A Field Guide to Digital Color

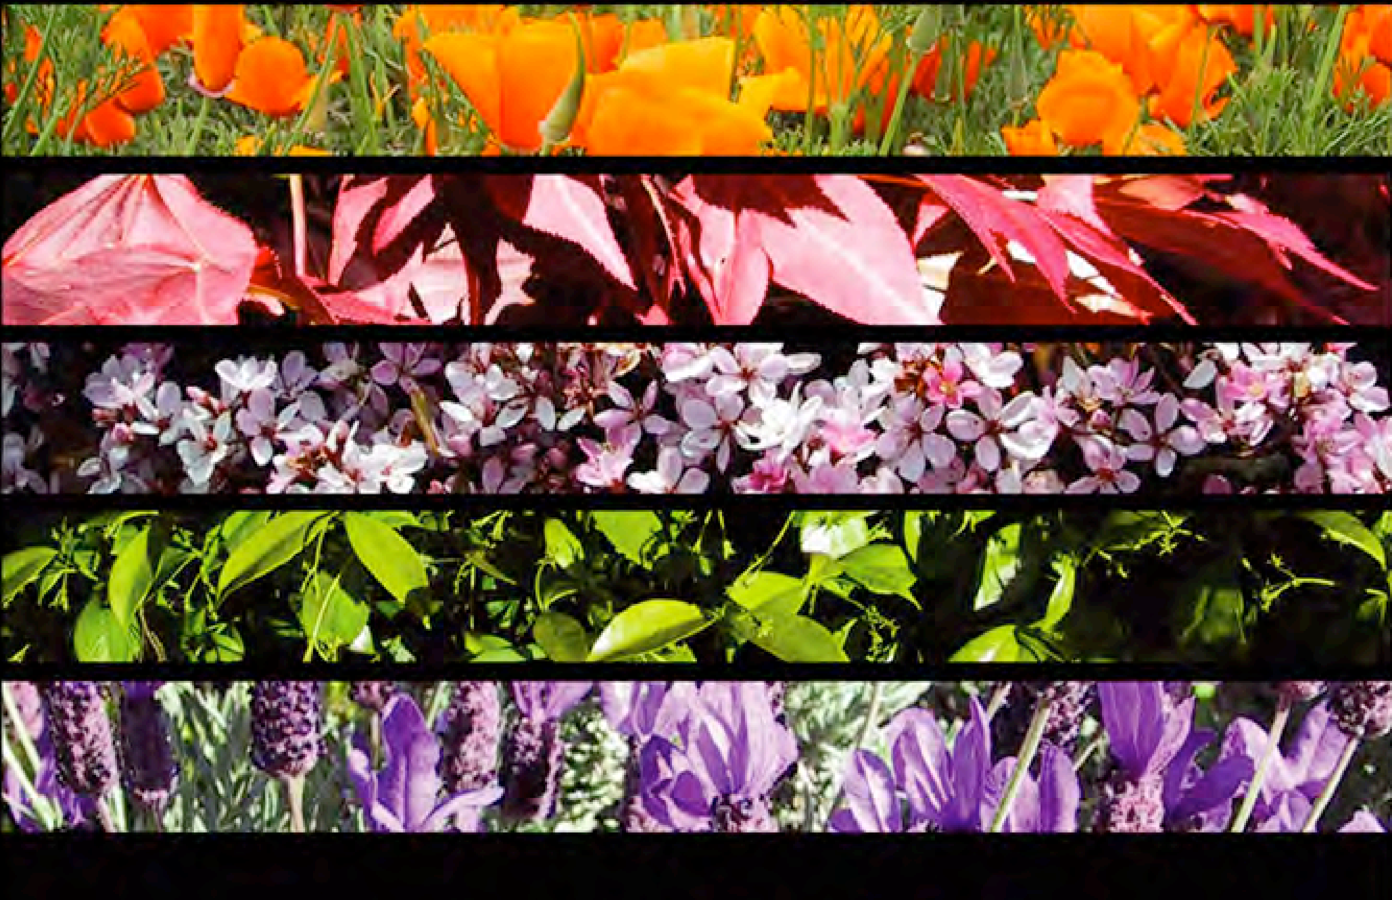

MAUREEN C. STONE

A Field Guide to Digital Color

A Field Guide to Digital Color

Maureen C. Stone

A K Peters
Natick, Massachusetts

CRC Press
Taylor & Francis Group
6000 Broken Sound Parkway NW, Suite 300
Boca Raton, FL 33487-2742

© 2003 by Taylor & Francis Group, LLC
CRC Press is an imprint of Taylor & Francis Group, an Informa business

No claim to original U.S. Government works
Version Date: 20140121

International Standard Book Number-13: 978-1-4665-7833-3 (eBook - PDF)

This book contains information obtained from authentic and highly regarded sources. Reasonable efforts have been made to publish reliable data and information, but the author and publisher cannot assume responsibility for the validity of all materials or the consequences of their use. The authors and publishers have attempted to trace the copyright holders of all material reproduced in this publication and apologize to copyright holders if permission to publish in this form has not been obtained. If any copyright material has not been acknowledged please write and let us know so we may rectify in any future reprint.

Except as permitted under U.S. Copyright Law, no part of this book may be reprinted, reproduced, transmitted, or utilized in any form by any electronic, mechanical, or other means, now known or hereafter invented, including photocopying, microfilming, and recording, or in any information storage or retrieval system, without written permission from the publishers.

For permission to photocopy or use material electronically from this work, please access www.copyright.com (http://www.copyright.com/) or contact the Copyright Clearance Center, Inc. (CCC), 222 Rosewood Drive, Danvers, MA 01923, 978-750-8400. CCC is a not-for-profit organization that provides licenses and registration for a variety of users. For organizations that have been granted a photocopy license by the CCC, a separate system of payment has been arranged.

Trademark Notice: Product or corporate names may be trademarks or registered trademarks, and are used only for identification and explanation without intent to infringe.

Visit the Taylor & Francis Web site at
http://www.taylorandfrancis.com

and the CRC Press Web site at
http://www.crcpress.com

To My Family

Contents

Preface

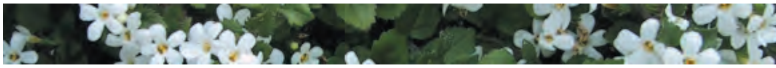

Color is a part of many different fields, from vision and perception to art and design. The subset of color, called *digital color,* in which color is encoded and manipulated as digital values, still spans many fields, including: image reproduction, computer graphics, and multi-media design. Each field has its own reference books, technical publications and conferences, making it easy to focus on one field without being aware of relevant work in others. This book is a "field guide," designed to provide a high-level summary of each field, with pointers to in-depth resources. For those interested in digital color, this is the book to read first.

