Color

CHAPTER 9

9.1 INTRODUCTION

So far, when we have talked about *color* we have usually been talking about *wavelength* of visible light. Although different wavelengths are associated with different colors, equating color and wavelength can often be very misleading—our intuitive idea of color is quite different from the result of a precise wavelength determination. In fact, while we associate a given wavelength with a particular color, we may nevertheless see that color when that wavelength is absent or not see it when the wavelength is present. Indeed, we see colors that are not even in the spectrum of visible light—that are not associated with any wavelength. We'll see that there is more to color than meets the eye, and, just as accurately, there is less to it than meets the eye.

Color is different for different people (Fig. 9.1). Consider the color green; it may evoke the freshness and life of lush vegetation or the decay of mold and slime. One dictionary gives one meaning as "fresh, youthful, vigorous," and the next meaning as "pale and sickly." We think of green as implying fertility (a green thumb), seasickness, money, gullibility, strangeness (little green men), immaturity, jealousy, poison, or life. Our individual perceptions of green vary widely. Thus, one of the authors, whose color vision is deficient, can barely distinguish green from gray, while the painter Kandinsky looks at green and sees "a fat, healthy, immovably resting cow, capable only of eternal rumination, while dull bovine eyes gaze forth vacantly into the world."

In this chapter, we'll discuss the important facts about color, on which we can generally agree. In Chapter 10, we'll try to understand how color phenomena are dependent on our perceptual mechanisms.

9.2 COLOR VERSUS WAVELENGTH, AND NONSPECTRAL COLORS

When we break up white light into its component wavelengths, say by means of a prism as in Plate 2.1a, we see all the colors of the spectrum spread out according to the wavelength of the light. Is there a simple relation between the colors and the wavelengths? If you ask a collection of observers to locate a given color in the spectrum, there will be general, but not total, agreement. Most observers will identify blue with a wavelength between 455 and 485 nm, green with a region between 500 and 550 nm, yellow between 570 and 590 nm, and red with a wavelength somewhere above 625 nm. Thus the colors are not generally identified with a unique wavelength. Further, the identification depends somewhat on the intensity of the light—a wavelength that appears somewhat red at low intensity may seem orange as the intensity is increased. And the naming of colors is hardly an exact science. For instance, ***monochromatic**** or ***spectral*** light (that is, ideally, light consisting of only one wavelength) with $\lambda = 600$ nm has been identified as orange chrome, golden poppy, spectrum orange, bitter sweet orange, oriental red, saturn red, cadmium red orange, red or-

*Greek *monos*, single, plus *chroma*, color.

FIGURE 9.1

ange, and yet other names in a list that continues to grow as advertisers take over the English language.

If you compare the monochromatic colors in a good spectrum with the colors you normally see around you, you'll find that most colors you see don't lie in the spectrum. For example, none of the colors in Plate 8.4 is identical to any of the colors seen in a spectrum cast by a prism. A few examples of nonspectral colors are purple, pink, brown, silver, fluorescent red, and iridescent green. Where do these other colors come from? How do we get the sensation of all these colors that do not appear in the spectrum?

9.3 THE INTENSITY-DISTRIBUTION CURVE AND THE CLASSIFICATION OF COLORS

Most colors around us are *not* monochromatic, but instead contain a *distribution of wavelengths*. Suppose we reflect white light from one of the greenish parts of Plate 8.4 and break up the reflected light with a prism. If we then measure the intensity at each visible wavelength and plot the result, we will get an ***intensity-distribution curve*** somewhat like the solid curve in Figure 9.2. Although there is a predominance of green light, there is also a little bit of every other visible wavelength present. Looking at the greenish part of the plate, then, your eye receives this entire distribution of different wavelengths.

