 15 ΔE^*_{uv} units but that have the same luminance as the standard. The choice of 15 ΔE^*_{uv} units magnifies the contours compared to those that would be obtained in a threshold experiment.

1976 L*a*b* (CIELAB) color coordinates. Both transformations stretch the *XYZ* color space so that the resulting Euclidean distance between color coordinates provides an approximation to the how well lights may be discriminated. Both the L*u*v* and L*a*b* systems are referred to as uniform color spaces. A more complete description of these color spaces is available elsewhere.[9,11] We provide the basic formulae and some discussion below.

Transformation to CIELUV Coordinates. The CIE 1976 L*u*v* color coordinates of a light may be obtained from its CIE *XYZ* coordinates according to the equations

$$L^* = \begin{cases} 116\left(\frac{Y}{Y_n}\right)^{1/3} - 16 & \frac{Y}{Y_n} > 0.008856 \\ 903.3\left(\frac{Y}{Y_n}\right) & \frac{Y}{Y_n} \le 0.008856 \end{cases}$$

$$u^* = 13L^*\left(\frac{4X}{(X+15Y+3Z)} - \frac{4X_n}{(X_n+15Y_n+3Z_n)}\right) \tag{30}$$

$$v^* = 13L^*\left(\frac{9Y}{(X+15Y+3Z)} - \frac{9Y_n}{(X_n+15Y_n+3Z_n)}\right)$$

In this equation, the quantities X_n, Y_n, and Z_n are the tristimulus coordinates of a white point. Little guidance is available as to how to choose an appropriate white point. In the case where the lights being judged are formed when an illuminant reflects from surfaces, the tristimulus coordinates of the illuminant may be used. In the case where the lights being judged are on a computer-controlled color monitor, the sum of the tristimulus coordinates of the three monitor phosphors stimulated at their maximum intensity may be used.

Transformation to CIELAB Coordinates. The CIE 1976 L*a*b* color coordinates of a light may be obtained from its CIE *XYZ* coordinates according to the equations

$$L^* = \begin{cases} 116\left(\frac{Y}{Y_n}\right)^{1/3} - 16 & \frac{Y}{Y_n} > 0.008856 \\ 903.3\left(\frac{Y}{Y_n}\right) & \frac{Y}{Y_n} \le 0.008856 \end{cases}$$

$$a^* = 500\left[f\left(\frac{X}{X_n}\right) - f\left(\frac{Y}{Y_n}\right)\right] \tag{31}$$

$$b^* = 200\left[f\left(\frac{Y}{Y_n}\right) - f\left(\frac{Z}{Z_n}\right)\right]$$

where the function $f(s)$ is defined as

$$f(s) = \begin{cases} (s)^{1/3} & s > 0.008856 \\ 7.787(s) + \frac{16}{116} & s \le 0.008856 \end{cases} \tag{32}$$

As with the CIELUV transformation, the quantities X_n, Y_n, and Z_n are the tristimulus coordinates of a white point.

Distance in CIELUV and CIELAB Spaces. In both the CIELUV and CIELAB color spaces, the Euclidean distance between the coordinates of two lights provides a rough guide to their discriminability. The symbols ΔE^*_{uv} and ΔE^*_{ab} are used to denote distance in the two uniform color spaces and are defined as

$$\Delta E^*_{uv} = \sqrt{(\Delta L^*)^2 + (\Delta u^*)^2 + (\Delta v^*)^2}$$
$$\Delta E^*_{ab} = \sqrt{(\Delta L^*)^2 + (\Delta a^*)^2 + (\Delta b^*)^2} \tag{33}$$

where the various Δ quantities on the right represent the differences between the corresponding coordinates of the two lights. Roughly speaking, a ΔE^*_{uv} or ΔE^*_{ab} value of 1 corresponds to a color difference that can just be reliably discerned by a human observer under optimal viewing conditions. A ΔE^*_{uv} or ΔE^*_{ab} value of 3 is sometimes used as an acceptable tolerance in industrial color reproduction applications.

Limits of the CIE Uniform Color Spaces. The two CIE color difference measures ΔE^*_{uv} and ΔE^*_{ab} provide only an approximate guide to the discriminability between two lights. There are a number of reasons why this is so. One major reason is that the formulae were designed not only to predict discrimination data but also certain suprathreshold judgments of color appearance.[90] A second important reason is that color discrimination thresholds depend heavily on factors other than the tristimulus coordinates. These factors include the adapted state of the observer,[87] the spatial and temporal structure of the stimulus,[91–93] and the task demands placed on the observer.[94–97] Therefore, the complete specification of a uniform color space must incorporate these factors. At present, a model of visual performance that would allow such incorporation is not available. The transformations to CIELUV and CIELAB spaces do, however, depend on the choice of white point. This dependence is designed to provide some compensation for the adapted state of the observer.

Limits of the Color-Matching Experiment

Specifying a stimulus using tristimulus coordinates depends on having an accurate set of color-matching functions. The various standard color spaces discussed earlier under "Common Color Coordinate Systems" are designed to be representative of an "average" or standard observer under typical viewing conditions. A number of factors limit the precision to which a standard color space can predict individual color matches. We describe some of these factors below. Wyszecki and Stiles[9] provide a more detailed treatment.

For most applications, standard calculations are sufficiently precise. When high precision is required, it is necessary to tailor a set of color-matching functions to the individual and observing conditions of interest. Once such a set of color-matching functions is available, the techniques described in this chapter may be used to compute corresponding color coordinates.

Individual Differences Among Color-Normal Observers. Standard sets of color-matching functions are summaries of color-matching results for a number of color-normal observers. There is small but systematic variability between the matches set by individual observers, and this variability limits the precision to which standard color-matching functions may be taken as representative of any given color-normal observer. A number of factors underlie the variability in color matching. Stiles and Burch carefully measured color-matching functions for 49 observers using 10° fields.[98,99] Webster and MacLeod analyzed individual variation in these color-matching functions.[100] They identified six primary factors that drive the variation in individual color matches. These are macular pigment density, lens pigment density, amount of rod intrusion into the matches (discussed shortly), and variability in the absorption spectra of the L, M, and S cone photopigments.

Lens pigment density is known to increase over the life span of an individual, resulting in systematic differences in color-matching functions between populations of different ages.[101] The nature of the mechanisms underlying the variability in cone photopigment absorption spectra is a matter of considerable current interest.

Color-Deficient and Color-Anomalous Observers. Some individuals require only two (or in rare cases only one) primaries in the color-matching experiment. These individuals are referred to as color-deficient or color-blind observers. Most forms of color deficiency can be understood by assuming that the individual lacks one (or more) of the normal three types of cone photopigment. For these individuals, use of standard color coordinates will produce acceptable results, since a match for all three cone types will also be a match for any subset of these types. In very rare cases, an individual has no cones at all and the vision of such an individual is mediated entirely by rods. A second class of individuals are trichomatic but set color matches substantially different from color-normal observers. These individuals are referred to as color anomalous. The leading hypothesis about the cause of color anomaly is that the individuals have photopigments with spectral sensitivities substantially different from individuals with normal color vision.[102] Our development of colorimetry can be used to tailor color specification for color-anomalous individuals if their color-matching functions are known. Estimates of the cone sensitivities of color-anomalous observers are available.[37] Simple standard tests exist for identifying color-blind and color-anomalous individuals. These include the Ishihara pseudo-isochromatic plates[103] and the Farnsworth 100 hue test.[104]

Cone Polymorphism. Recent genetic and behavioral evidence suggests that there are multiple types of human L and possibly human M cone photopigments.[105–107] This possibility is referred to as cone polymorphism. Moreover, Neitz, Neitz, and Jacobs[108] argue that some individuals possess more than three types of cone photopigments. This claim challenges the conventional explanation for trichromacy and is controversial. The interested reader is referred directly to the current literature. Because the purported difference in spectral sensitivity of different subclasses of human L and M cone photopigments is quite small, the possibility of cone polymorphism is not of practical importance for most applications. The most notable exception is in certain psychophysical experiments where precise knowledge of the relative excitation of different cone classes is crucial.

