21
Color

The chief topics discussed in this chapter are the nature of light, the nature of color, color and human vision, various color spaces (or models), additive and subtractive colors, complementary colors, and the CIE diagram.

21.1 Light

At a certain point in history, people started asking how the world is constructed and what are its constituents. For a while, it was widely believed that earth, fire, water, and air were the basic constituents of the world, but what about light? Light is very different from these constituents, It is fleeting. It is everywhere and nowhere. It cannot be grabbed, collected, saved, touched, felt, or smelled, and yet it definitely exists. It is something that makes it possible to see other objects.

Today, we know that light is indeed strange because it can be considered both as a wave and as a stream of particles. As a young man, Isaac Newton experimented with passing light through a prism and spreading a beam of white light into a rainbow of colors. He therefore assumed that light was a stream of particles (corpuscles). In the early 1800s, Thomas Young experimented with passing narrow beams of light through slits and observed interference patterns, similar to those of water waves, that convinced him that light must be a wave.

Modern science considers light as either an electromagnetic wave or as a stream of particles (termed photons) that have energy and momentum, but no mass.

What "waves" (or undulates) in light is the electric and magnetic fields. When a region of space is flooded with light, those fields vary periodically as we move from point to point in space. If we stay at one point, the fields also vary periodically with time. Thus, visible light is a (small) part of the electromagnetic spectrum (Figure 21.1) that

D. Salomon, *The Computer Graphics Manual*, Texts in Computer Science,
DOI 10.1007/978-0-85729-886-7_21, © Springer-Verlag London Limited 2011

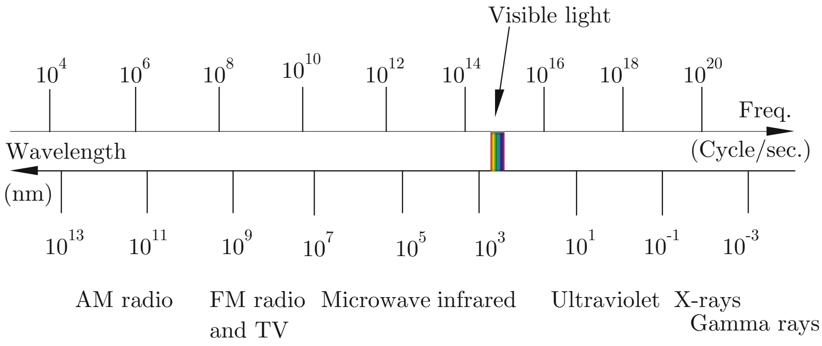

Figure 21.1: The Electromagnetic Spectrum.

includes radio waves, ultraviolet, infrared, X-rays, and other types of electromagnetic radiation.

Precisely how light features itself depends on how we look at it. When we perform an experiment that tests the wave nature of light, we see interference patterns or other features of waves, but when an experiment tests for the particle nature of light, we observe results consistent with light being a stream of photons. The most important attribute of a photon is its frequency, because the photon's energy is proportional to it. Photons are useful in physics to explain a multitude of phenomena (such as the interaction of matter and light).

In computer graphics, the most important property of light is its color, which is why the wave interpretation of light is used in this chapter (see also Section 26.3).

Light Makes Right.

21.2 Color and the Eye

The most important properties of a wave are its frequency f, its wavelength λ, and its speed. Light moves, obviously, at the speed of light (in vacuum, it is $c \approx 3 \times 10^{10}$cm/s). The three quantities are related by $f\lambda = c$. It is important to realize that the speed of light depends on the medium in which it moves (see discussion of refraction in Section 17.2.2). As light moves from vacuum to air to glass, it slows down (in glass, the speed of light is about $0.65c$). Its wavelength also decreases, but its frequency remains constant. Nevertheless, it is customary to relate colors to the wavelength and not to the frequency. Visible light has very short wavelengths, which is why a convenient unit for its wavelength is the nanometer (1 nm=10^{-9} m).