When you simultaneously play two different notes on a piano, you usually hear the two separate notes. But when you shine two different wavelengths onto the same place on a screen, your eye doesn't separate the resulting light into two colors; rather you see some sort of *mixture*. How can your *sensations* of these mixtures be characterized? For example, consider what happens if we shine a little extra red light along with the greenish light of Fig. 9.2:

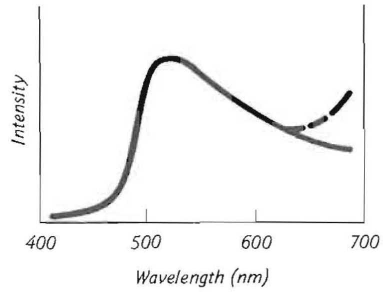

FIGURE 9.2

Intensity-distribution curve. Solid line: the intensity of light, at each visible wavelength, obtained when white light is reflected from a greenish region of Plate 8.4. Dashed line: the same light with a little extra red light mixed in.

the color mixture is changed—indeed, there are an infinite number of other ways that we could modify the intensity-distribution curve, and most people can distinguish a million or so such different color mixtures. Recognizing this huge variety of colors, the author Nabokov felt that a phrase like "the sky is blue" conveyed little information. Thus in *Speak, Memory* he attempts to express the variety of the color blue with such terms as "blue-white," "misty-blue," "purplish-blue," "silvery blue," "cobalt blue," "indigo blue," "azure," "china-blue," "dove-blue," "crystal blue," and "ice-bright." These are certainly more poetic than intensity-distribution curves, but may not convey the same image to every reader. What we want is a straightforward way of classifying all the distinguishable color sensations that is simpler than giving the intensity at each visible wavelength, but still contains all the necessary information.

We'll see that rather than specifying *all* the numbers in the intensity-distribution curve, for many purposes it is sufficient to specify only *three* numbers. Your eye cannot extract all the information contained in the curve by looking at the light; many lights with very different curves appear the same to you. The three qualities of the colored light that determine how the light appears are hue, saturation, and brightness.

Hue corresponds to the *main color* or *color name;* it is what distinguishes one spectral color from another. For example, all yellows differ in hue from all blues, regardless of any other possible similarities. Hue is specified by the *dominant wavelength* in an intensity-distribution curve. Thus, the greatest intensities in the light represented by Figure 9.2 lie between 500 and 530 nm, so the hue of that light is some sort of green. (The term dominant wavelength is defined more precisely in Sec. 9.4C—the "dominant wavelength" may actually not be present, even though the color looks like light of that wavelength.)

Saturation corresponds to the *purity* of the color—a very saturated color generally has almost all its intensity fairly close to the dominant wavelength, while an unsaturated color would have contributions from many other wavelengths. The monochromatic, spectral colors have the highest saturation. White light, which generally consists of all wavelengths with no dominant one, is completely *un*saturated (Fig. 9.3a). (White is often thought of symbolically as the essence of purity—brides wear it. From our point of view, it is the *least* pure color you can get.) Other colors may be thought of as a mixture of white light with a saturated color. Figure 9.3b shows the intensity-distribution curve of a saturated red light. A mixture of that saturated red with white (Fig. 9.3c) gives a desaturated red, or pink. The saturation is a measure of how much dominant wavelength there is compared to the amount of white mixed in. Most objects around you have unsaturated colors because their dominant hue comes from absorption, which takes place somewhat below their surface, but there is also some surface reflection, which is independent of wavelength and thus mixes in some white. (See the TRY IT.) Artists use the terms *chroma* or (unfortunately) *intensity* to mean something similar to saturation. (Because the definitions come from