Retinal Inhomogeneity. Most standard colorimetric systems are based on color-matching experiments where the bipartite field was small and viewed foveally. The distribution of photoreceptors and of inert visual pigment is not homogeneous across the retina. Thus color-matching functions that are accurate for the fovea do not necessarily describe color matching in the extra fovea. The CIE 1964 10° *XYZ* color-matching functions are designed for situations where the colors being judged subtend a large visual angle. These functions are provided in Table 3.

Rod Intrusion. Both outside the fovea and at low light levels, rods can play a role in color matching. Under conditions where rods play a role, there is a shift in the color-matching functions due to the contribution of rod signals. Wyszecki and Stiles[9] discuss approximate methods for correcting standard sets of color-matching functions when rods intrude into color vision.

Chromatic Aberrations. By some standards, even the small (roughly 2°) fields used as the basis of most color coordinate systems are rather coarse. The optics of the eye contain chromatic aberrations which cause different wavelengths of light to be focused with different accuracy. These aberrations can cause a shift in the color-matching functions if the stimuli being matched have a fine spatial structure. Two stimuli which are metameric at low spatial frequencies may no longer be so at high spatial frequencies. Such effects can be quite large.[45,63] It is possible to correct color coordinates for chromatic aberration if enough side information is available. Such correction is rare in practice but can be important for stimuli with a fine spatial structure. Some guidance is available from the

literature.[63,109] Another strategy available in the laboratory is to correct the stimulus for the chromatic aberration of the eye.[110]

Pigment Self-Screening. One of the factors that determines the cone sensitivities is that the photopigment itself acts as an inert filter. In any individual cone. the spectral properties of this filter depend on the fraction of photopigment that has recently been isomerized by light quanta. As the overall intensity of the stimulus changes. this fraction changes, which changes the cones' sensitivity functions. Although such shifts may generally be neglected, they can become quite important under circumstances where very intense adapting fields are employed.[18]

Calculating the Effect of Errors in Color-Matching Functions. Given that there is some variation between different standard estimates of color-matching functions. between the color-matching functions of different individuals, and between the color-matching functions that mediate performance for different viewing conditions, it is of interest to determine whether the magnitude of this variation is of practical importance. There is probably no general method for making this determination, but here we outline one approach.

Consider the case of rendering a set of illuminated surfaces on a color monitor. If we know the spectral power distribution of the monitor's phosphors it is possible to compute the appropriate weights on the monitor phosphors to produce a light metameric to each illuminated surface. The computed weights will depend on the choice of color-matching functions. Once we know the weights, however, we can find the CIE 1976 L*u*v* (or CIE 1976 L*a*b*) coordinates of the emitted light. This suggests the following method to estimate the effect of differences in color-matching functions. First, we compute the CIE 1976 L*u*v* coordinates of surfaces rendered using the first set of color-matching functions. Then we compute the corresponding coordinates when the surfaces are rendered using the second set of color-matching functions. Finally, we compute the ΔE^*_{uv} difference between corresponding sets of coordinates. If the ΔE^*_{uv} are large, then the differences between the color-matching functions are important for the rendering application.

We have performed this calculation for a set of 462 measured surfaces[50,51] rendered under CIE Illuminant D_{65}. The two sets of color-matching functions used were the 1931 CIE *XYZ* color-matching functions and the Judd-Vos modified *XYZ* color-matching functions. The monitor phosphor spectral power distributions were measured by the author.[111] The results are shown in Fig. 17. The plot shows a histogram of the ΔE^*_{uv}

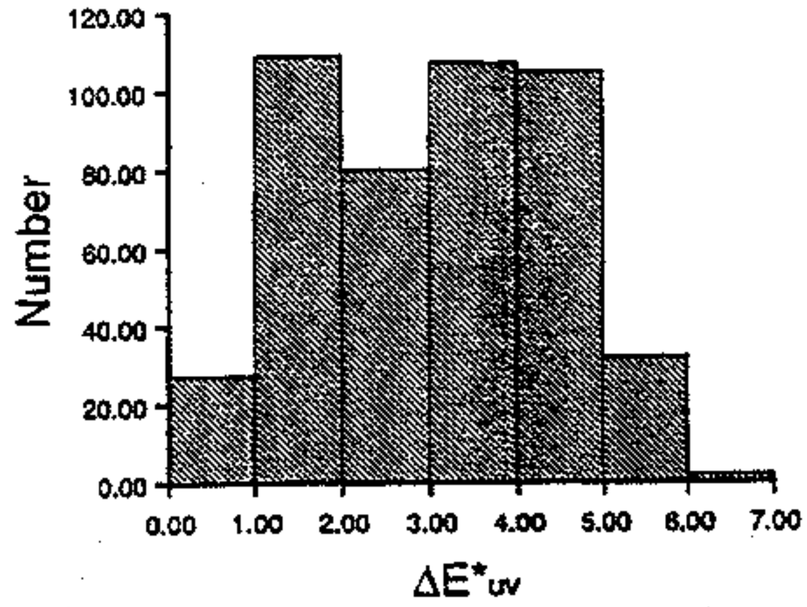

FIGURE 17 Effect of changes in color matching functions. The plot shows a histogram of the ΔE^*_{uv} differences between two sets of lights, each of which is a monitor rendering of the same set of illuminated surfaces. The two renderings were computed using different sets of color matching functions.

differences. The median difference is 3.1 ΔE^*_{uv} units. This difference is quite close to discrimination threshold and, for most applications, the differences between the two sets of color-matching functions will probably not be of great consequence.

Color Appearance

Context Effects. A naive theory of color appearance would predict that once we know the tristimulus coordinates of a light, it would be straightforward to predict its color appearance. Unfortunately, such a theory does not work. The color appearance of a light depends not only on its spectral properties but also on the context in which it is seen. By context, we refer to other attributes of the light itself, such as its size, shape, and temporal profile, to the properties of the lights that surround the light of interest, and to the adapted state of the observer. The adapted state of the observer is presumably determined by the stimuli viewed in the recent past. Some are even willing to argue that more intangible factors, such as the observer's expectations, also influence color appearance.

Wyszecki[3] and Evans[2] each review color context effects. Among the context phenomenon given the most study are color contrast, where the properties of nearby image regions affects the color appearance of a test. and observer adaptation, where the images that an observer has viewed influences his or her judgment of a subsequently viewed test. Empirically, adaptation and contrast are somewhat difficult to separate. Indeed, the two are difficult to separate conceptually, since adaptation can be viewed as a delayed contrast effect and contrast can be considered as rapidly acting adaptation. Although phenomena with names such as simultaneous color contrast, successive color contrast, color assimilation, and observer adaptation are often discussed and considered separately, there is some potential for clarity by considering the possibility that they are all manifestations of similar processes.

The fact that the color-matching experiment, from which color coordinates are derived, does not predict color appearance may seem surprising. After all, the fundamental judgment of the color-matching experiment concerns the identity of color appearance. Recall, however, that the test and matching regions are always juxtaposed, so that they are seen in identical context. As long as the context in which two lights are seen is identical, color coordinates do correctly predict color appearance.

Context effects make color reproduction difficult. If we wish to reproduce an image so that it appears identical to another image when the two will be viewed in different contexts, we must take the effect of context into account. This sort of consideration is of particular importance in color photography and television, since in these cases the reproduced image has a vastly different size from the image it depicts.[2,5,112,113] Two points are worth noting. First, if we can arrange two images so that they have identical tristimulus coordinates at each image location, then we can be assured that each presents the same context to the viewer and that the two images will match. This principle is the basis of many applied color reproduction systems and is appropriate when the source and destination images have roughly equal sizes and will be viewed under similar conditions.[6] Second, tristimulus coordinates are still of great importance when the context does vary. This is because the persistence law of color matching tells us that the context itself may be specified through the use of tristimulus coordinates. We do not need to specify the context in terms of the full spectral power distributions of the lights. Thus the basic color-matching experiment provides the foundation on which we can build theories of color appearance. A theory that successfully links tristimulus coordinates with some sort of appearance coordinates, as a function of the viewing context, would provide the quantitative framework necessary to take context into account in colorimetric applications. The search for such a theory has been at the center of color science since at least the time of von Kries.[1] Wyszecki[3] and Wyszecki and Stiles[9] review of the measurement of the effect of context on color appearance.