Who Should Read this Book?

- Students learning about color, especially those in computer science, engineering, and digital media.
- Engineers and computer scientists who need to learn about digital color.
- Graphic artists and digital photographers, who are tired of books that present only "how to" without describing "why?"
- Experts in one field of color who want to learn about other fields.
- Anyone wanting a broad introduction to digital color.

How this Book Is Structured

Each chapter provides a summary of a color field, with a focus on digital color—many fields spread widely beyond pixels. One of the challenges of a book such as this is to balance between superficial and overwhelming. Ideally, experts will find the description of their field an accurate but high-level summary. For those unfamiliar with the field, the description should seem detailed to the point of needing careful reading. Along with providing information specific to the field, the presentation cross-references other fields, and emphasizes common concepts.

Each chapter identifies one or more significant books for its field towards the end of the introduction. Other books may be mentioned throughout the chapter as needed. All books are referenced informally by name, with the full citation information included in the bibliography. Within the chapters, there are occasional references to technical papers that provide detail not found in any book. For these, the full citation information is contained in the margin near where the paper is mentioned. There was a deliberate decision to be selective, rather than inclusive, in the choice of these references, as it is impossible to reference all potentially relevant literature in a book of this breadth without overwhelming the reader. Many of the named papers are good starting points for finding other papers in their area.

There is a resources section at the end of the book that includes a bibliography and pointers to professional organizations, journals, and standards bodies. The list of books is comprehensive, but far from exhaustive, and contains primarily those books I personally own and have found useful. The bibliography is introduced by a discussion that includes a sentence or two about each book. This is not a true annotated bibliography, with a detailed paragraph for each entry, but an effort to qualitatively describe the strengths, weaknesses and application of each book. The description of each organization and standards body includes the journals and magazines published by the organization, and a website for further information.

The book does include mathematical equations to precisely describe concepts such as the transformation between different color representations. Formulas are valuable for implementers, and for those for whom the math speaks more clearly than words. The math is presented, however, in such a way that it can be skimmed or ignored by those not interested in that level of detail.

Overview of the Chapters

Chapter 1: Color Vision. The basic principles of color vision, in which light strikes the eye and is encoded in the retina as exactly three signals, underlie the use of RGB for encoding color. These same principles are used as the basis for measuring color, including representing color as points on the CIE chromaticity diagram, examples of which are used throughout the book.

Chapter 2: Color Appearance. The basic principles presented in Chapter 1 provide a simple model for color perception that expands the principles of color appearance. This is best demonstrated by the various optical illusions showing that what a color "looks like" depends on context; for example, in the right circumstances, the identical color (as measured) can appear either black or white. Included is an overview of the computational models that are being used to incorporate these principles into digital color systems. Many color appearance principles underlie color design principles.

Chapter 3: RGB and Brightness. In digital color, colors are defined as numerical values, which have no inherent visual meaning. A fundamental requirement for understanding digital color as presented in this book is knowing how to bind these numbers to specific physical and psychophysical metrics. The principles presented in this chapter, even without the mathematical detail, provide a critical foundation for the rest of the book.

Chapter 4: Color in Nature. Color in the natural world is created by the interaction of light with matter. This chapter is included to provide a reference for color synthesis in such digital color fields as computer graphics, and to provide a deeper understanding of some of the physical processes that underlie color reproduction technologies such as printing and photography.

Chapters 5–9 are a set of 5 related chapters on the topic of color reproduction. In this book, the color reproduction fields are presented with an emphasis on digital color reproduction and the use of "device-independent color," which is color defined with respect to the psychophysical metrics described in Chapters 1 and 3.

Chapter 5: Color Reproduction. Traditionally, the field of color reproduction is focused on the production of images that are pictures of the natural world. This chapter provides an overview of the general principles of color image reproduction as it is traditionally practiced in photography, television, and printing.

Chapter 6: Image Capture. In image capture, the "natural world" becomes image pixels. The technology of image capture includes cameras and scanners.

Chapter 7: Additive Color Systems. Additive color systems output image pixels as colored light. The technologies of additive color include displays and digital projectors. Originally designed for television, additive color systems are the most familiar physical form for digital color, due to their use in computer displays.

Chapter 8: Subtractive Color Systems. Digital color images must be converted to subtractive color dyes and inks to be printed on paper or projected as slides or movies. This chapter describes the principles and technologies that underlie printing and photography.

Chapter 9: Color Management Systems. Color management systems, which are based on using device-independent color representations for digital color, can be used to create reliable color transformations from one digital medium to another. This chapter provides insight, for example, on how to create prints that are a good reproduction of the pixels displayed on a monitor. It is not a cookbook for any particular system, but a presentation of the basic principles behind all color management systems, including why problems occur and ways to maintain some degree of control.

Chapter 10: Computer Graphics. The field of computer graphics blends the physical models of Chapter 4 with the principles of color reproduction presented in Chapter 5. This chapter focuses on 3D rendering, which is the process that creates images from numerically-defined 3D models.

Chapter 11: Color Selection and Design. The principles of color selection and design presented here are primarily those that define the aesthetic use of color as taught in graphic design. These principles are more algorithmic and rule-based than those outside of the field may appreciate. Other color design principles, such as text legibility, are readily defined in terms of the color appearance principles defined in Chapter 2.

Chapter 12: Color in Information Display. The effective use of color for information display can be described by Tufte's principles, the primary of which is to "do no harm." This field includes concepts from color perception, color reproduction, and color design that can be applied to visualization, illustration, and user interface design.

Acknowledgements

My education in digital color began at the Xerox Palo Alto Research Center, where I first tried to print illustrations that had been designed on a computer display. This led to work with Bill Cowan and John Beatty that was one of the first efforts to apply color science to the problem of printing computer-generated images. Bill, who was working with Günter Wyszecki at the National Research Council of Canada, provided expertise in color science and key insights about merging color and digital technology. We received the invaluable help of Dusty Rhodes, who was one of Xerox's resident experts on color printing, with a long and venerable career in the graphic arts. He patiently taught us about high-quality printing and color image reproduction, whose subtleties are not easily modeled in digital color systems. This experience gave me my initial expertise in digital color and my on-going enthusiasm for the study of color.

Throughout my career at PARC, I had opportunity to learn from other color folks at Xerox, including: Gary Starkweather, Mik Lamming, Dale Green, Chuck Haines, and Rob Buckley. Another impor-

tant part of my professional experience has come from my participation in the SID/IS&T Color Imaging Conferences, where I've had the opportunity to meet and learn from Bob Hunt, Mark Fairchild, Joann Taylor, and Lindsay MacDonald, among many others.