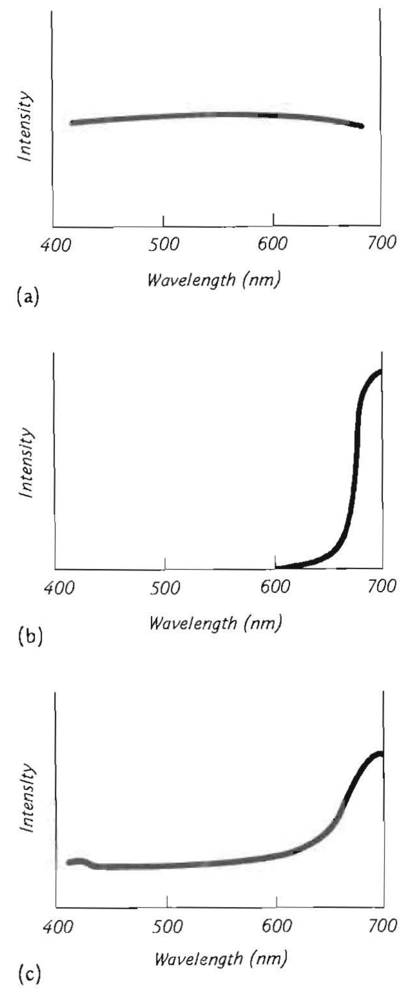

FIGURE 9.3

Saturation. **(a)** White light is completely unsaturated. **(b)** A saturated red light. **(c)** A less saturated red light—pink.

different fields, the meanings are slightly different.)

The third attribute of color depends on whether we're talking about colored lights or colored surfaces. ***Brightness*** refers to the sensation of overall intensity of a *light*, ranging from dark, through dim, to bright, and dazzling (Fig. 9.4a, see also Sec. 7.3A). ***Lightness,*** on the other hand, is related to the percentage of incident light reflected by a *surface* and refers to the *blackness, grayness, or whiteness* of a color (Fig. 9.4b).

We know how to change the brightness of a light (even a colored

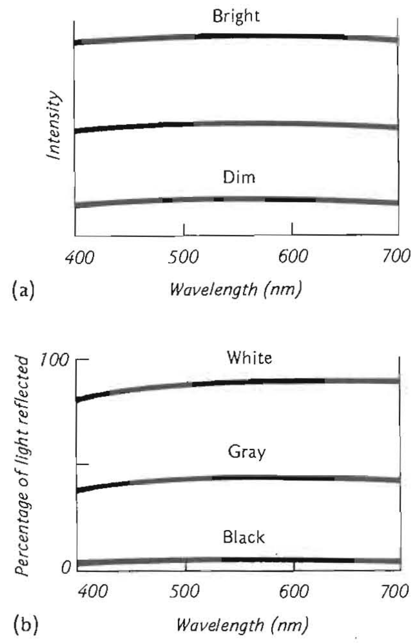

FIGURE 9.4

(a) Brightness of a light. The intensity-distribution curves for three different brightnesses. **(b)** Lightness of a surface. The curves correspond to the percentage of incident light reflected at each wavelength.

one); we need only adjust the intensity of the source. To change the lightness of a surface you must make it less reflecting, say by lightly rubbing pencil all over it. The surface will then appear grayer—you have decreased its lightness (Fig. 9.4b). Thus, the more reflecting the surface (in the visible), the lighter it is. We often think of lightness as being related to the total amount of light reflected at all visible wavelengths—but we must be careful. We can change the amount of reflected light in two ways. One is to make the surface *less reflecting*, say by the method just described. This indeed lowers the lightness. The other is to change the overall *illumination*—the incident light. But if you shine less light on a piece of white paper, it doesn't appear gray, it continues to look white. That is, the lightness of any surface is nearly independent of the overall illumination (lightness constancy—Sec. 7.3B).

To apply these ideas to color, take an orange surface and blacken it, lowering each point on the intensity-distribution curve of its reflected light by a factor of two, perhaps, without changing the relative heights. You have then decreased the surface's lightness without changing its hue. The surface will then look *brown*, providing you compare it to, say, the original orange, so you can tell that it is the amount of reflected light, not the illumination, that has changed. Thus, surprisingly, the sensation of lightness corresponds more closely to the surface's ability to reflect light than to properties of the light actually entering your eye from that surface. (The artist's term *value* is somewhat like lightness, but may also involve saturation.)

Colors may be arranged according to their hue, saturation, and lightness in a ***color tree.*** There are several different schemes, but they are all similar to those shown in Plate 9.1 and Figure 9.5. The "trunk" of the tree consists of the completely unsaturated colors—it ranges from black at the bottom, through grays, to white at the top. That is, the *height* is a measure of *lightness.* Out from the trunk, you find different hues in different directions—the *hue* varies *around* the tree.