Color Constancy. A hypothesis about the origin of color context effects is that the visual system attempts to discount the properties of the illuminant so that color appearance is correlated with object surface reflectance. To the extent that the visual system does this, it is referred to as color constant. It has long been known that the human visual system exhibits partial color constancy.[2,114,115] The phenomenon of color constancy suggests that the spectral power distribution of the illuminant is an important parameter of the viewing context for predicting color appearance.[116] Land's retinex theory was an early account of how this parameter might affect appearance.[117–119] More recent work has examined the relation between physics-based color constancy algorithms and color appearance.[62,79–81,116,120]

Sensitivity Regulation. Another framework for thinking about color context effects has to do with the visual system's tendency to adjust its sensitivity to the ambient viewing conditions. If, as seems likely, the signals mediating the visual system's sensitivity also mediate color appearance, then we would expect sensitivity regulation to cause context effects in color appearance. Regulation of visual sensitivity is discussed in Chap. 00. It is possible that similar mechanisms could subserve both maximization of sensitivity and the achievement of color constancy.[40]

Brightness Matching and Photometry

The foundation of colorimetry is the human observer's ability to judge identity of color appearance. It is sometimes of interest to compare certain perceptual attributes of lights that do not, as a whole, appear identical. In particular, there has been a great deal of interest in developing formulae that predict when two lights with different relative spectral power distributions will appear equally bright. Colorimetry provides a partial answer to this question, since two lights that match in appearance must appear equally bright. Intuitively, however, it seems that it should be possible to set the relative intensities of any two lights so that they match in brightness.

In a heterochromatic brightness-matching experiment, observers are asked to scale the intensity of a matching light until its brightness matches that of an experimentally controlled test light. Although observers can perform the heterochromatic brightness-matching task, they often report that it is difficult and their matches tend to be highly variable.[9] For this reason, more indirect methods for equating the overall effectiveness of lights at stimulating the visual system have been developed. The most commonly used method is that of flicker photometry. In a flicker photometric experiment, two lights of different spectral power distributions are presented alternately at the same location. At moderate flicker rates (about 20 Hz), subjects are able to adjust the overall intensity of one of the lights to minimize the apparent flicker. The intensity setting that minimizes apparent flicker is taken to indicate that the two lights match in their effectiveness as visual stimuli. Two lights equated in this way are said to be equiluminant or to have equal luminance. Other methods are available for determining when two lights have the same luminance.[9]

Because experiments for determining when lights have the same luminance obey linearity properties similar to Grassmann's laws, it is possible to determine a luminous efficiency function that allows the assignment of a luminance value to any light. A luminous efficiency function specifies, for each sample wavelength, the relative contribution of that wavelength to the overall luminance. We can represent a luminous efficiency function as an N_λ dimensional row vector $\mathbf{v}$. Each entry of the matrix specifies the relative luminance of light at the corresponding sample wavelength. The luminance v of an arbitrary spectral power distribution $\mathbf{b}$ may be computed by the equation

$$v = \mathbf{vb} \tag{34}$$

The CIE has standardized four luminous efficiency functions. The most commonly used of these is the standard photopic luminous efficiency function. This is identical to the 1931 XYZ $\bar{y}$ color-matching function. For lights that subtend more than 4° of visual angle, a luminous efficiency function given by the 1964 10° XYZ $\bar{y}$ color-matching function is preferred. More recently, the Judd-Vos modified $\bar{y}$ color-matching function has been made a supplemental standard.[121] A final standard luminous efficiency function is available for use at low light levels when the rods are the primary functioning photoreceptors. The symbol V_λ is often used in the literature to denote luminous efficiency functions. Note that Eq. (34) allows the computation of luminance in arbitrary units. Chapter 24, Vol. 2 of this *Handbook* discusses standard measurement units for luminance.

It is important to note that luminance is a construct derived from flicker photometric and related experiments. As such, it does not directly predict when two lights will be judged to have the same brightness. The relation between luminance and brightness is quite complicated and is the topic of active research.[9,10] It is also worth noting that there is considerable individual variation in flicker photometric judgments, even among color-normal observers. For this reason, it is common practice in psychophysical experiments to use flicker photometry to determine isoluminant stimuli for individual subjects.

Opponent Process Model

We conclude with remarks on what mechanisms might process color information after the initial transduction of light by the cone photoreceptors. Figure 18 shows a framework that in its general form underlies most current thinking. This framework is generally referred to as an opponent process model. It was first proposed in its modern form by Jameson and Hurvich.[122–124] The first stage of the model is cone photoreceptor transduction. We have already considered the implications of this stage in some detail. The L, M, and S cone responses are indicated by the triangles in the figure. After transduction, the opponent process model postulates that the outputs of the cones at each location are recombined to produce a luminance response and two chromatic responses. This recombination is often referred to as an opponent transformation, because the two chromatic responses are postulated to receive excitatory responses from some classes of cones and inhibitory responses from other classes of cones. The luminance and chromatic responses are illustrated by squares in the figure. We have denoted the luminance responses as "Lum" and the two chromatic responses as "R/G" and "B/Y," respectively. The opponent transformation itself is illustrated by a network of connections.

The opponent process framework provides a way to understand the results of flicker

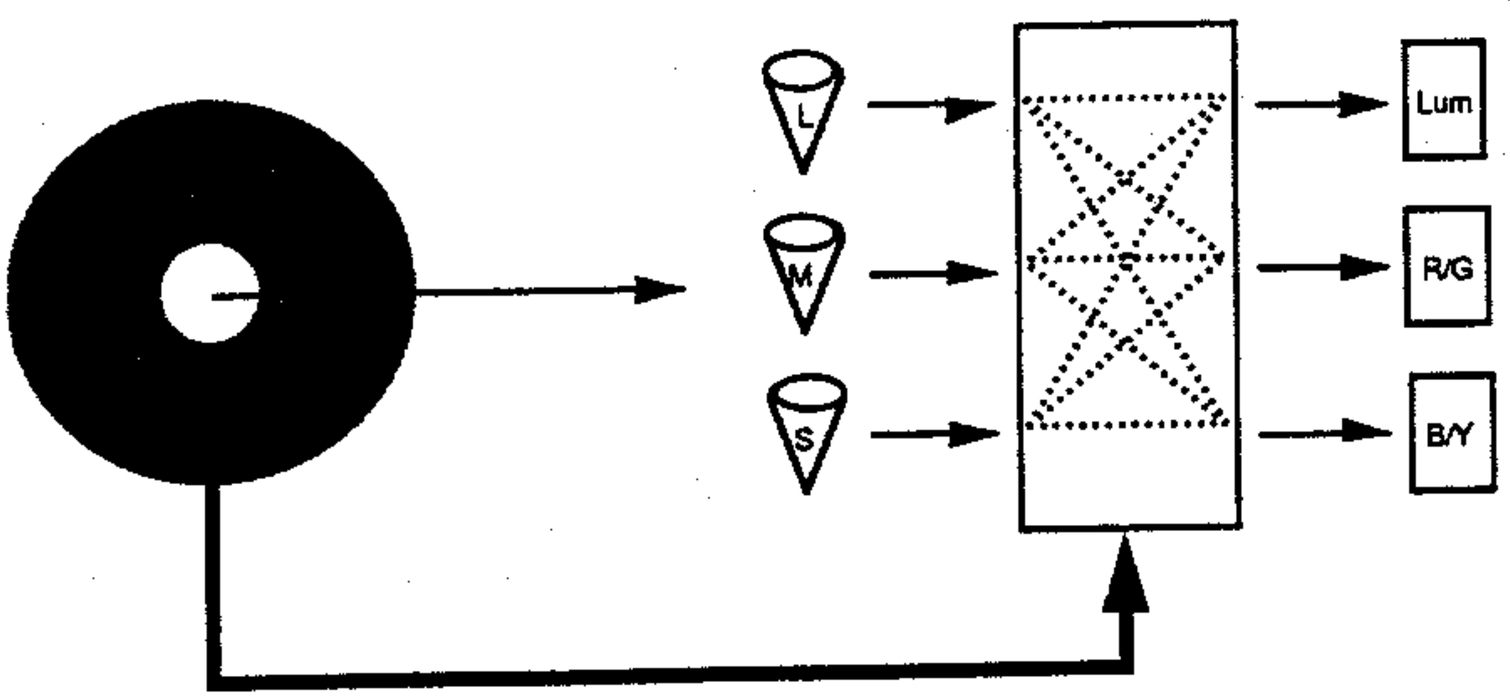

FIGURE 18 Opponent process framework for understanding color vision. See description in the text.

photometric and related experiments. Within the opponent process framework, these experiments are explained as indicating a judgment that depends solely on the luminance response, separately from the two chromatic responses. This interpretation is lent credence by the fact that observers can also make judgments that may be interpreted as tapping solely the chromatic mechanisms.[122,124–127] The opponent process framework has also been employed heavily to account for the effect of context on color appearance.[62,128–132] The idea is that the parameters of the opponent transformation depend on the viewing context. This is illustrated in the figure by the heavy arrow between the annulus that surrounds the test light and the network representing the transformation.