They say it takes a village to raise a child. It can also take a "village" to create a book, especially if one is a working parent. Special thanks goes to Joel Bartlett, who read each chapter as I wrote it and helped me decide when it was "good enough." Joel also helped with technical information about photography and digital cameras. Joann Taylor checked my color science and set me straight about color difference formulae. Bill van Melle proofread the entire book, correcting the grammer and improving the presentation. Bethany Kanui created the index, and Mark Wong helped with the illustrations. Russ Atkinson printed many drafts and lent a sympathetic ear. Jan Hardenbergh offered regular encouragement, feedback, and a helping hand.

Iris, Alice, and the staff at A K Peters provided flexible support, hard work, and a patient tolerance with my perfectionism and resulting delays. They were a pleasure to work with.

Finally, I want to thank my husband, Doug Wyatt, and my sons, Robin and Kevin, for their loving and patient support. Like many books, this one took far more time than expected. Thanks, guys—mom has *finally* finished her book.

Contributors

There are many people who provided images or data used in this book: Arthur Amezcua, Jason Babcock, Joel Bartlett, Steve Bennet, Cynthia Brewer, Michael Cohen, Elaine Cohen, Mike Cook, Jeffrey DiCarlo, James Ferwerda, Andrew Glassner, Amy Gooch, Bruce Gooch, Steven Gortler, Donald Greenberg, Radek Grzeszczuk, Chet Haase, Eric Haines, Pat Hanrahan, Jan Hardenbergh, Jack Holm, Paul Hubel, Henrik Wann Jenson, Brad Johanson, Jeremy Johnson, Sterling Johnson, Joe Kane, Russell Kerschmann, Marc Levoy, Bruce Lindbloom, Steve Marschner, John McCann, Lindsay MacDonald, Aditi Majumder, Greg McNutt , Gary Meyer, David Nadeau, Thor Olson, Mary Orr, Sumanta Pattanaik, Arjun Ramamurthy, Erik

Reinhard, David Salesin, Peter Shirley, Michael Stark, Richard Szeliski, Joann Taylor, Kenneth Torrance, Ben Trumbore, Pauline Ts'o, Brian Wandell, and George Winkenbach.

Special Contributors

Bruce Lindbloom created custom renderings of color gamuts for this book and provided helpful information about the practice of color management in the graphic arts. Bruce's website is a valuable resource for digital color: www.brucelindbloom.com.

Pauline Ts'o scanned in a large selection of her lovely photographs for use in this book, and also taught me about color management at Rhythm & Hues.

Andrew Glassner created several custom illustrations of basic graphics techniques for this book. Andrew's website is: www.glassner.com.

Sterling Johnson sent me an entire CDROM full of stunning bubble pictures. I wish I could have used more of them. Sterling's website is: www.bubblesmith.com.

David Nadeau, from the San Diego Supercomputer Center, told me the fascinating story of how they created the colors for the emmision nebula simulation produced for the Hayden Planetarium in New York City.

Pat Hanrahan and the Stanford University Graphics Lab contributed the use of their color measurement equipment, file server space and the opportunity to learn about projection displays.

1

Color Vision

Color vision starts with light striking the retina, a photosensitive organ in the back of the eye. The cells in the retina encode the light into signals, which are transmitted to the brain. There, the signals are interpreted to provide the perception of color. This chapter focuses on color vision at the encoding level, where it is represented by three signal values generated by the retina. This encoding is the foundation for color measurement, and for the representation of color as a triple of red, green, and blue values.

Introduction

Light enters the eye as a spectrum of colors distributed by wavelength. The spectrum impinges on the retina in the back of the eye, and for color vision, is absorbed by special cells called *cones*. Human beings have three types of cones, which respond to different wavelengths of light. These are called long, medium, and short wavelength cones (LMS), or, more informally, red, green, and blue cones (RGB).

Each cone absorbs light and sends a signal along the optic nerve that indicates the total amount of light energy "seen" by the cone. This signal is interpreted in the brain to create the perception of color (Figure 1). This process has two important implications for understanding color vision. First, the spectral distribution is collapsed into exactly three values, one for each cone. Second, these values, not the spectrum, are what represents color at its earliest perception. These principles are called *trichromacy* and *metamerism*, respectively.

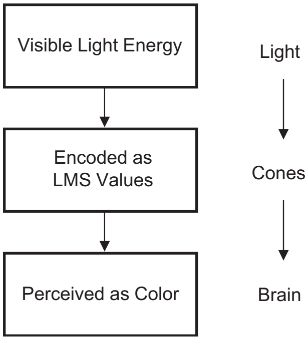

Figure 1.

Process diagram of basic color vision. Light strikes the cones, which encode signals for the brain.

Color measurement, or *colorimetry,* can use these principles to encode color as three numbers. Similarly, digital color imaging uses these principles to encode color as RGB pixels, which are also represented by three numbers. This is not coincidence: color vision at the retinal encoding level is essentially an RGB encoding system. Basic vision, as presented in this chapter, is wonderfully orderly, involving concise mathematical operations on spectra and their encodings. However, as the next chapter will show, encoding is only the beginning of color perception. The full perception of color is much less tidy.

There are many excellent books on color vision, several of which are listed in the bibliography. For understanding digital color, my favorite is Brian Wandell's *Foundations of Vision,* which is structured in three parts: encoding, representation, and response. What this chapter describes is essentially the encoding part of his taxonomy. *Color Science*, by Wyszecki and Stiles, is considered the standard reference for quantitative color science and colorimetry. Roy Berns' update of the classic Billmeyer and Saltzman *Principles of Color Technology* is an up-to-date reference on color, its perception, and its measurement.

The first sections in this chapter define the basics of color vision: describing color as a spectrum, the retinal sensors, and the principles of trichromacy and metamerism. The second half describes the principles behind colorimetry, especially those specified by the Commission Internationale de l'Eclairage (CIE).

Light

Visible light can be described as a function of power versus wavelength. This function is called a *spectral distribution function*, or *spectrum*. Colored light contains wavelengths in the range 370–730 nanometers (nm). The color of the light depends on the distribution of the energy over the visible spectrum; different wavelengths appear different colors.

Each wavelength is associated with a color. Wavelength colors vary from purple (violet) and purple-blue (indigo) through blue, green, yellow, and orange to red, which are the familiar colors of the rainbow. Longer than visible wavelengths are infrared and shorter ones are ultraviolet light. Figure 2 shows the spectral distribution for a fluorescent light bulb. The wavelength colors are roughly illustrated by the colors in the band under the figure. The true colors would be much more intense.