Visible light ranges from about 400 nm to about 700 nm and the color we observe is determined by the wavelength. A wavelength of 420 nm, for example, corresponds to pure violet, while 620 nm is perceived by the eye and brain as pure red. Using special lasers, it is possible to create pure (monochromatic) light consisting of one wavelength

(Figure 21.2a). Most light sources, however, output light that's a mixture of several (or even many) wavelengths, normally with one wavelength dominating (Figures 21.2b and 21.19).

The colors of the spectrum that are most visible to the human eye are (Figure 21.4) violet (390–430), blue-violet (460–480), cyan, green (490–530), yellow (550–580), orange (590–640), and red (650–800).

White light is a mixture of all wavelengths, but what is gray light? It turns out that the wavelength of light is not its only important attribute. The *intensity* is another attribute that should be considered. Gray light is a mixture of all wavelengths, but at a low intensity. When doing computer graphics, the main problem is how to specify the precise color of each pixel to the hardware. In many real-life situations, it is sufficient to say "I think I would like a navy blue suit," but computer hardware requires much more precise color specification. It therefore may be a surprise to discover that color can be completely specified by just three parameters. Their meanings depend on the particular *color model* (or color space) used. The RGB model is popular in computer graphics. In the printing industry, the CMYK is normally used. Many artists use the HLS model. These models are discussed here.

Figure 21.2b shows a simplified diagram of light smeared over the entire range of visible wavelengths, with a spike at about 620 nm (red), where it has a much higher intensity. This light can be described by specifying its *hue*, *saturation*, and *luminance*. The hue of the color is its dominant wavelength—620 nm in our example. We can think of the hue of a color as the position of the color in the rainbow. The luminance is related to the intensity of the light. It is defined as the total power included in the spectrum and it is proportional to the area under the curve, which is $L = (700 - 400)A + B(D - A)$. The saturation is defined as the percentage of the luminance that resides in the dominant wavelength. In our case, it is $B(D-A)/L$. When there is no dominant wavelength (i.e., when $D = A$ or $B = 0$), the saturation is zero and the light is white. Large saturation means either large $D - A$ or small L. In either case, there is less white in the color and we see more of the red hue. Large saturation therefore corresponds to pure color.

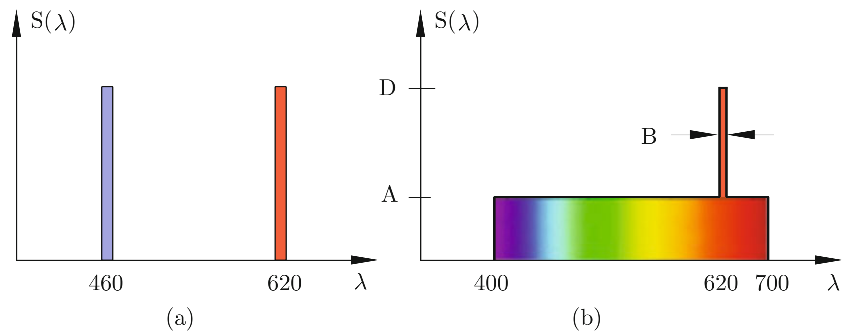

Figure 21.2: (a) Pure Colors. (b) A Dominant Wavelength.

21.3 Color and Human Vision

An object has intrinsic and extrinsic attributes. They are defined as follows:

- An intrinsic attribute is innate, inherent, inseparable from the thing itself and independent of the way the rest of the world is. Examples of intrinsic attributes are length, shape, mass, electrical conductivity, and rigidity.

- Extrinsic is any attribute that is not contained in or belonging to an object. Examples are the monetary value, esthetic merit, pleasing color, and usefulness of an object to us.

In complete darkness, we cannot see. With light around us, we see objects and they have colors. Is color an intrinsic or an extrinsic attribute of an object? The surprising answer is neither.