Although the opponent process model provides a useful framework, a number of difficult issues need to be resolved before the framework may be used to extend current colorimetric standards. One major difficulty is that the details of the opponent transformations required to explain different classes of data are not in good agreement. A second difficulty is that physiological data do not seem to support such a simple conception of biological mechanisms that mediate the processing of color information.[133,134]

26.5 APPENDIX A. MATRIX ALGEBRA

This appendix provides a brief introduction to matrix algebra. The development emphasizes the aspects of matrix algebra that are used in this chapter and is somewhat idiosyncratic. In addition, we do not prove any of the results we state. Rather, our intention is to provide the reader unfamiliar with matrix algebra with enough information to make this chapter self-contained.

Basic Notions

Vectors and matrices. A vector is a list of numbers. We use lowercase bold letters to represent vectors. We use single subscripts to identify the individual entries of a vector. The entry a_i refers to the ith number in **a**. We call the total number of entries in a vector its dimension.

A matrix is an array of numbers. We use uppercase bold letters to represent matrices. We use dual subscripts to identify the individual entries of a matrix. The entry a_{ij} refers to the number in the ith row and jth column of **A**. We sometimes refer to this as the ijth entry of **A**. We call the number of rows in a matrix its row dimension. We call the number of columns in a matrix its column dimension. We generally use the symbol N to denote dimensions.

Vectors are a special case of matrices where either the row or the column dimension is 1. A matrix with a single column is often called a column vector. A matrix with a single row is often called a row vector. By convention, all vectors used in this chapter should be understood to be column vectors unless explicitly noted otherwise.

It is often convenient to think of a matrix as being composed of vectors. For example, if a matrix has dimensions N_r by N_c, then we may think of the matrix as consisting of N_c column vectors, each of which has dimension N_r.

Addition and Multiplication. A vector may be multiplied by a number. We call this scalar multiplication. Scalar multiplication is accomplished by multiplying each entry of

the vector by the number. If **a** is a vector and b is a number, then $\mathbf{b} = \mathbf{a}b$ is a vector whose entries are given by $c_j = b_j a$.

Two vectors may be added together if they have the same dimension. We call this vector addition. Vector addition is accomplished by entry-by-entry addition. If **a** and **b** are vectors with the same dimension, the entries of $\mathbf{c} = \mathbf{a} + \mathbf{b}$ are given by $c_j = a_j + b_j$.

Two matrices may be added if they have the same row and column dimensions. We call this matrix addition. Matrix addition is also defined as entry-by-entry addition. Thus, if **A** and **B** are matrices with the same dimension, the entries of $\mathbf{C} = \mathbf{A} + \mathbf{B}$ are given by $c_{ij} = a_{ij} + b_{ij}$. Vector addition is a special case of matrix addition.

A vector may be multiplied by a matrix if the column dimension of the matrix matches the dimension of the vector. If **A** has dimensions N_r by N_c and **b** has dimension N_c, then $\mathbf{c} = \mathbf{A}\mathbf{b}$ is an N_r dimensional vector. The ith entry of **c** is related to the entries of **A** and **b** by the equation.

$$c_i = \sum_{j=1}^{N_c} a_{ij} b_j \tag{A.1}$$

It is also possible to multiply a matrix by another matrix if the column dimension of the first matrix matches the row dimension of the second matrix. If **A** has dimensions N_r by N and **B** has dimensions N by N_c, then $\mathbf{C} = \mathbf{A}\mathbf{B}$ is an N_r by N_c dimensional matrix. The ik^{th} entry of **C** is related to the entries of **A** and **B** by the equation

$$c_{ij} = \sum_{j=1}^{N} a_{ij} b_{jk} \tag{A.2}$$

By comparing Eqs. (A.1) and (A.2) we see that multiplying a matrix by a matrix is a shorthand for multiplying several vectors by the same matrix. Denote the N_c columns of **B** by $\mathbf{b}_1, \ldots, \mathbf{b}_{N_c}$ and the N_c columns of **C** by $\mathbf{c}_1, \ldots, \mathbf{c}_{N_c}$. If $\mathbf{C} = \mathbf{A}\mathbf{B}$, then $\mathbf{c}_k = \mathbf{A}\mathbf{b}_k$ for $k = 1, \ldots, N_c$.

It is possible to show that matrix multiplication is associative. Suppose we have three matrices **A**, **B**, and **C** whose dimensions are such that the matrix products (**AB**) and (**BC**) are both well defined. Then $(\mathbf{AB})\mathbf{C} = \mathbf{A}(\mathbf{BC})$. We often write **ABC** to denote either product. Matrix multiplication is not commutative. Even when both products are well defined, it is not in general true that **BA** is equal to **AB**.

Matrix transposition. The transpose of an N_r by N_c matrix **A** is an N_c by N_r matrix **B** whose ijth entry is given by $b_{ij} = a_{ji}$. We use the superscript "T" to denote matrix transposition: $\mathbf{B} = \mathbf{A}^T$. The identity $(\mathbf{AB})^T = \mathbf{B}^T\mathbf{A}^T$ always holds.

Special Matrices and Vectors. A diagonal matrix **D** is an N_r by N_c matrix whose entries d_{ij} are zero if $i \neq j$. That is, the only nonzero entries of a diagonal matrix lie along its main diagonal. We refer to the nonzero entries of a diagonal matrix as its diagonal entries.

A square matrix is a matrix whose row and column dimensions are equal. We refer to the row and column dimensions of a square matrix as its dimension.

An identity matrix is a square diagonal matrix whose diagonal entries are all one. We use the symbol $\mathbf{I}_N$ to denote the N by N identity matrix. Using Eq. (A.2) it is possible to show that for any N_r by N_c matrix **A**, $\mathbf{A}\mathbf{I}_{N_c} = \mathbf{I}_{N_r}\mathbf{A} = \mathbf{A}$.

An orthogonal matrix **U** is a square matrix that has the property that $\mathbf{U}^T\mathbf{U} = \mathbf{U}\mathbf{U}^T = \mathbf{I}_N$, where N is the dimension of **U**.

A zero vector is a vector whose entries are all zero. We use the symbol $\mathbf{0}_N$ to denote the N dimensional zero vector.

Linear Models

Linear Combinations of Vectors. Equation (A.1) is not particularly intuitive. A more useful way to think about the effect of multiplying a vector **b** by matrix **A** is as follows. Consider the matrix **A** to consist of N_c column vectors $\mathbf{a}_1, \ldots, \mathbf{a}_{N_c}$. Then from Eq. (A.1) we have that the vector $\mathbf{c} = \mathbf{Ab}$ may be obtained by the operations of vector addition and scalar multiplication by

$$\mathbf{c} = \mathbf{a}_1 b_1 + \cdots + \mathbf{a}_{N_c} b_{N_c} \tag{A.3}$$

where the numbers $b_1, \ldots, b_{N_c}$ are the entries of **b**. Thus the effect of multiplying a vector by a matrix is to form a weighted sum of the columns of the matrix. The weights that go into forming the sum are the entries of the vector. We call any expression that has the form of the right-hand side of Eq. (A.3) a linear combination of the vectors $\mathbf{a}_1, \ldots, \mathbf{a}_{N_c}$.

Independence and Rank. Consider a collection of vectors $\mathbf{a}_1, \ldots, \mathbf{a}_{N_c}$. If no one of these vectors may be expressed as a linear combination of the others, then we say that the collection is independent. We define the rank of a collection of vectors $\mathbf{a}_1, \ldots, \mathbf{a}_{N_c}$ as the largest number of linearly independent vectors that may be chosen from that collection. We define the rank of a matrix **A** to be the rank of the vectors $\mathbf{a}_1, \ldots, \mathbf{a}_{N_c}$ that make up its columns. It may be proved that the rank of a matrix is always less than or equal to the minimum of its row and column dimensions. We say that a matrix has full rank if its rank is exactly equal to the minimum of its row and column dimensions.