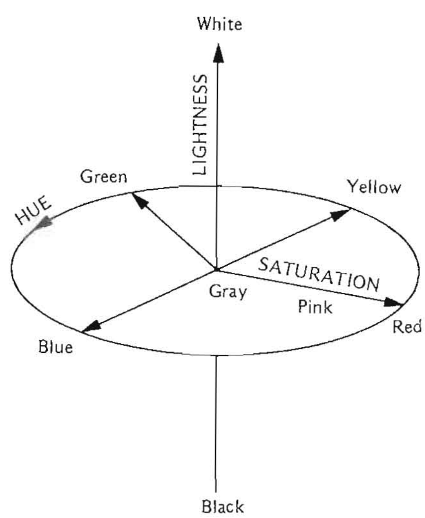

FIGURE 9.5

Schematic drawing of a color tree (compare with Plate 9.1).

Finally, as you move *away* from the trunk in some direction, the color becomes more *saturated.* For example, moving out in the red direction, you go from gray through pink, and only reach a saturated red at the tip of that "branch." Such standard color trees, or atlases, are of great importance to artists, paint manufacturers, printers, and anyone who must reproduce colors reliably, because any surfaces with the *same hue, saturation, and lightness* will *look the same color,* no matter what the intensity-distribution curves of their reflected light happen to be. Similarly, colored lights with the *same hue, saturation, and brightness* will *look the same color,* no matter what their intensity-distribution curves happen to be. As we've said, your eye cannot detect the wavelength distribution—rather it is sensitive to these three attributes. Since a given sensation can be produced by a variety of different stimuli, the color tree tells us very little directly about the physical processes that occur in color perception. We can get such information by exploring the various ways we can mix colors that appear identical to the eye.

TRY IT

FOR SECTION 9.3
Surface reflections and saturation

To see the effects of surface reflections on the saturation of colors, examine a shiny red apple or a piece of colored plastic in a room with several lights. Notice that no matter how intense the red of the apple, or how saturated the color of the plastic, the highlights (surface reflections of the lights) still appear white—completely unsaturated. View the apple outside on an overcast day, when there are no highlights. Does this affect the saturation? Crumple a piece of colored paper and compare the saturation of a part of the paper that is partially shaded with an adjacent part that is not.

9.4 COLOR MIXING BY ADDITION

One way to learn something about our perception of color is to start with *lights* for which we know both the intensity-distribution curves and the sensations that these lights produce, and then to ask what sensations are produced by *combinations* of these lights. In this section, we'll discuss combining two lights by shining them at the same spot on a ***white screen,*** a screen that diffusely reflects all visible wavelengths equally well. The light reflected to your eyes from that spot then contains both lights—it is an ***additive mixture;*** at every wavelength the intensities of the two lights add, so the combined intensity-distrubution curve is the *sum* of the two individual curves. (We'll refer to other techniques for mixing colors later.)

If we shine white light on the screen, that portion of the screen looks white to us. If, instead, we reflect the white light off a silver object and onto the screen, the screen again looks white, not silver. Had we used a copper object instead of the silver, the light on the screen would look orange; a fluorescent red object would reflect light to the screen that looks no different than ordinary red light, and so on. Thus, by concentrating on these lights, rather than on the objects themselves, we reduce the gamut of colors; but we're still left with the pure colors of the spectrum, washed-out colors, purples, and a host of other nonspectral colors, which remain for us to analyze.

A. The simple additive rules

Suppose the intensity-distribution curves of a blue, a green, and a red light look like those shown in Figures 9.6a, b, and c. What will the combination of green plus red look like? Adding the two curves gives the result shown in Figure 9.6d—a rather broad, flat-topped curve with no particular wavelength dominating, centered about λ = 575 to 600 nm, that is, in the yellow. We cannot yet tell by looking at this curve what the sensation produced by the corresponding light will be, but if you actually do the experiment you find that the resultant light does look *yellow.* This is a fairly strange result; in no way does yellow appear to be a mixture of green and red. Thus we see that your eye can interpret rather different intensity distributions as the same color—the mixture of Figure 9.6d looks very much like the monochromatic yellow of Figure 9.6g. Indeed, you don't have to use the broad colors of Figures 9.6b and c. If you mix roughly equal amounts of monochromatic green and monochromatic red (Fig. 9.6h), the result *also* looks yellow, even though the spectral yellow (λ = 580 nm) is completely absent.