Stated another way, colors do not exist in nature. What does exist is light of different wavelengths and materials that absorb and reflect different wavelengths. When such light enters our eye, the light-sensitive cells in the retina (Figure 21.3) send signals that our brain interprets as color. Thus, colors exist only in our minds. This may sound strange, because we see colors all the time, but consider the problem of describing a color to another person. All we can say is something like “this object is red,” but a color-blind person has no idea of redness and is left uncomprehending. The reason we cannot describe colors is that they do not exist in nature. We cannot compare a color to any known attribute of the objects around us.

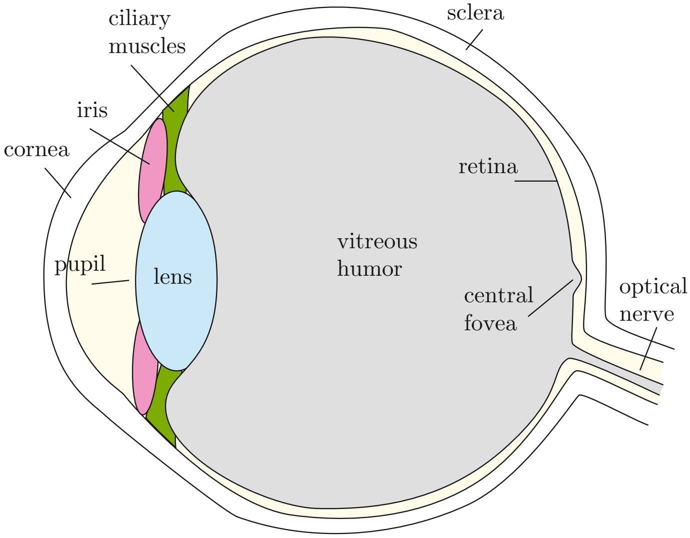

Figure 21.3: The Eye.

The retina is the back part of the eye. It contains the light-sensitive (photoreceptor) cells which enable us to sense light and color. There are two types of photoreceptors, rods and cones.

The rods are most sensitive to variations of light and dark, shape and movement. They contain one type of light-sensitive pigment and respond to even a few photons. They are therefore responsible for black and white vision and for our dark-adapted (scotopic) vision. In dim light we use mainly our rods, which is why we don't see colors under such conditions. There are about 120 million rods in the retina.

The cones are much less numerous. There are only about 6–7 million of them and they are concentrated at the center of the retina, the fovea (or yellow spot, see Figure 21.3 and Page 875) and are much less sensitive to light (it takes hundreds of photons for a cone to respond). On the other hand, the cones are sensitive to wavelength. Thus, it is the cones that send color information to the brain and that are responsible for our color vision (more accurately, the cones are responsible for photopic vision, vision of the eye under well-lit conditions). There are three types of cones and they send three different types of stimuli to the brain. This is why any color specification requires three numbers (we perceive a three-dimensional color space). The three types of cones are sensitive to red, green, and blue, which is why these colors are a natural choice for primary colors. (More accurately, the three types of cones feature maximum sensitivity at wavelengths of about 420 nm (blue), 534 nm (Bluish-Green), and 564 nm (Yellowish-Green).)

(Why do humans and other mammals have so many more rods than cones? Current theory maintains that mammals spent the first part of their evolutionary history as nocturnal animals, to avoid dinosaurs, so natural selection did not favor mammals with many cones (sensitive color vision). Birds, on the other hand, have many cones because they had nothing to fear from dinosaurs and have always been diurnal.)

The importance of red, green, and blue as primary colors was suspected long before anyone knew about the rods, cones, and their sensitivities. As early as 1801, the polymath Thomas Young developed a trichromatic theory of color where he argued that red, green, and blue are three components of any color. (It is currently believed that he was preceded by George Palmer who had similar ideas around 1786.) This theory was later extended by Helmholtz. The Young–Helmholtz theory of color vision states that there are three types of receptors in the retina and they are responsible for the perception of color. These receptors are sensitive to red, green, and blue.