Linear Models. When a vector **c** has the form given in Eq. (A.3), we say that **c** lies within a linear model. We call N_c the dimension of the linear model. We call the vectors $\mathbf{a}_1, \ldots, \mathbf{a}_{N_c}$ the basis vectors for the model. Thus an N_c dimensional linear model with basis vectors $\mathbf{a}_1, \ldots, \mathbf{a}_{N_c}$ contains all vectors **c** that can be expressed exactly using Eq. (A.3) for some choice of numbers $b_1, \ldots, b_{N_c}$. Equivalently, the linear model contains all vectors **c** that may be expressed as $\mathbf{c} = \mathbf{Ab}$ where the columns of the matrix **A** are the vectors $\mathbf{a}_1, \ldots, \mathbf{a}_{N_c}$ and **b** is an arbitrary vector. We refer to all vectors within a linear model as the subspace defined by that model.

Simultaneous Linear Equations

Matrix and Vector Form. Matrix multiplication may be used to express a system of simultaneous linear equations. Suppose we have a set of N_r simultaneous linear equations in N_c unknowns. Call the unknowns $b_1, \ldots, b_{N_c}$. Conventionally, we would write the equations in the form

$$\begin{aligned} a_{11} b_1 + \cdots + a_{1N_c} b_{N_c} &= c_1 \\ a_{21} b_1 + \cdots + a_{2N_c} b_{N_c} &= c_2 \\ &\cdots \\ a_{N_r 1} b_1 + \cdots + a_{N_r N_c} b_{N_c} &= c_{N_r} \end{aligned} \tag{A.4}$$

where the a_{ij} and c_i represent known numbers. From Eq. (A.1) it is easy to see that we may rewrite Eq. (A.4) as a matrix multiplication:

$$\mathbf{Ab} = \mathbf{c} \tag{A.5}$$

In this form, the entries of the vector **b** represent the unknowns. Solving Eq. (A.5) for **b** is equivalent to solving the system of simultaneous linear equations in Eq. (A.4).

Solving Simultaneous Linear Equations. A fundamental topic in linear algebra is finding solutions for systems of simultaneous linear equations. We will rely on several basic results in this chapter, which we state here.

When the matrix **A** is square and has full rank, it is always possible to find a unique matrix $\mathbf{A}^{-1}$ such that $\mathbf{A}\mathbf{A}^{-1} = \mathbf{A}^{-1}\mathbf{A} = \mathbf{I}_N$. We call the matrix $\mathbf{A}^{-1}$ the inverse of the matrix **A**. The matrix $\mathbf{A}^{-1}$ is also square and has full rank. Algorithms exist for determining the inverse of a matrix and are provided by software packages that support matrix algebra.

When the matrix **A** is square and has full rank, a unique solution **b** to Eq. (A.5) exists. This solution is given simply by the expression $\mathbf{b} = \mathbf{A}^{-1}\mathbf{c}$. When the rank of **A** is less than its row dimension, then there will not in general be an exact solution to Eq. (A.5). There will, however, be a unique vector **b** that is the best solution in a least-squares sense. That is, there is a unique vector **b** that minimizes the expression $\sum_{i=1}^{N_r} ((Ab)_i - c_i)^2$. We call this the least-squares solution to Eq. (A.5). Finding the least-squares solution to Eq. (A.5) is often referred to as linear regression. Algorithms exist for performing linear regression and are provided by software packages that support matrix algebra.

A generalization of Eq. (A.5) is the matrix equation

$$\mathbf{A}\mathbf{B} = \mathbf{C} \tag{A.6}$$

where the entries of the matrix **B** are the unknowns. From our interpretation of matrix multiplication as a shorthand for multiple multiplications of a vector by a matrix, we can see immediately that this type of equation may be solved by applying the above analysis in a columnwise fashion. If **A** is square and has full rank, then we may determine **B** uniquely as $\mathbf{A}^{-1}\mathbf{C}$. When the rank of **A** is less than its row dimension, we may use regression to determine a matrix **B** that satisfies Eq. (A.6) in a least-squares sense. It is also possible to solve matrix equations of the form $\mathbf{B}\mathbf{A} = \mathbf{C}$ where the entries of **B** are again the unknowns. An equation of this form may be converted to the form of Eq. (A.6) by applying the transpose identity (discussed earlier in this appendix). That is, we may find **B** by solving the equation $\mathbf{A}^T\mathbf{B}^T = \mathbf{C}^T$ if $\mathbf{A}^T$ meets the appropriate conditions.

Null Space. When the column dimension of a matrix **A** is greater than its row dimension N_r, it is possible to find nontrivial solutions to the equation

$$\mathbf{A}\mathbf{b} = \mathbf{0}_{N_r} \tag{A.7}$$

Indeed, it is possible to determine a linear model such that all vectors contained in the model satisfy Eq. (A.7). This linear model will have dimension equal to the difference between N_r and the rank of the matrix **A**. The subspace defined by this linear model is called the null space of the matrix **A**. Algorithms to find the basis vectors of a matrix's null space exist and are provided by software packages that support matrix algebra.

Singular Value Decomposition

The singular value decomposition allows us to write any N_r by N_c matrix **X** in the form

$$\mathbf{X} = \mathbf{U}\mathbf{D}\mathbf{V}^T \tag{A.8}$$

where **U** is an N_r by N_r orthogonal matrix, **D** is an N_r by N_c diagonal matrix, and **V** is an N_c by N_c orthogonal matrix.[25] The diagonal entries of **D** are guaranteed to be nonnegative. Some of the diagonal entries may be zero. By convention, the entries along this diagonal

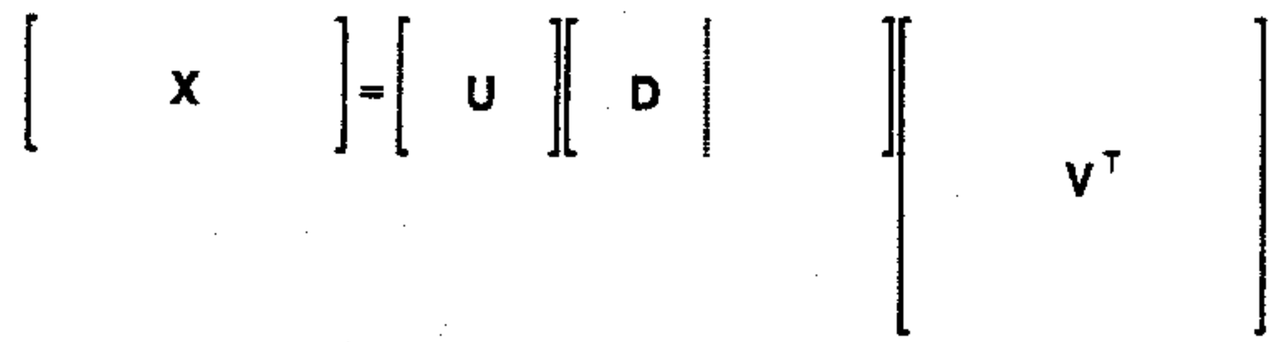

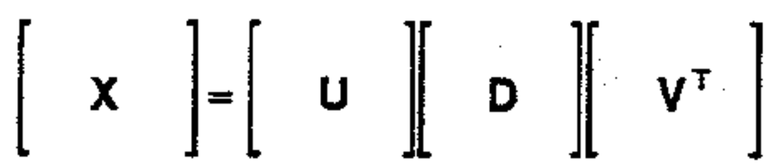

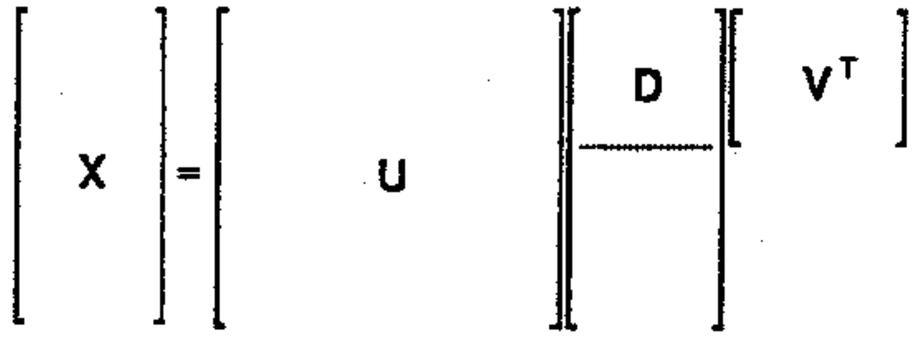

FIGURE A-1 The figure depicts the singular value decomposition (SVD) of an **N** by **M** matrix **X** for three cases $N_c > N_r$, $N_c = N_r$, and $N_c < N_r$.

are arranged in decreasing order. We illustrate the singular value decomposition in Fig. A.1. The singular value decomposition has a large number of uses in numerical matrix algebra. Routines to compute it are generally provided as part of software packages that support matrix algebra.