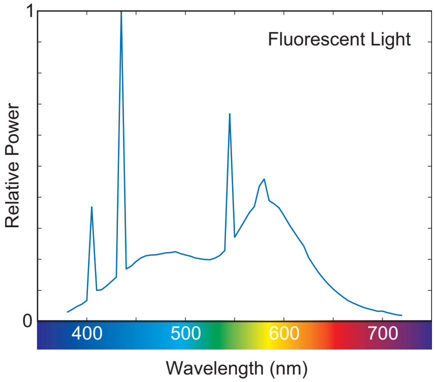

Figure 2.
The spectral distribution function for a fluorescent light. The colored band shows approximate wavelength colors.

The height of a spectral distribution function indicates its power—the taller the brighter. Two spectra that are simply scaled multiples of each other will have the same color, one brighter than the other. For example, Figure 3 is the same purple color at different brightness levels. The area under the spectrum is a measure of brightness called the *intensity*.

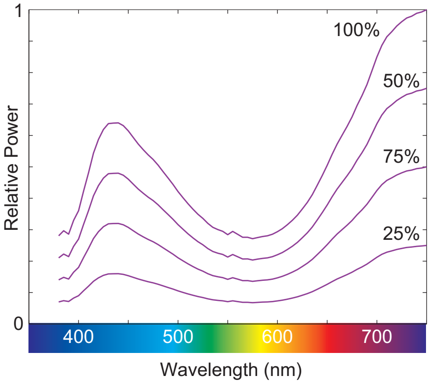

Figure 3.
The same purple color at 100%, 75%, 50%, and 25% intensity.

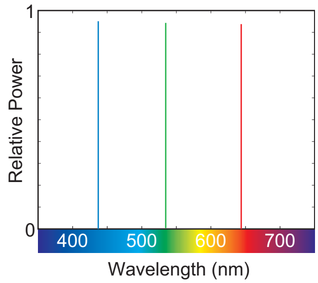

Figure 4.

Three monochromatic colors. Viewed individually, each is a pure, saturated blue, green, or red, as indicated. Viewed all together, the color would look white.

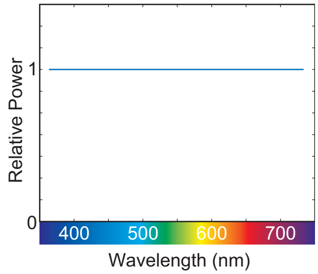

Figure 5.

Equal energy white—the spectrum is the same at all wavelengths.

Spectral distributions that include only a single wavelength describe an intense, pure, saturated color, as in laser light. Such colors are plotted as spikes, as shown in Figure 4 (the colors are approximate).

Spectral distributions that contain all colors in approximately equal amounts will appear white. *Equal energy white* is the color created by presenting equal energy at all wavelengths (Figure 5). While not readily found in nature, it is often used in computer graphics and other computational processes as a canonical white light.

Different mixes of wavelengths will give different colors. For example, the monochromatic colors in Figure 4 are each a pure, saturated color. Viewed all together as a mixture, they will appear white. The precise color seen depends on both the spectrum and the cones, as described below.

Note that it is more common in signal processing and engineering to plot spectra as a function of frequency, which is the inverse of wavelength. This means that red is a low-frequency color and appears on the left of the graph, and blue-violet on the right. This is also the derivation of the terms infrared and ultraviolet. However, color and vision scientists consistently use wavelength plots, where red is on the right.

Colored Objects

The color of an object is defined by two spectra: the surface reflectance of the object, and the light source shining on it. The product of these two spectra is the light that enters the eye and stimulates the cones. Surface reflectances are often specified as percentages, where 100% reflects all of the available light.

Changing the light creates a different spectral result, resulting in a different color. This is why paint samples in the store, for example, look different than when viewed at home. This is illustrated in Figure 6, which shows the result of multiplying the same spectrum by two different light sources, one representing daylight, and the other incandescent light. In this plot, both surface reflectance

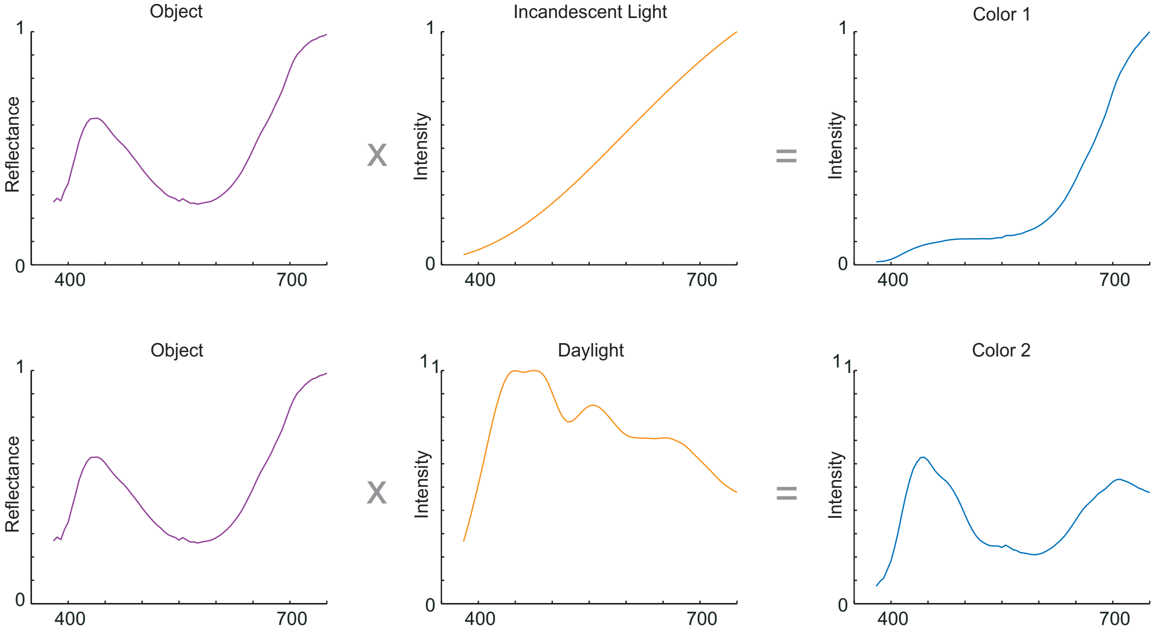

Figure 6.
Different lights reflecting off the same surface create different spectra, and hence different colors.

and intensity are normalized to the range 0…1, to make them easier to evaluate.

The full perception of object color is complex. For example, a painted wall in a room with the light from both a lamp and a window shining on it is perceived to be all one color, even though the light varies significantly across it. Some of these effects are treated in Chapter 2, which describes color appearance. Within this chapter, the color of an object is simply a spectral distribution, which changes as the light changes.