What about other color combinations? Mixing blue and green lights gives a broad peak centered in the blue-green, a color called ***cyan**** (Fig. 9.6e). This is not too surprising, but again the result looks nearly the same whether we use broad or narrow blues and greens.

Mixing equal amounts of blue and red gives a double-humped distribution (Fig. 9.6f)—*no* one wavelength dominates. The sensation one gets is ***magenta,***† a kind of purple. (Again, we must face the problem of color names. In this book, we'll reserve the name violet for the shortest wavelength visible light, the name purple for combinations of short and long wavelength visible lights, and the name magenta for this particular purple.) Now this is something new; purples aren't in the spectrum at all—a new hue! Shakespeare tells us: "To . . . add another hue/Unto the rainbow . . ./Is wasteful and ridiculous excess," and yet your eye does just that, producing a purple not present in any rainbow.

*Greek *kuanos,* dark blue(!).

†Neither Greek nor Latin, but the name of a town in northern Italy.

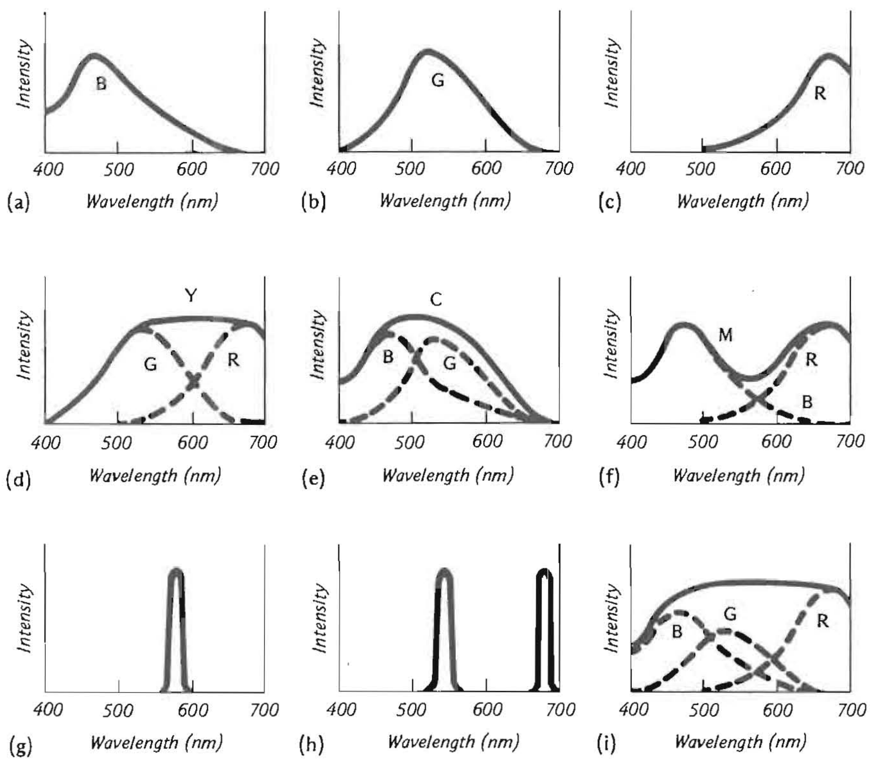

FIGURE 9.6

Additive color mixing. Intensity-distribution curves of **(a)** blue (B), **(b)** green (G), and **(c)** red (R) lights. Intensity-distribution curves of the additive mixtures (in equal amounts) of **(d)** G + R ≡ Yellow (Y), **(e)** B + G ≡ Cyan (C), and **(f)** B + R ≡ Magenta (M). Intensity-distribution curves of **(g)** monochromatic yellow and **(h)** a yellow made of an additive mixture of monochromatic green plus monochromatic red. **(i)** Intensity-distribution curves of the additive mixture of B + G + R ≡ White (W).