26.6 ACKNOWLEDGMENT

I thank J. Jacobs, F. Li, J. Loomis, P. Lennie, D. MacLeod, L. Maloney, J. Pokorny. J. Speigle, A. Stockman, J. Tietz, B. Wandell, and D. Williams.

26.7 REFERENCES

1. J. von Kries, "Chromatic Adaptation," originally published in *Festschrift der Albrecht-Ludwigs-Universitat,* 1902, pp. 145–148; D. L. MacAdam (ed.), *Sources of Color Vision,* MIT Press. Cambridge, 1970.
2. R. M. Evans, *An Introduction to Color,* Wiley, New York, 1948.
3. G. Wyszecki, "Color Appearance," *Handbook of Perception and Human Performance,* K. R. Boff, L. Kaufman, and J. P. Thomas (eds.), John Wiley & Sons, New York, 1986.
4. B. K. P. Horn, "Exact Reproduction of Colored Images," *Computer Vision, Graphics and Image Processing* **26**:135–167 (1984).
5. R. W. G. Hunt, *The Reproduction of Colour,* 4th ed., Fountain Press, Tolworth, England, 1987.
6. W. F. Schreiber, "A Color Prepress System Using Appearance Variables," *Journal of Imaging Technology* **12** (1986).
7. R. Hall, *Illumination and Color in Computer Generated Imagery,* Springer-Verlag, New York. 1989.
8. D. H. Brainard and B. A. Wandell, "Calibrated Processing of Image Color," *Color Research and Application* **15**:266–271 (1990).

9. G. Wyszecki and W. S. Stiles, *Color Science—Concepts and Methods, Quantitative Data and Formulae,* 2d ed., John Wiley & Sons, New York, 1982.
10. J. Pokorny and V. C. Smith, "Colorimetry and Color Discrimination," *Handbook of Perception and Human Performance,* K. R. Boff, L. Kaufman, and J. P. Thomas (eds.), John Wiley & Sons, 1986.
11. CIE *Colorimetry,* 2d ed., Bureau Central de la CIE, 1986.
12. ISO/CIE "CIE Standard Colorimetric Observers," 1991.
13. ISO/CIE "CIE Standard Colorimetric Illuminants," 1991.
14. D. H. Krantz, "Color Measurement and Color Theory: I. Representation Theorem for Grassmann Structures," *Journal of Mathematical Psychology* **12**:283–303 (1975).
15. P. Suppes, D. H. Krantz, R. D. Luce, and A. Tversky, *Foundations of Measurement,* Academic Press, San Diego, 1989.
16. D. L. McAdam, *Sources of Color Science,* MIT Press, Cambridge, Mass., 1970.
17. D. Judd and G. Wyszecki, *Color in Business, Science and Industry,* John Wiley & Sons, New York, 1975.
18. G. S. Brindley, *Physiology of the Retina and the Visual Pathway,* 2d ed., Williams and Wilkins, Baltimore, 1970.
19. R. M. Boynton, *Human Color Vision,* Holt, Reinhart and Winston, New York, 1979.
20. B. A. Wandell, "The Synthesis and Analysis of Color Images," *IEEE Transactions on Pattern Analysis and Machine Intelligence* **PAMI-9**:2–13 (1987).
21. J. B. Cohen, "Color and Color Mixture: Scalar and Vector Fundamentals," *Color Research and Application* **13**:5–39 (1988).
22. H. J. Trussell, "Applications of Set Theoretic Methods to Color Systems," *Color Research and Application* **16**:31–41 (1991).
23. B. Noble and J. W. Daniel, *Applied Linear Algebra,* 2d ed., Prentice-Hall, Inc., Englewood Cliffs, N.J. 1977.
24. G. H. Golub and C. F. van Loan, *Matrix Computations,* Johns Hopkins University Press, Baltimore, 1983.
25. J. M. Chambers, *Computational Methods for Data Analysis,* John Wiley & Sons, New York, 1977.
26. W. K. Pratt, *Digital Image Processing,* John Wiley & Sons, New York, 1978.
27. R. A. Johnson and D. W. Wichern, *Applied Multivariate Statistical Analysis,* Prentice-Hall, Englewood Cliffs, N.J., 1988.
28. J. Little and C. Moler, *MATLAB User's Guide.* The Matchworks, Natick, Mass., 1991.
29. W. H. Press, B. P. Flannery, S. A. Teukolsky, and W. T. Vetterling, *Numerical Recipes in C: The Art of Scientific Computing,* Cambridge University Press, Cambridge, 1988.
30. R. A. Becker, J. M. Chambers, and A. R. Wilks. *The New S Language,* Wadsworth & Brooks/Cole, Pacific Grove, Calif., 1988.
31. E. Anderson, Z. Bai, C. Bishof, J. Demmel, J. Dongarra, J. D. Croz, A. Greenbaum, S. Hammarling, A. McKenney, S. Ostrouchov, and D. Sorensen, *LAPACK User's Guide,* Society for Industrial and Applied Mathematics, Philadelphia, 1992.
32. J. G. Grassmann, "Theory of Compound Colors," originally published in *Annalen der Physik und Chemie* **89**:69–84 (1853); D. L. McAdam (ed.), *Sources of Color Vision,* MIT Press, Cambridge, 1970.
33. W. A. H. Rushton, "Visual Pigments in Man," *Handbook of Sensory Physiology,* H. J. A. Dartnall (ed.), Springer, Berlin, 1972.
34. A. Stockman, D. I. A. MacLeod, and N. E. Johnson, "The Spectral Sensitivities of Human Cones," *JOSA A* **10**:2491–2521 (1993).
35. D. B. Judd, "Report of U.S. Secretariat Committee on Colorimetry and Artificial Daylight," *CIE Proceedings* **1** (1951).

36. J. J. Vos, "Colormetric and Photometric Properties of a 2 Degree Fundamental Observer." *Color Research and Application* **3**:125–128 (1978).

37. P. DeMarco, J. Pokorny, and V. C. Smith, "Full-Spectrum Cone Sensitivity Functions for X-Chromosome-Linked Anomalous Trichromats," *Journal of the Optical Society* **A9**:1465–1476 (1992).

38. J. L. Schnapf, T. W. Kraft, and D. A. Baylor, "Spectral Sensitivity of Human Cone Photoreceptors," *Nature* **325**:439–441 (1987).

39. V. Smith and J. Pokorny, "Spectral Sensitivity of the Foveal Cone Photopigments Between 400 and 500 nm," *Vision Research* **15**:161–171 (1975).

40. J. Walraven, C. Enroth-Cugell, D. C. Hood, D. I. A. MacLeod and J. Schnapf. "The Control of Visual Sensitivity: Receptoral and Postreceptoral Processes," *Visual Perception: The Neurophysiological Foundations,* L. Spillmann and J. Werner (eds.), Academic Press, New York, 1989.

41. A. M. Derrington, J. Krauskopf, and P. Lennie, "Chromatic Mechanisms in Lateral Geniculate Nucleus of Macaque," *Journal of Physiology* **357**:241–265 (London, 1984).

42. D. I. A. MacLeod and R. M. Boynton, "Chromaticity Diagram Showing Cone Excitation by Stimuli of Equal Luminance," *Journal of the Optical Society of America* **69**:1183–1186 (1979).

43. J. Krauskopf, D. R. Williams, and D. W. Heeley, "Cardinal Directions of Color Space," *Vision Research* **22**:1123–1131 (1982).