The Retina, Rods, and Cones

The retina covers the back of the eye. Light enters through the pupil, and is focused by the cornea and the lens on the retina, similar to the way an image is focused on film or on the imaging array in a digital camera. There are two types of imaging sensors on the retina—the *rods* and the *cones*. Only the cones are used for normal color vision.

We see color when the light is bright. As the total light dims (dim room light, or twilight), color vision begins to simplify towards the gray, colorless view one sees "in the dark." Both rods and cones contain photopigments, which generate signals on the optic nerve when exposed to light. Rod photopigments are sensitive at low light levels, but saturate at about the point at which cone photopigments begin to respond. Rod and cone vision are formally called *scotopic* and *photopic* vision, respectively. There is a narrow range of lighting where both rods and cones are active, which is called *mesopic*. Color vision as described here is strictly photopic, involving only the cones.

The response of each cone can be encoded as a function of wavelength called the *spectral response curve* for the cone, as shown in Figure 7. Note that cone response curves overlap significantly, especially the medium (green) and long (red) curves.

Multiplying the spectrum of the incoming light, called the *stimulus*, by the response curve and integrating to get the resulting intensity defines the signal that is sent from the eye to the brain. This is illustrated in Figure 8. The viewed light is multiplied by the three cone response curves producing three filtered "views" of the spectrum, one for each cone. The area under each of these curves is the signal sent on the optic nerve, as illustrated by the height of the bars.

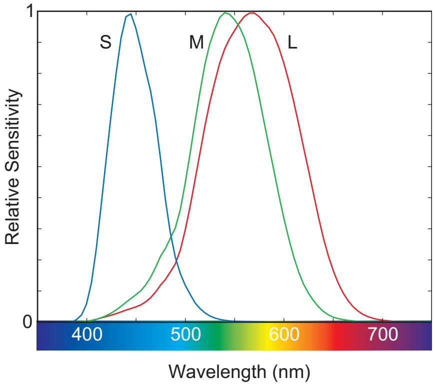

Figure 7.
The spectral response curves for the short (blue), medium (green), and long (red) cones.

Unlike manufactured arrays of light sensors, such as those used in digital cameras, which are laid out in a uniform rectangular array, the rods and cones are unevenly distributed in the retina. The cones are concentrated in the center of the eye, in the region called

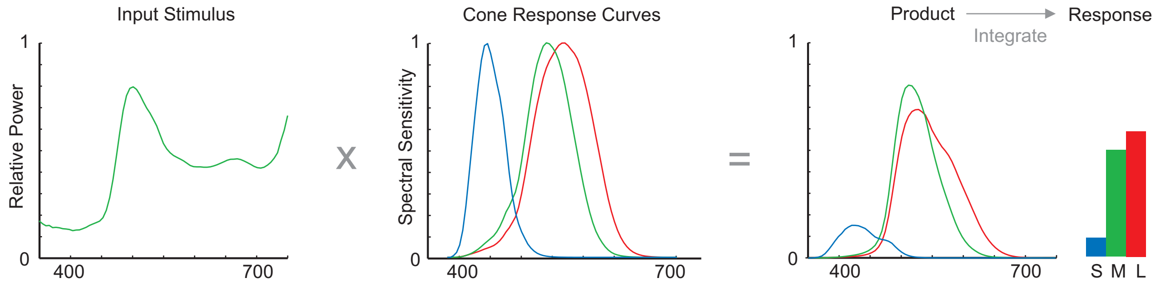

Figure 8.
Multiplying a spectrum by the cone response curves and integrating creates the basic color signal to the brain. The height of the bars reflects the strength of the three signals.

the *fovea*. The rods, plus a few, scattered cones, cover the rest of the retina. There are approximately 5 million cones and 100 million rods in each human eye. There is also a blind spot in the retina, a region that has no rods or cones, where the retinal receptors attach to the optic nerve at the back of the eye. The fovea covers less than 10% of the retina, yet is responsible for all of the color signals sent from the eye to the brain.

How does this uneven mosaic of receptors create a continuous image in our brains? The answer is that our eyes constantly move, in a pattern of rapid jumps called *saccades*. It is left to the brain to stitch together all the images to create a continuous color image.

Trichromacy and Metamerism

The encoding created by the cones means that every spectrum is represented by exactly three signals—this is the principle of trichromacy. That color can be described by exactly three values is a result that appears over and over again in color science and engineering. Colored pixels in digital color are encoded as three values (red, green, and blue) due to trichromacy.

Different animals have different models for color vision. While most primates are trichromats, birds and fish often have more sensors, indicating that their visual system operates in four or more dimensions. Animal cones may respond to different wavelengths than human cones. Many insects, for example, have sensors that operate in the ultraviolet range, which is invisible to people. Other animals, for example, cats and dogs, appear to have fewer cones than humans, and current thinking suggest that they may see color only in the blue-green part of the spectrum. Previously, it was thought that they saw only shades of gray, though probably with more precision than humans do.

The principle of metamerism states that different spectra that produce the same encoded signals look the same color. Color is defined by the product of the spectrum with the cone response, not the spectrum alone. The principle of metamerism underlies all color reproduction technologies. Instead of reproducing the original spectral distribution that described the color, they create an equivalent response, or *metameric match* by mixing three colors in a controlled way. Figure 9 shows a pair of spectra that create (approximately) the same purple color as shown: one is smooth, taken from a colored object lit by white light source, and the other is a mixture of red, green, and blue phosphor emissions from a computer display. These spectral distributions are very different, but appear the same.

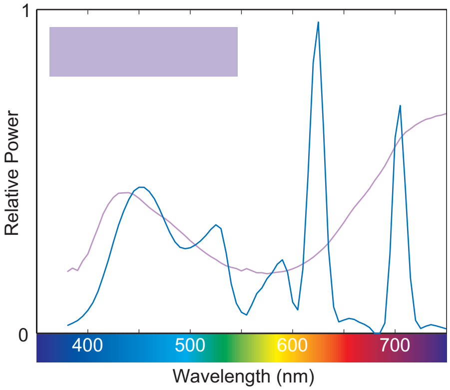

Figure 9.
Two different spectra create the same color, shown as a purple bar, illustrating the property of metamerism.

Trichromacy and metamerism can be used to create instruments that measure color. Creating an instrument that responds in the same way as the human eye is critical for many industries. It allows colored materials to be described impartially and quantitatively, and creates metrics that define when colors match. From the discussion above, it might seem obvious to fit an instrument with filters and sensors that behave like the cones. However, the precise definition of the cone response was not known until relatively recently. The science of color measurement is much older, and is based on experiments involving matching colors with sets of three primary lights.