In addition to scotopic and photopic visions, there is also mesopic vision. This type of vision is a combination of photopic and scotopic visions in low light but not full darkness.

Nothing in our world is perfect, and this includes people and other living beings. Color blindness in humans is not rare. It strikes about 8% of all males and about 0.5% of all females. It is caused by having just one or two types of cones (the other types may be completely missing or may just be weak). A color blind person with two types of cones can distinguish a limited range of colors, a range that requires only two numbers to specify a color. A person with only one type of cone senses even fewer colors and perceives a one-dimensional color space (grayscale). With this in mind, can there be persons, animals, or aliens with four or more types of cones in their retina? The answer is yes, and such beings would perceive color spaces of more than three dimensions and would sense more colors than we do! (Naturally, we cannot imagine what those colors are.) This surprising conclusion stems from the fact that colors are not attributes of objects and exist only in our minds.

Note. In living beings with many types of cones, certain types of cones in the retina may be responsible for more perceived colors, but may also be specialized. Cones of a certain type may widen the visible spectrum to include ultra-violet and infrared. There may be specialized cones that help their owner to detect motion or that act as sun glasses by filtering out certain light frequencies. Thus, not every additional type of cone will increase the number of dimensions of the color space perceived by a creature.

> I can see it clearly. I don't know what it is. I can describe it. Up to a point. A lot of the colors don't have a name.
>
> —A. S. Byatt, *Angels and Insects* (1986).

The famous "blind spot" also deserves mentioning at this point. The blind spot is a point where the retina does not have any photoreceptors. Any image in this region is not detected by the eye and is invisible. The blind spot is the point where the optic nerves come together and exit the eye on their way to the brain. To find your blind spot, look at this image,

close your left eye, and hold the page about 20 inches (50 cm) away. Look at the square with your right eye. Slowly bring the page closer while looking at the square. At a certain distance, the cross will disappear. This occurs when the cross falls on the blind spot of your retina. Reverse the process. Close your right eye and look at the cross with your left eye. Bring the image slowly closer to you and the square will disappear.

Each of the photoreceptors sends a light sensation to the brain that's essentially a pixel, and the brain combines these pixels into a continuous image. Thus, the human eye is similar to a digital camera. Once this is understood, we naturally want to compare the resolution of the eye to that of a modern digital camera. Current (2010) digital cameras feature from 3–4 Mpixels (for a cell phone camera) to about 12 Mpixels (for a good-quality camera).

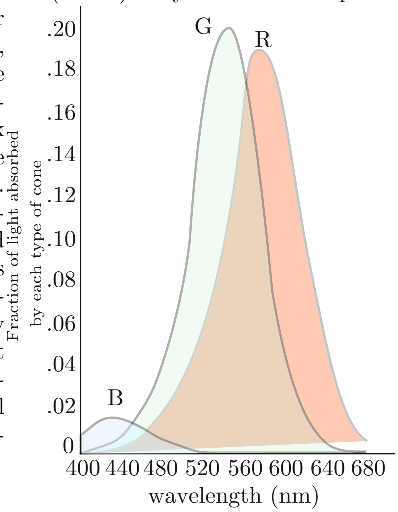

Figure 21.4: Sensitivity of the Cones.

Thus, the eye features a much higher resolution, but its effective resolution is even higher if we consider that the eye can move and refocus itself about three to four times a second. This means that in a single second, the eye can sense and send to the brain about half a billion pixels. Assuming that our camera takes a snapshot once a second, the ratio of the resolutions is about 100.

Certain colors—such as red, orange, and yellow—are psychologically associated with heat. They are considered *warm* and cause a picture to appear larger and closer than it actually is. Other colors—most notably blue, violet, and green—are associated with cool things (air, sky, water, ice) and are therefore called *cool* colors. They cause a picture to appear smaller and farther away.

Many psychologists, artists, and designers agree that certain colors create or enhance emotions in people. The following table may raise some objections, but is generally considered true.