44. J. Krauskopf, D. R. Williams, M. B. Mandler, and A. M. Brown, "Higher Order Color Mechanisms," *Vision Research* **26**:23-32 (1986).

45. D. J. Flitcroft, "The Interactions Between Chromatic Aberration, Defocus and Stimulus Chromaticity: Implications for Visual Physiology and Colorimetry," *Vision Research* **29**:349–360 (1989).

46. A. B. Poirson and B. A. Wandell, "The Ellipsoidal Representation of Spectral Sensitivity." *Vision Research* **30**:647–652 (1990).

47. J. D. Foley, A. van Dam, S. K. Feiner, and J. F. Hughes, *Computer Graphics: Principles and Practice,* 2d ed., Addison-Wesley, Reading, Mass., 1990.

48. D. B. Judd, D. L. MacAdam, and G. W. Wyszecki, "Spectral Distribution of Typical Daylight as a Function of Correlated Color Temperature," *Journal of the Optical Society of America* **54**:1031–1040 (1964).

49. J. Cohen, "Dependency of the Spectral Reflectance Curves of the Munsell Color Chips." *Psychon. Sci.* **1**:369–370 (1964).

50. K. L. Kelly, K. S. Gibson, and D. Nickerson, "Tristimulus Specification of the Munsell Book of Color from Spectrophotometric Measurements," *Journal of the Optical Society of America* **33**:355–376 (1943).

51. D. Nickerson, "Spectrophotometric Data for a Collection of Munsell Samples," U.S. Department of Agriculture, 1957.

52. L. T. Maloney, "Evaluation of Linear Models of Surface Spectral Reflectance with Small Numbers of Parameters," *Journal of the Optical Society of America* **A3**:1673–1683 (1986).

53. E. L. Krinov, "Surface Reflectance Properties of Natural Formations," *National Research Council of Canada: Technical Translation* **TT-439** (1947).

54. J. P. S. Parkkinen, J. Hallikainen, and T. Jaaskelainen, "Characteristic Spectra of Munsell Colors," *Journal of the Optical Society of America* **6**:318–322 (1989).

55. T. Jaaskelainen, J. Parkkinen, and S. Toyooka, "A Vector-Subspace Model for Color Representation," *Journal of the Optical Society of America* **A7**:725–730 (1990).

56. J. R. Magnus and H. Neudecker, *Matrix Differential Calculus with Applications in Statistics and Econometrics,* Wiley, Chichester, 1988.

57. T. W. Anderson, *An Introduction to Multivariate Statistical Analysis,* 2d ed., John Wiley & Sons. New York, 1971.

58. L. T. Maloney, "Computational Approaches to Color Constancy," unpublished Ph.D. thesis, Stanford University, 1984.

59. D. H. Marimont and B. A. Wandell, "Linear Models of Surface and Illuminant Spectra," *Journal of the Optical Society of America* **A9:**1905–1913 (1992).

60. B. A. Wandell and D. H. Brainard, "Towards Cross-Media Color Reproduction," *Proceedings OSA Applied Vision Topical Meeting*, 1989.

61. D. Saunders and A. Hamber, "From Pigments to Pixels: Measurement and Display of the Colour Gamut of Paintings." *Proceedings of the SPIE: Perceiving, Measuring, and Using Color* **1250:**90–102 (1990).

62. D. H. Brainard and B. A. Wandell, "Asymmetric Color-Matching: How Color Appearance Depends on the Illuminant," *Journal of the Optical Society of America* **A9:**1433–1448 (1992).

63. D. R. Williams, N. Sekiguchi, W. Haake, D. H. Brainard, and O. Packer. "The Cost of Trichromacy for Spatial Vision," *From Pigments to Perception*, B. B. Lee and A. Valberg (eds.), Plenum Press, New York, 1991.

64. D. H. Brainard, B. A. Wandell, and W. B. Cowan, "Black Light: How Sensors Filter Spectral Variation of the Illuminant," *IEEE Transactions on Biomedical Engineering* **36:**140–149 (1989).

65. G. Wyszecki, "Evaluation of Metameric Colors," *Journal of the Optical Society of America* **48:**451–454 (1958).

66. S. A. Burns, J. B. Cohen, and E. N. Kuznetsov, "Multiple Metamers: Preserving Color Matches Under Diverse Illuminants," *Color Research and Application* **14:** 16–22 (1989).

67. J. B. Cohen and W. E. Kappauf, "Metameric Color Stimuli, Fundamental Metamers, and Wyszecki's Metameric Blacks: Theory, Algebra, Geometry, Application," *American Journal of Psychology* **95:**537–564 (1982).

68. J. B. Cohen and W. E. Kappauf, "Color Mixture and Fundamental Metamers: Theory, Algebra, Geometry, Application," *American Journal of Psychology* **98:**171–259 (1985).

69. B. A. Wandell. "Color Rendering of Camera Data," *Color Research and Application, Supplement* **11:**S30–S33 (1986).

70. P. Sallstrom, "Color and Physics: Some Remarks Concerning the Physical Aspects of Human Color Vision," University of Stockholm, Institute of Physics, 1973.

71. M. H. Brill, "A Device Performing Illuminant-Invariant Assessment of Chromatic Relations," *J. Theor. Biol.* **71:**473–478 (1978).

72. G. Buchsbaum, "A Spatial Processor Model for Object Colour Perception," *Journal of the Franklin Institute* **310:**1–26 (1980).

73. L. T. Maloney and B. A. Wandell, "Color Constancy: A Method for Recovering Surface Spectral Reflectances," *Journal of the Optical Society of America* **A3:**29–33 (1986).

74. M. D'Zmura and P. Lennie, "Mechanisms of Color Constancy," *Journal of the Optical Society of America* **A3:**1662–1672 (1986).

75. H. Lee, "Method for Computing the Scene-Illuminant Chromaticity from Specular Highlights," *Journal of the Optical Society of America* **A3:**1694–1699 (1986).

76. B. Funt and J. Ho. "Color from Black and White," *Proceedings of the Second International Conference Comp. Vis.*, 1988.

77. B. V. Funt and M. S. Drew, "Color Constancy Computation in Near-Mondrian Scenes Using a Finite Dimensional Linear Model," *IEEE Comp. Vis & Pat. Recog.*, 1988.

78. M. D'Zmura and G. Iverson, "Color Constancy: Basic Theory of Two-Stage Linear Recovery of Spectral Descriptions for Lights and Surfaces," *Journal of the Optical Society of America,* **A10:**2148–2165.

79. M. H. Brill and G. West, "Chromatic Adaptation and Color Constancy: A Possible Dichotomy," *Color Research and Application* **11:**196–204 (1986).

80. L. T. Maloney, "Color Constancy and Color Perception: The Linear Models Framework," *Attention and Performance XIV: Synergies in Experimental Psychology, Artificial Intelligence, and Cognitive Neuroscience,* D. E. Meyer and S. E. Kornblum (eds.), MIT Press, Cambridge, Mass., 1992.

81. J. Pokorny, S. K. Shevell, and V. C. Smith, "Colour Appearance and Colour Constancy," *Vision and Visual Dysfunction,* P. Gouras (ed.), Macmillan, London, 1991.

82. D. H. Brainard, B. A. Wandell, and E. J. Chichilnisky, "Color Constancy: From Physics Appearance," *Current Directions in Psychological Science* **2:**165–170 (1993).

83. D. L. MacAdam, "Visual Sensitivities to Color Differences in Daylight," *Journal of the Opt Society of America* **32:**247–274 (1942).

84. D. L. MacAdam, "Colour Discrimination and the Influence of Colour Contrast on Acuit *Documenta Ophthalmologica* **3:**214–233 (1949).

85. G. Wyszecki and G. H. Fielder, "New Color-Matching Ellipses," *Journal of the Optical Soci of America* **61:**1135–1152 (1971).

86. G. Wyszecki, "Matching Color Differences," *Journal of the Optical Society of Amer* **55:**1319–1324 (1965).

87. W. S. Stiles, "Color Vision: The Approach Through Increment Threshold Sensitivity *Proceedings National Academy of Sciences* **45:**100–114 (U.S.A., 1959).

88. C. Noorlander and J. J. Koenderink, "Spatial and Temporal Discrimination Ellipsoids in Col Space," *Journal of the Optical Society of America* **73:**1533–1543 (1983).