Colorimetry

Colorimetry is the science of color measurement. It is based on empirical studies of humans matching colors. This results are used to create three *color-matching functions* that can be used convert any spectrum into a standard encoding, just as the cones convert any stimulus into LMS signals.

The color-matching experiments that underlie colorimetry are constructed as follows. Choose three primary lights (call them red, green, and blue), which can vary in intensity. Then, take a set of reference colors such as the monochromatic colors of the spectrum, or colors generated by filtering a white light. The job of the observer is to combine and adjust the primary colors to create a result that matches the reference color. This is shown schematically in Figure 10. Once the match is made, the color can then be defined simply by describing the amount of each primary needed to match it.

This is a remarkably useful result. It means that any color can be described as three numbers, which correspond to the amount of each primary used to create the color. These three numbers are called the *tristimulus values* for the color.

What about colors that cannot be matched? The answer is to use "negative light." For example, if the blend is too reddish, even with

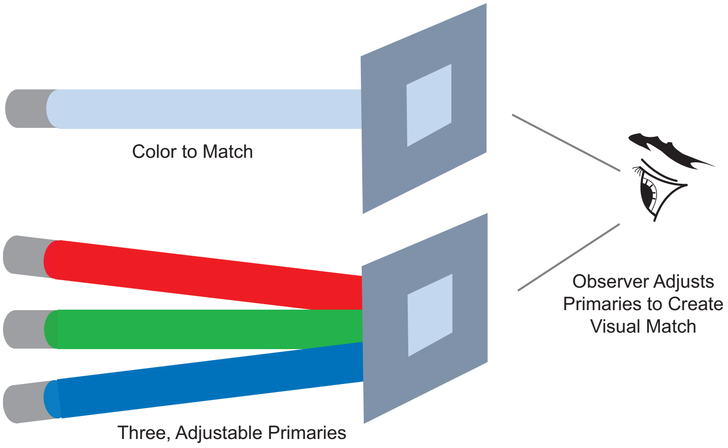

Figure 10.
Color-matching experiments are the basis for colorimetry. An observer adjusts three primary lights to match each sample color.

no red shining on the blend, the observer is allowed to shine some of the red primary on the reference to make the match. The amount of red shone on the reference color is the "negative red" included in the match. It is therefore possible to match all other colors with *any* set of three, distinct colors, assuming you can use "negative light" as described above. In this context, "distinct" means that no primary color is simply a brighter version than another primary.

Defining a color by three numbers plus a known set of primaries is clearly a substantial simplification over the full spectral representation of the color. Color matches defined this way, however, are defined only for colors described as spectrum. For colored objects, the color must be viewed on a standard background under standard lighting to produce a single spectrum as a stimulus. The resulting match, therefore, is valid only for those lighting and viewing conditions.

It would be very inconvenient to have to perform a color-matching experiment to define the tristimulus values for each new color. Fortunately, color matching can be defined more generally because colored light and its perception work as an additive system—the result of shining several lights on the same point is simply the sum of their spectra. A spectrum, therefore, can be considered as the sum of a number of monochromatic colors, one per wavelength. This was alluded to in Figure 4, where the sum of the three monochromatic colors is perceived as white.

Color matching is also additive; the match to a sum of spectra is the sum of the matches. To spell this out mathematically, for a given spectrum, S, let the tristimulus values be R, G, and B (RGB). If RGB_1 matches S_1, and RGB_2 matches S_2, then $RGB_1 + RGB_2$ will match $S_1 + S_2$. This principle was first formalized by Grassmann, and it is called Grassmann's additivity law. This principle can be applied to predict the colors produced by additive color systems such as digital color displays, as described in Chapters 3 and 7.

Now, let us return to the problem of color matching. To find the tristimulus values for an arbitrary color, we use Grassmann's additivity law to construct a set of color-matching functions for our primaries. First, we perform the color-matching experiment on each of the monochromatic spectral colors (sampled every nanometer, for example). The result is a set of three functions of wavelength, one for each primary, whose height is the amount of the primary needed to match that particularly wavelength.

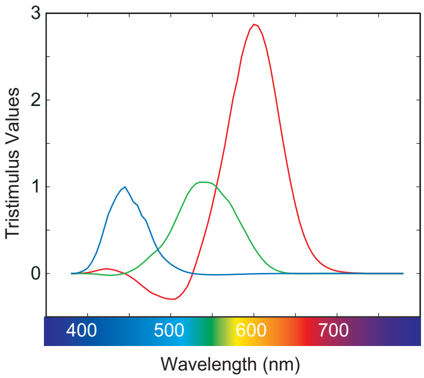

Figure 11. The color-matching functions defined by Stiles and Burch using three monochrome light sources: R = 645 nm, G = 525 nm, and B = 444 nm.

Figure 11 shows a set of color-matching functions constructed by Stiles and Burch in 1959 using a set of three monochromatic colors, as indicated. The red curve is negative in the shorter wavelengths, indicating that some "negative red" is needed to create a match in those regions. All color-matching functions measured from real lights include negative lobes because the long and medium cone response curves overlap. Therefore, there is no way to fully stimulate the red and green cones independently of each other.

Given a set of color-matching functions and a color defined by a spectrum, the tristimulus values are determined by multiplying the spectrum by the color-matching functions and integrating the results. This is the same process shown in Figure 8 to illustrate the cone response, which suggests that color-matching functions are equivalent to cone response curves. To explore this, we need to look further at the mathematical properties of color mixture.

Given a set of color-matching functions, it is possible to create equivalent functions for a new set of primary colors by linear transformation. All we need to know is the definition of the new primaries in terms of the old ones. For those familiar with linear algebra, this is simply a change of basis, and can accomplished for a three dimensional system with a 3 × 3 transformation matrix. Applying this matrix to each of the samples in the color-matching functions in turn creates the color-matching functions for the new set of primaries. As a result, all color-matching functions are linear transformations of each other.

The cone response functions can be shown to be a linear transformation of the color-matching functions once the optical properties of the eye, such as chromatic aberration, are properly included. This result was finally demonstrated in the late 80s, over 50 years after the first color-matching functions were measured.

CIE Colorimetry

In 1931, the Commission Internationale de l'Eclairage (CIE) standardized a set of primaries and color-matching functions that are the basis for most color measurement instruments used today. They transformed a set of measured color-matching functions, similar to those of Stiles and Burch, to create a set of curves that were more convenient to use. These are notated $\bar{x}$, $\bar{y}$, and $\bar{z}$ and are shown in Figure 12. This set is positive throughout the entire visible spectrum, and one of the curves ($\bar{y}$) can be used to compute the perceived brightness of the measured color, called its *luminance*. Luminance will be discussed in more detail in Chapter 2. The function z is zēro for most wavelengths because it was more convenient for the hand calculations that were common at the time. There are no physical lights that can be used to create these functions; the primaries are now mathematical abstractions.