Finally, if you add all three colors (blue, green, and red) together in the proper proportions, you get the flat intensity-distribution curve shown in Figure 9.6i. This appears *white*, as we would expect from Figure 9.3a.

Figure 9.7 summarizes the rules we have developed so far for additive mixing of colored lights. The marvelous thing about these rules is that they are valid no matter what the intensity distributions of the three colors are. If two colors *look* the same, they will give identical results in an *additive* mixture with another color. This great simplicity, which allows us to forget about the details of the distribution curves, is what makes additive mixtures so appealing to physicists and is the reason that we can get away with talking about only three properties of a colored light.

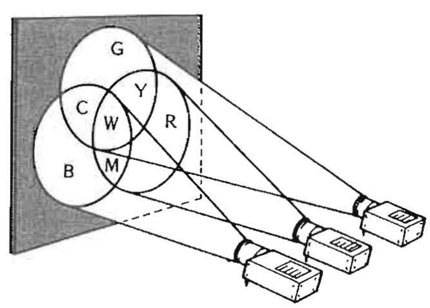

FIGURE 9.7

Simple additive mixing rules. The drawing shows three partially overlapping light beams, which combine additively. B = Blue, G = Green, R = Red, Y = Yellow, C = Cyan, M = Magenta, and W = White:

G + R ≡ Y
B + G ≡ C
B + R ≡ M
B + G + R ≡ W

B. Complementary colors

We can use the rules summarized in Figure 9.7 to discover some new additive combinations. We've seen that adding all three colors, blue, green, and red (with proper intensities), gives white:

B + G + R ≡ W*

But green and red together give yellow:

G + R ≡ Y

Hence, combining these two results:

B + Y ≡ W

blue and yellow lights added together give white. Two colors that, when added together, give white are called ***complementary colors.*** Thus *blue and yellow* are colors complementary to each other. Similarly, because we can think of cyan as being the same as blue and green mixed together:

C ≡ B + G

we can write:

C + R ≡ W

—*cyan and red* are complementary. Likewise:

M + G ≡ W

—*magenta and green* are complementary. (Note: these are not the complementary pairs your art teacher taught you.)

Again, it does not matter to your

*We follow the standard notation of the C.I.E. (*Commission Internationale de l'Eclairage*, the International Commission on Illumination) and use ≡ to mean "looks the same as," rather than = which would imply equality. Two colors that look the same need not be physically equal (see Fig. 9.8). These equations are valid only for the proper intensity balance of the colors being added.

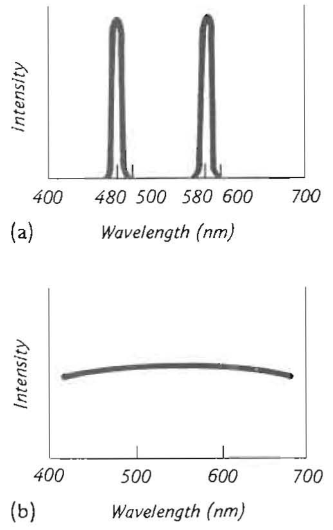

FIGURE 9.8

The lights with these two intensity-distribution curves look alike to your eye, even though one has only two wavelengths present while the other has all visible wavelengths.
(a) Monochromatic blue plus monochromatic yellow. **(b) Broad-band white** (all visible wavelengths).