89. A. B. Poirson, B. A. Wandell, D. Varner, and D. H. Brainard, "Surface Characterizations o Color Thresholds," *Journal of the Optical Society of America* **7:**783–789 (1990).

90. A. R. Robertson, "The CIE 1976 Color-Difference Formulae," *Color Research and Applicatio* **2:**7–11 (1977).

91. H. deLange, "Research into the Dynamic Nature of the Human Fovea-Cortex Systems with Intermittent and Modulated Light. I. Attenuation Characteristics with White and Coloured Light," *Journal of the Optical Society of America* **48:**777–784 (1958).

92. H. deLange, "Research into the Dynamic Nature of the Human Fovea-Cortex Systems with Intermittent and Modulated Light. II. Phase Shift in Brightness and Delay in Color Perception," *Journal of the Optical Society of America* **48:**784–789 (1958).

93. N. Sekiguchi, D. R. Williams, and D. H. Brainard, "Efficiency for Detecting Isoluminant and Isochromatic Interference Fringes," *Journal of the Optical Society of America* **A10:**2118–2133 (1993).

94. E. C. Carter and R. C. Carter, "Color and Conspicuousness," *Journal of the Optical Society of America* **71:**723–729 (1981).

95. L. D. Silverstein and R. M. Merrifield, "The Development and Evaluation of Color Systems for Airborne Applications," *DOT/FAA/PM-85-19, U.S. Dept. of Transportation*, Federal Aviation Administration, 1985.

96. A. B. Poirson and B. A. Wandell, "Task-Dependent Color Discrimination," *Journal of the Optical Society of America* **A7:**776–782 (1990).

97. A. L. Nagy and R. R. Sanchez, "Critical Color Differences Determined with a Visual Search Task," *Journal of the Optical Society* **A7:** 1209–1217 (1990).

98. W. S. Stiles and J. M. Burch, "Interim Report to the Commission Internationale de l'Eclairage, Zurich, 1955, on the National Physical Laboratory's Investigation of Color-Matching (1955)," *Optica Acta* **2:**168–181 (1955).

99. W. S. Stiles and J. M. Burch, "NPL Colour-Matching Investigation: Final Report (1958)," *Optica Acta* **6:**1–26 (1959).

100. M. A. Webster and D. I. A. MacLeod, "Factors Underlying Individual Differences in the Color Matches of Normal Observers," *Journal of the Optical Society of America* **A5:**1722–1735 (1988).

101. J. Pokorny, V. C. Smith, and M. Lutze, "Aging of the Human Lens," *Applied Optics* **26:**1437–1440 (1987).

102. J. Pokorny, V. C. Smith, G. Verriest, and A. J. L. G. Pinckers, *Congenital and Acquired Color Vision Defects*, Grune and Stratton, New York, N.Y., 1979.

103. S. Ishihara, *Tests for Colour-Blindness*, Kanehara Shuppen Company, Ltd., Tokyo, 1977.

104. D. Farnsworth, "The Farnsworth-Munsell 100 Hue and Dichotomous Tests for Color Vision," *Journal of the Optical Society of America* **33:**568–578 (1943).

105. S. C. Merbs and J. Nathans, "Absorption Spectra of Human Cone Photopigments," *Nature* **356:**433–435 (1992).

106. J. Winderickx, D. T. Lindsey, E. Sanocki, D. Y. Teller, A. G. Motulsky, and S. S. Deeb, "Polymorphism in Red Photopigment Underlies Variation in Colour Matching," *Nature* **356:**431–433 (1992).

107. M. Neitz, J. Neitz, and G. H. Jacobs, "Spectral Tuning of Pigments Underlying Red-Green Color Vision," *Science* **252:**971–973 (1991).

108. J. Neitz, M. Neitz, and G. H. Jacobs, "More Than 3 Different Cone Pigments Among People with Normal Color Vision," *Vision Research* **33:**117–122 (1993).

109. L. N. Thibos, M. Ye, X. X. Zhang, and A. Bradley, "The Chromatic Eye—A New Reduced-Eye Model of Ocular Chromatic Aberation in Humans," *Applied Optics* **31:**3594–3667 (1992).

110. I. Powell, "Lenses for Correcting Chromatic Aberration of the Eye," *Applied Optics* **20:** 4152–4155 (1981).

111. D. H. Brainard, "Calibration of a Computer Controlled Color Monitor," *Color Research and Application* **14:**23–34 (1989).

112. R. Evans, "Visual Processes and Color Photography," *Journal of the Optical Society of America* **33:**579–614 (1943).

113. R. W. G. Hunt, "Chromatic Adaptation in Image Reproduction," *Color Research and Application* **7:**46–49 (1982).

114. H. Helmholtz, *Helmholtz's Physiological Optics,* translation from the 3d German edition, Optical Society of America, New York, 1909.

115. E. G. Boring, *Sensation and Perception in the History of Experimental Psychology,* D. Appleton Century, New York, 1942.

116. D. H. Brainard and B. A. Wandell, "A Bilinear Model of the Illuminant's Effect on Color Appearance," *Computational Models of Visual Processing,* M. S. Landy and J. A. Movshon (eds.), MIT Press, Cambridge, Mass., 1991.

117. E. H. Land and J. J. McCann, "Lightness and Retinex Theory," *Journal of the Optical Society of America* **61:**1–11 (1971).

118. E. H. Land, "Recent Advances in Retinex Theory and Some Implications for Cortical Computations: Color Vision and the Natural Image," *Proceedings of the National Academy of Sciences* **80:**5163–5169 (U.S.A., 1983).

119. D. H. Brainard and B. A. Wandell, "Analysis of the Retinex Theory of Color Vision," *Journal of the Optical Society of America* **A3:**1651–1661 (1986).

120. A. C. Hurlbert, H. Lee, and H. H. Bulthoff "Cues to the Color of the Illuminant," *Investigative Ophthalmology and Visual Science,* supplement **30:**221 (1989).

121. CIE, "CIE 1988 2° Spectral Luminous Efficiency Function for Photopic Vision," 1990.

122. D. Jameson and L. M. Hurvich, "Some Qualitative Aspects of an Opponent-Colors Theory. I. Chromatic Responses and Spectral Saturation," *Journal of the Optical Society of America* **45:**546–552 (1955).

123. L. M. Hurvich and D. Jameson, "An Opponent-Process Theory of Color Vision," *Psychological Review* **64:**384–404 (1957).

124. D. Jameson and L. M. Hurvich, "Theory of Brightness and Color Contrast in Human Vision," *Vision Research* **4:**135–154 (1964).

125. J. Larimer, D. H. Krantz, and C. M. Cicerone, "Opponent-Process Additivity—I. Red/Green Equilibria," *Vision Research* **14:**1127–1140 (1974).

126. J. Larimer, D. H. Krantz, and C. M. Cicerone, "Opponent Process Additivity—II. Yellow/Blue Equilibria and Non-Linear Models," *Vision Research* **15:**723–731 (1975).

127. D. H. Krantz, "Color Measurement and Color Theory: II. Opponent-Colors Theory," *Journal of Mathematical Psychology* **12:**304–327 (1975).

128. C. M. Cicerone, D. H. Krantz, and J. Larimer, "Opponent-Process Additivity—III. Effect of Moderate Chromatic Adaptation," *Vision Research* **15:**1125–1135 (1975).

129. J. Walraven, "Discounting the Background—The Missing Link in the Explanation of Chromatic Induction," *Vision Research* **16**:289–295 (1976).

130. S. K. Shevell, "The Dual Role of Chromatic Backgrounds in Color Perception," *Vision Research* **18**:1649–1661 (1978).

131. J. S. Werner and J. Walraven, "Effect of Chromatic Adaptation on the Achromatic Locus: The Role of Contrast, Luminance and Background Color," *Vision Research* **22**:929–944 (1982).

132. S. K. Shevell and M. F. Wesner, "Color Appearance Under Conditions of Chromatic Adaptation and Contrast," *Color Research and Application* **14** (1989).

133. P. Lennie, J. Krauskopf, and G. Scalar, "Chromatic Mechanisms in Striate Cortex of Macaque," *Journal of Neuroscience* **10**:649–669 (1990).

134. R. L. DeValois and K. K. DeValois, "A Multi-Stage Color Model," *Vision Research* **33**:1053–1065 (1993).