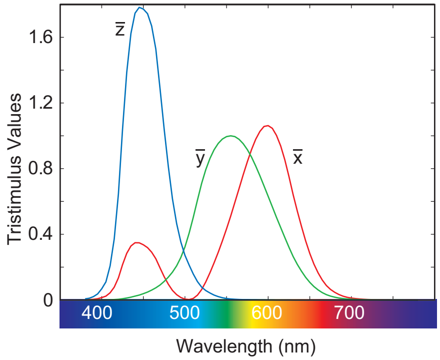

Figure 12.
CIE recommended color-matching functions (1931), also called the CIE 1931 standard observer. $\bar{y}$ can also be used to calculate luminance.

The CIE tristimulus values are computed from the CIE color-matching functions and are notated X, Y, and Z. The color-matching functions have been normalized so that an input spectrum with a constant value of one (equal energy unit spectrum) would create (X, Y, Z) = (1, 1, 1), as shown in Figure 13. If all input spectra are normalized to lie in the range (0, 1), the resulting tristimulus values

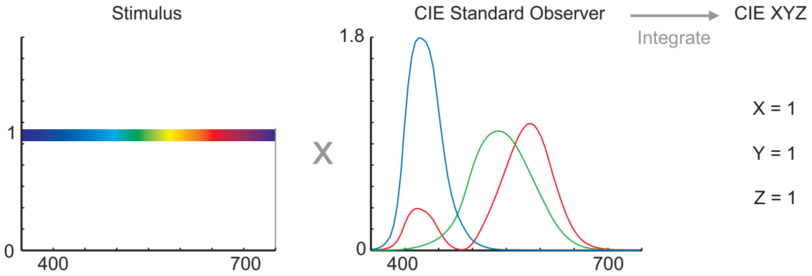

Figure 13.
The CIE tristimulus values are calculated by multiplying the stimulus times the color-matching functions and integrating the result. They are normalized such that equal energy white gives XYZ = (1, 1, 1), as shown.

all fall inside the unit cube—no tristimulus value would exceed 1.0. Both absolute and normalized tristimulus values are used in color measurement.

It is difficult to visualize and compare three-dimensional values, and in many applications, it is useful to be able to define a color independently from its brightness. The CIE *chromaticity diagram*, shown in Figure 14, is a two-dimensional projection of the tristimulus values that is the most common way of visualizing them. It is derived from the CIE XYZ values by dividing them by their sum, as shown in the equations. The resulting chromaticity values, notated x, y and z, have the property that x + y + z = 1. The chromaticity diagram is a plot of x versus y.

This diagram is worth some study. Any color (a spectrum of light) can be plotted as a point, independent of its brightness (power). Also, all metameric matches to the spectrum, which by definition have the same tristimulus values, will plot at the same point.

The chromaticity diagram illustrates how color is additive. Take two colors, each of which plots as a point in the chromaticity diagram. Any colors that are mixtures of these two colors will plot along the line connecting them. The distance along this line is proportional to the relative brightness of the two colors. For three colors, all mixtures will lie within the triangle formed by their chromaticity coordinates, and so on.

The monochromatic spectral colors lie along the horse-shoe shaped path, which is called the *spectrum locus*. All visible colors lie within the shape bounded by this path and the line that connects

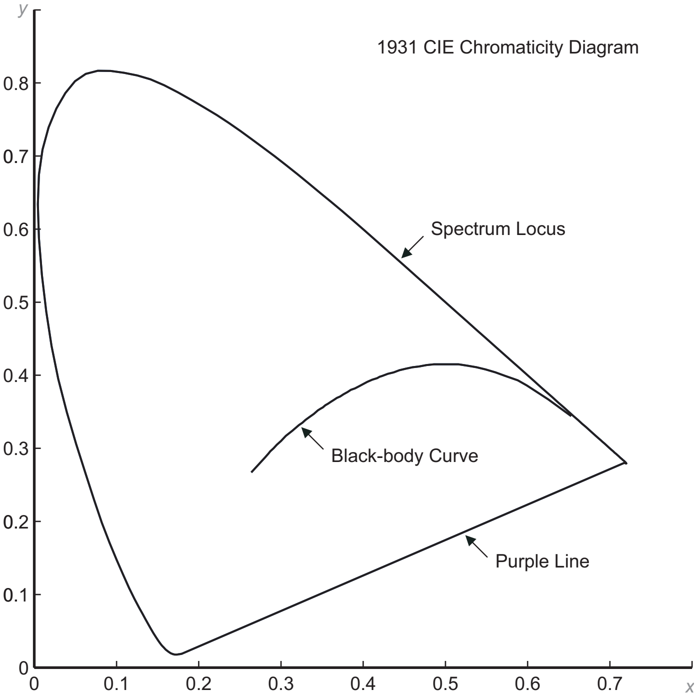

Figure 14.
CIE 1931 chromaticity diagram showing the spectrum locus, the purple line, and the black-body curve. The transformation from CIE XYZ is shown in the equations.

$$x = \frac{X}{X+Y+Z}$$

$$y = \frac{Y}{X+Y+Z}$$

$$1 = x+y+z$$

the two ends (called the *purple line*) because all spectra can be created by sums of monochromatic colors. The closer a color lies to the spectrum locus, the more saturated (vivid, colorful) it is. Take a point near the center of the diagram as a reference white and draw a line out to the spectrum locus. The line defines a family of colors that have the same basic hue, but become progressively more pure (saturated) as they approach the spectrum locus. The colorimetric term *purity* describes the relative distance along such a line.

Colors that appear white lie near the center of the chromaticity diagram. Most of these colors are clustered near the *black-body curve*, which describes colors caused by heating an object that reflects no light (called a black-body radiator) until it glows white-hot. These colors range from yellowish-red to blue-white.