Black	Classy, serious, dramatic
Blue	Secure, loyal, comfortable
Brown	Older, natural
Gray	Distinctive, business-like, cold
Green	Nature, food, healthy, money
Orange	Warm, energy
Pink	Babyish, health, soft
Purple	Sophisticated, royal
Red	Strong, aggressive, heavy
White	Pure, simple, clean
Yellow	Careful, bright

21.3.1 Color Temperature

We see an object because of light emitted from it or reflected by it that reaches our eye. An object is perceived as blue because it reflects blue light and absorbs (or transmits) all the other colors. Similarly, an object is black because it absorbs every color (all the wavelengths of electromagnetic radiation) that falls on it. However, an object cannot simply absorb energy without limit, so any object also emits energy, in the form of electromagnetic radiation.

An ideal black object (termed black body) absorbs all the wavelengths, but it also emits radiation. Normally, we don't see this radiation because at room temperature it is in the infrared range, but if the black body is heated to higher temperatures, it emits radiation of shorter wavelengths and we see it first red, then orange, yellow, white, and finally blue (which then turns to ultra violet). Notice that colors that we consider warm (such as red and orange) have low temperatures, while cool colors (blue and white) have high temperatures.

Thus, there is a simple relation between color and temperature. We say that the temperature of a given color C is T° Kelvin if a black body appears to have color C when heated to temperature T. This relation is illustrated by Figure 21.5.

The concept of color temperature has important applications in fields such as lighting, photography, videography, publishing, manufacturing, and astrophysics.

The light emitted by an incandescent bulb has a typical color temperature of 2,700 K, and is somewhat yellowish. The new, compact fluorescent light bulbs have higher color temperatures. Color temperatures of 3,000–3,500 K produce a neutral white light, while higher temperatures (above 4,000 K) result in artificial light that is very close to daylight.

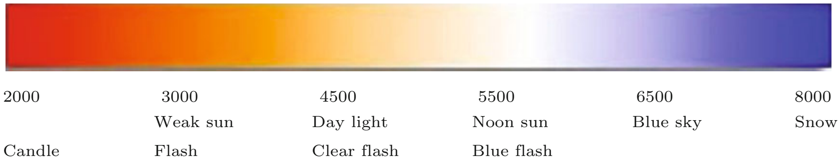

Figure 21.5: Scale of Color Temperatures.

21.4 The HLS Color Model

This model was introduced in 1978 by Tektronix, aiming for an intuitive way to specify colors. The term HLS stands for hue, lightness, and saturation. Lightness (or value) refers to the amount of black in the color. It controls the brightness of the color. Maximum lightness always creates white, regardless of the hue. Minimum lightness results in black. Saturation (or chroma) refers to the amount of white in the color. It controls the purity or vividness of the color. Low saturation means more white in the color, resulting in a pastel color. Very low saturation results in a washed-out color. For a pure, vivid color, the saturation should be maximum. The achromatic colors black, white, and gray have zero saturation and differ in their values (Figure 21.7).

The HLS model is summarized in the map of Figure 21.6 and in the double cone of Figure 21.8. The vertical axis corresponds to L (lightness). It starts at zero (black) at the bottom and ends at 1 (white) at the top. The distance from the central axis corresponds to S (saturation). All points on the axis have zero saturation, so they correspond to shades of gray. Points farther away from the axis have more saturation; they correspond to more vivid colors. The H parameter (hue) corresponds to the hue of the color. This parameter is measured as an angle of rotation around the hexagon.

Figure 21.6: The HLS Map.