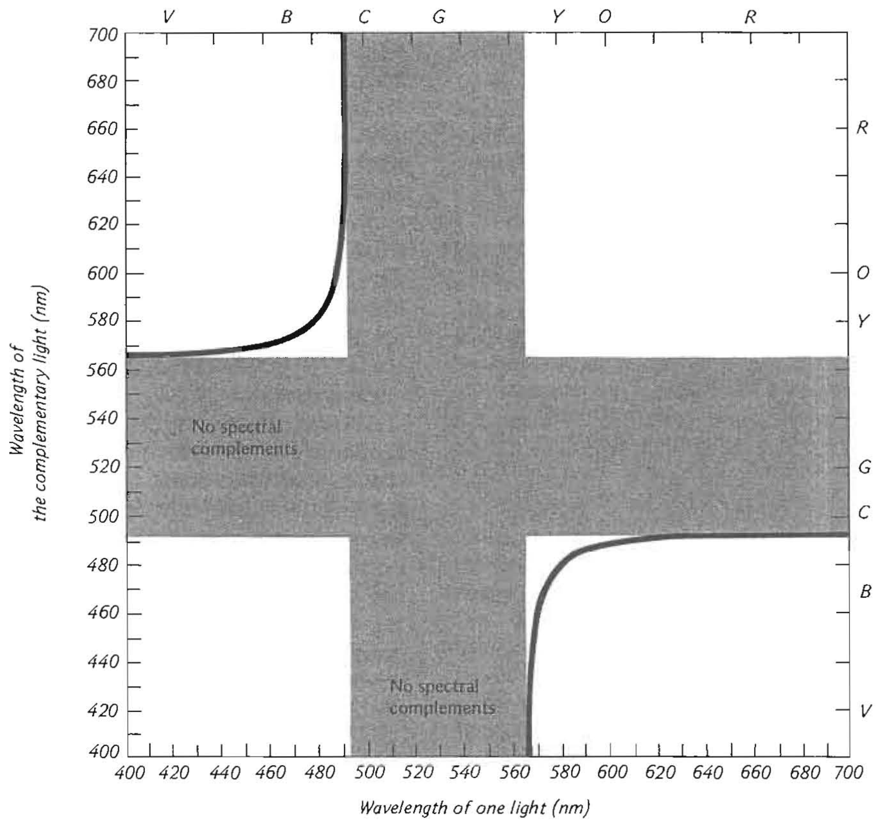

FIGURE 9.9

Wavelengths of complementary pairs of monochromatic colors. To find the complement of a given wavelength, say λ = 600 nm, draw a horizontal line from the 600-nm mark on the vertical axis. Find the point where this line intersects the curve and drop a vertical line from that point to the horizontal axis to determine the wavelength of the complement, λ = 489 nm. Thus, the complement of orange (600 nm) is bluish cyan (489 nm). (V = Violet, B = Blue, C = Cyan, G = Green, Y = Yellow, O = Orange, and R = Red.) The curves depend somewhat on the observer and on the choice of white.

eye whether the intensity distributions are broad or monochromatic; if two colors are complementary, their additive mixture is white. So if we add monochromatic yellow (λ = 580 nm) to monochromatic blue (λ = 480 nm), the result will look just as white as ordinary white light, which has all the wavelengths in it. That is, to your eye, the lights represented by the two intensity-distribution curves of Figure 9.8 look the *same!* Two such colors, which look alike even though they have different intensity-distribution curves, are called ***metamers.****

We can plot the combinations of monochromatic colors that are complementary (Fig. 9.9). Notice that there are no complementary monochromatic colors for the greens (roughly between 495 and 565 nm). This is because, as we already know, the complement of green is the nonspectral magenta—a double-humped intensity distribution, a combination of red and blue. This gives us a way to describe the nonspectral purple hues by wavelength, namely by the wavelength of their complements. For example, we can speak of a magenta 535c, meaning a color that gives white when added to 535-nm green.

The TRY IT uses negative afterimages to show you that complementary colors are intimately involved with the inner workings of your eye.

*Greek *meta*, with, plus *meros*, part.

TRY IT

FOR SECTION 9.4B
Complementary colors and negative afterimages

Pick out one point on Plate 5.2 and stare at it for a while (30 seconds), then look at a blank piece of white paper. Just as with Figure 7.16, you will see a negative afterimage. Here, however, because the plate you stared at is colored, the afterimage involves your color perception. When you stare at, say, the red region, your mechanism for detecting red (whatever it may be) becomes desensitized. When you then look at white light, you see less red in it with your desensitized red mechanism, so the light will look like white with red removed. Since white light can be considered as a mixture of cyan and red, when you remove the red from white the light looks cyan. Where you stared at red in the plate, the afterimage is cyan. Similarly, in every other region of the afterimage, you see the complement of the color in that region of the plate.