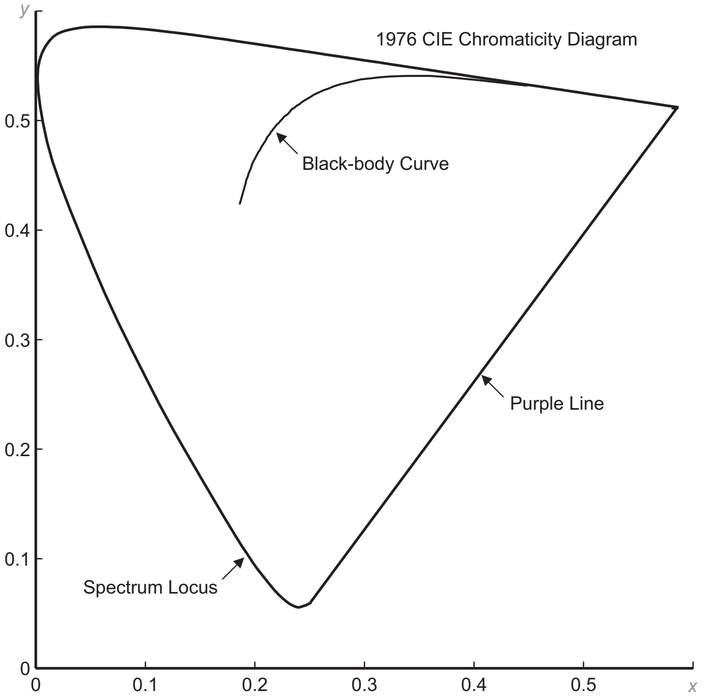

Figure 15.
CIE 1976 chromaticity diagram showing the spectrum locus, the purple line, and the black-body curve. The transformation from CIE XYZ is shown in the equations.

$$u' = \frac{4X}{X + 15Y + 3Z}$$
$$v' = \frac{9Y}{X + 15Y + 3Z}$$

One major concern with the 1931 chromaticity diagram is that distances are not proportional to perceptual differences in color. For example, two points that are within, say, 0.1 units of each other in the reddish (lower right) part of the chart will be perceptually more different than two colors the same distance away in the greenish (upper left) part of the chart. Therefore, there have been many variations on this diagram that attempt to make it more uniform. One of the most commonly used is the u′ v′ system, shown in Figure 15, along with its supporting mathematical transformation. Which chromaticity diagram to use is often very much a matter of practice and technical "taste." In spite of the many improvements proposed and adopted throughout the years, the 1931 diagram is still the most common.

Colorimetry and Colored Objects

Colorimetry can only be applied directly to the light reflecting from a colored object, which can be described as a spectral distribution. Colorimetry applied to objects, therefore, will always depend on the spectral description of the light. Many instruments for measuring colored objects contain a light source, and are held against the surface of the object to eliminate all other light.

Some authors multiply the surface reflectance by the color-matching functions to define "tristimulus values" for the surface. This is a rather suspect practice that stretches the definition of colorimetry beyond its origins in visual matching. Physically, this is equivalent to illuminating the surface with a uniform white light (equal energy at all wavelengths) to create a spectral stimulus. The result, however, is different than any color the object will produce in practice, because such light sources do not exist in nature. Furthermore, such tristimulus values do not follow the additive principles described above. What does it mean to blend two object colors? The potential for confusion hardly seems to compensate for the possible convenience.

Color Measurement

Instrumentation for tristimulus-based color measurement comes in two forms: instruments that sample spectra and compute tristimulus values based on the samples, and instruments that use filters to approximate the color-matching functions. The first usually have "spectro-" in their name, and are either *spectroradiometers* (which measure light directly and absolutely) or *spectrophotometers* (which measure relative reflectance, and often contain their own light source). The second form are *colorimeters*, and are often specifically designed for monitors, reflection prints, or transmission films. In general, colorimeters are cheaper, faster, and less accurate than systems that sample spectra.

All color instrumentation comes with a specification that tells its precision and the range of brightness it can accurately measure. It is important to pay attention to this information: there is no point in agonizing over a 1% luminance variation, for example, if the instrument is only accurate to 2%. Similarly, measuring something that is too bright may simply generate bad data, not an error message, and may also damage the instrument.

All instrumentation needs routine calibration and evaluation to be sure it is working correctly. In addition to following the manufacturer's recommendations, it is useful to create some standard measurements, and make them routinely to see if the instrument has drifted. Similarly, for any task, it is important to find some "reality check" measurements to ensure you don't waste time gathering garbage. For example, when measuring a CRT display, the sum of the tristimulus values for red, green, and blue should equal the tristimulus values for white, because white is the sum of red, green, and blue.

Most color measurement equipment can be connected to a computer via a serial port. The vendors often offer both a way to interface the instrument to a program and some end-user software for color measurement. My experience is that this software is often not as good as the equipment it runs—these folks understand color measurement, not software. Similarly, be prepared for some quirks in the serial port driver, as the "standards" in this area seem to be rather loosely interpreted, a problem that is not unique to color measurement equipment. It is most convenient if the interface can be driven by sending text codes down the serial port line. This is generally robust, and can be easily interfaced to any system, or even driven by hand from a terminal emulator.

Before color management systems created a need for inexpensive and convenient color measurement, most colorimetric instruments were expensive and difficult to use. The use of color measurement in color management systems (described in Chapter 9) for the graphic arts and digital photography has created a market for less expensive and easier to use equipment. While often not as precise as the high-end instruments, such equipment is usually adequate for digital color imaging needs.

Densitometric Measurement

Densitometry is such a common form of color measurement that it is described here, even though it is not based on color-matching functions and the principles of CIE colorimetry. A *densitometer* measures the amount of light transmitted through or reflected from a material and encodes it logarithmically. *Density* is defined as $-\log_{10}(T)$ or $-\log_{10}(R)$, where T (transmittance) or R (reflectance) lie in the range (0,1). Therefore, a transmittance of 1 is a density of 0, a transmittance of 0.001 (0.1%) is a density of 3.0, and the density for a transmittance of 0 is undefined.

Color densitometers measure density through narrow-band color filters, such as those shown in Figure 16. These are designed to selectively measure colorants in photographic or print media. Their spectral response functions are not color-matching functions, and cannot be used to compute tristimulus values. True color-matching functions always overlap, as do the curves in Figure 11 and Figure 12.

Densitometers are often used in digital color to measure the dyes and inks used in print and film, the subtractive color processes. The peaks of their color filters are designed to match the specific colorants being measured.

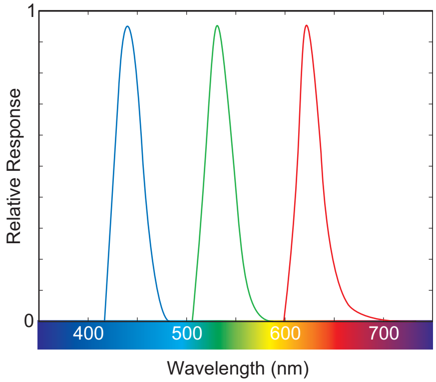

Figure 16.
Color filters for densitometry (ISO Status A), recommended for measuring the colors in photographic prints. These are not color-matching functions.

Summary

Light in the visible spectrum stimulates the cones in the retina, which encode and send messages along the optic nerve to produce the sensation of color. Different spectra can produce the same response, which is the property called metamerism. The encoding can be represented by three numbers because the human visual system is basi-