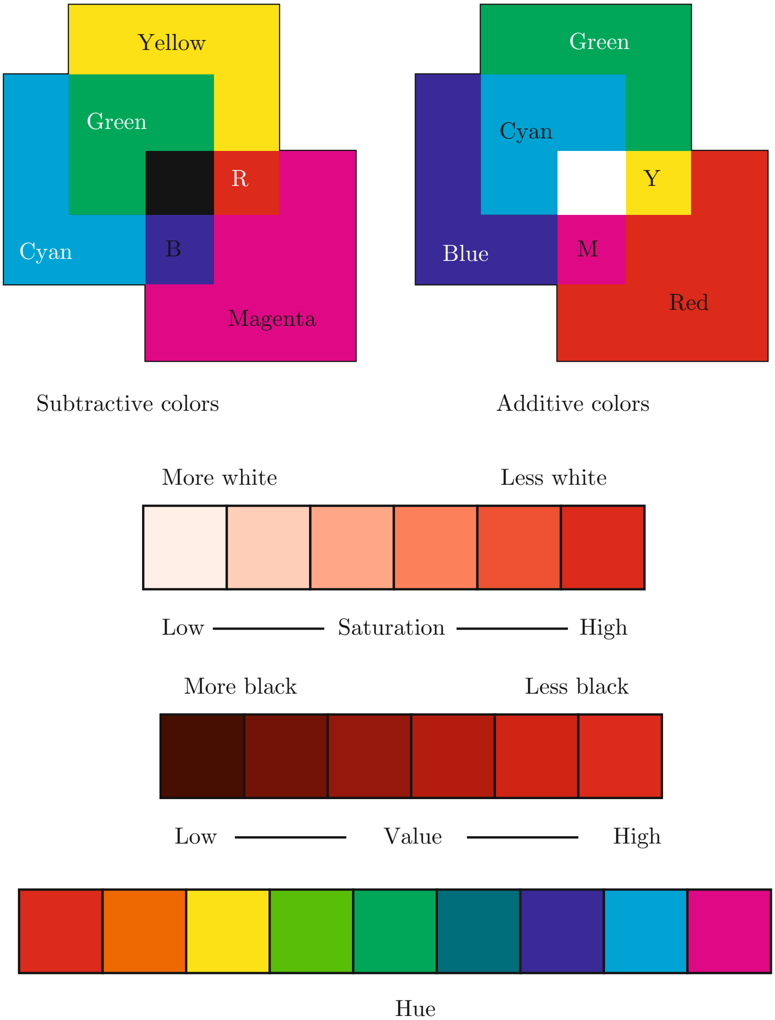

Figure 21.7: Examples in Color.

21.5 The HSV Color Model

The HSV model also uses hue, saturation, and value (lightness). It is summarized in the cone of Figure 21.9. This is a single cone where the value V, which corresponds to lightness, goes from 0 (black) at the bottom, to 1 (white) at the flat top. The S and H parameters have the same meanings as in the double HLS cone.

21.6 The RGB Color Space

A *primary hue* is a color in a color model that cannot be made from the other colors included in that model. Primary hues serve as a basis for mixing and creating all other colors in the color model. Any color created by mixing *two* primary hues in a color model is a *secondary hue* in that model.

In the RGB color model, the three primaries are red, green, and blue. They can be combined, two at a time (Figure 21.7) to create the secondary hues. Magenta (pinkish hue) is (R + B), cyan (bluish hue) is (B + G), and yellow is (R + G). There are two reasons for using the red, green, and blue colors as primaries: (1) the cones in the eye are very sensitive to these colors (Figure 21.4) and (2) adding different amounts of red, green, and blue can produce many colors (although not all colors, see discussion of RGB color gamut on Page 999).

The RGB color model is useful in computer graphics because of the way color displays work. CRTs create different colors by light emitted from phosphors of different types, whereas LCDs create colors from LCDs covered by colored filters (Section 26.3). The colors are then mixed in the eye of the observer, creating the impression of a perfect mixture. Assuming a range of $[0, 255]$ for each RGB color component, here are some examples of mixed colors:

$$\text{red} = (255, 0, 0),\ \text{magenta} = (255, 0, 255),\ \text{white} = (255, 255, 255),$$
$$50\%\ \text{gray} = (127, 127, 127),\ \text{light gray} = (25, 25, 25).$$

> It's the weird color-scheme that freaks me. Every time you try to operate one of these weird black controls, which are labeled in black on a black background, a small black light lights up black to let you know you've done it!
>
> —Mark Wing-Davey (as Zaphod Beeblebrox) in *The Hitchhiker's Guide to the Galaxy* (1981).