*This phenomenon is used by artists who place a fairly saturated color next to its complement to give **vibrating colors;** for example, when you stare at a* *boundary between red and cyan, your eye tremors cause the afterimage of the cyan to overlap the red, and vice versa (see Sec. 7.6). Since the afterimage of cyan is already red, when it is seen on the red background it enhances the saturation of the red. Thus you see intense flashes of color at the boundary. (Kelly uses green rather than cyan. Examine the red-green border in the plate carefully. The effect here is different from the chromatic aberration effect at the red-blue border, referred to in Sec. 5.2B. The eye has very little dispersion between red and green.) Larry Poons has painted huge canvases of one color, scattered throughout with dots of another color. If you are fortunate enough to see such a work, stare at one dot for a half-minute and then let your eye roam over the painting. The afterimages of the dots roam with your eye and provide an exciting interplay with the painted dots.*

Try producing the complementary afterimage in different circumstances. For example, place a small rectangular piece of colored paper on a white background, stare at a dot in the center of the colored paper for half a minute, then rotate the colored paper by 90°, and look at the dot again. The colored paper and its complementary afterimage form a cross. Notice that where they cross, the colored paper looks very desaturated.

C. *Chromaticity diagrams*

We have seen that different additive combinations of blue, green, and red lights produce yellow, cyan, magenta, and white. This suggests that we might be able to produce *any* color as an additive mixture of three properly chosen lights. A way of testing this idea may be thought of as a game. You choose three different colored lights and we pick an arbitrary fourth colored light. You then mix your three lights, adjusting their relative amounts, to see if you can make a mixture that matches our light in hue, saturation, and brightness. Suppose you pick three pure, monochromatic colors, one each in the blue, green, and red parts of the spectrum. Can you then match *any* color that we pick?

Well, almost any. Although this is your best choice, and there are many colors you *can* match, there are some that you can't. For instance, you can use your blue and green lights to make cyan, but if we choose a monochromatic cyan, you're in trouble. The best cyan you can make from blue and green lights will not be as saturated. You can make good matches to many unsaturated colors, but you cannot make perfect matches when we choose saturated colors. Strictly speaking, you've lost the game. However, if we change the rules slightly, you *can* win.

Although you can't match the monochromatic cyan perfectly, you can convert it to something you can match; if you add a little of your red to our monochromatic cyan, the cyan will become desaturated. Since red is complementary to cyan, mixing a little red with a little of the cyan gives a little white, which, when combined with the rest of the cyan, gives a desaturated cyan. You can then match the resulting desaturated cyan with a mixture of your blue and green lights. That is, using only your three lights, you have produced a match:

$$\text{blue} + \text{green} \equiv \text{monochromatic cyan} + \text{a little red}$$

If we treated this as a mathematical equation with an equals sign, we could rewrite it as:

$$\text{blue} + \text{green} - \text{a little red} \equiv \text{monochromatic cyan}$$

That is, if we allow you *negative* amounts, you can match our monochromatic cyan. In fact, if we allow negative contributions, you can match *any* color we choose, no matter how saturated—you can always win the game. (Remember, a "negative contribution" is just a shorthand way of saying that you mix some of one of *your* colors with *our* color and *then* match the resultant color with your other two colors.)

If you choose as your three colors a monochromatic red with $\lambda = 650$ nm, a monochromatic green with $\lambda = 530$ nm, and a monochromatic blue with $\lambda = 460$ nm, the relative amounts you will need to match any monochromatic color we choose are

FIGURE 9.10

The relative amounts of your three colors (460-nm blue, 530-nm green, and 650-nm red) needed to match any monochromatic (spectral) color that we choose. Notice that the required amounts of red and green colors are zero at 460 nm. This is because you can match our 460 nm using only your 460-nm blue. The relative amount of the blue, then, is 100% at that point. Similarly, the blue and red amounts vanish at 530 nm, while the blue and green vanish at 650 nm. (For historical reasons, "equal amounts of blue, green, and red" means that the intensity ratios of blue/green/red are about 1.3/1.0/1.8. The curves shown here and in succeeding figures are standardized; the actual data vary somewhat with observer, intensity of light, etc.)