21.6.1 The RGB Cube

The *color gamut* of a color model is the entire range of colors that can be produced by the model. The color gamut of the RGB model can be summarized in a diagram shaped like a cube. Figure 21.10 shows the three fully saturated faces and a small projection of the three other faces of this cube (see also Figure 21.12). Every point in the cube has three coordinates (r, g, b)—each in the range $[0, 1]$ (in the figure, the ranges are 0–255, corresponding to eight bits per color component)—which give the intensities of

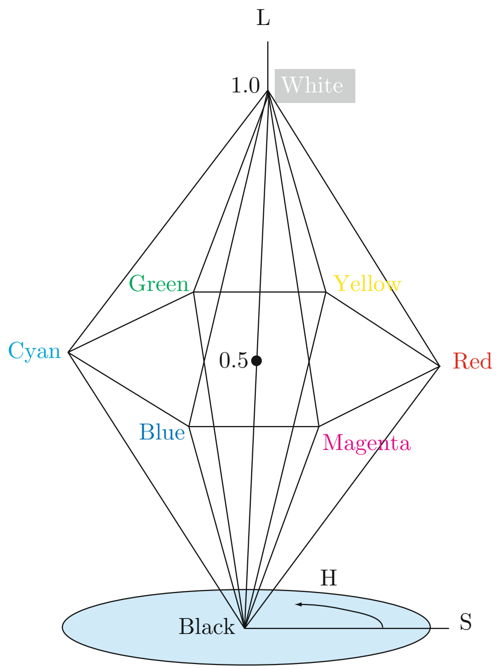

Figure 21.8: The HLS Double Hexcone.

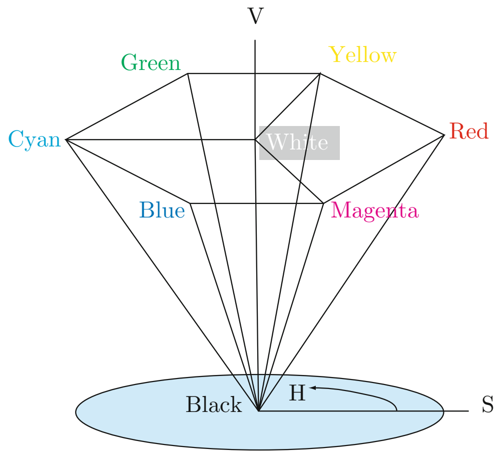

Figure 21.9: The HSV Hexcone.

red, green, and blue of the point. Small values, close to $(0,0,0)$, correspond to a dark shade, whereas anything close to $(1,1,1)$ is very bright. Point $(1,1,1)$ itself corresponds to pure white. Point $(1,0,0)$ corresponds to red and point $(0,1,0)$ corresponds to green. Therefore, point $(1,1,0)$ describes a mixture of red and green, that is, yellow. Notice that the 256 levels of a primary color often do not represent equally-spaced intensities, because of gamma correction (Section 26.4.4).

Reference [colorcube 10] is a demonstration of the RGB color cube.

◇ **Exercise 21.1:** What are the RGB coordinates of 50% gray and where is it located in the RGB cube?

The RGB cube is useful because coordinates of points in it can readily be translated into values stored in the color lookup table of the computer.

◇ **Exercise 21.2:** (Easy.) What colors correspond to the diagonal line connecting the black and white corners of the RGB cube (this line is known as the neutral axis of the cube)?

21.7 Additive and Subtractive Colors

Colors can be mixed by adding or subtracting them and these mixtures appear to us as new colors. Imagine a white wall in a dark room. There is no light for the wall to reflect, so it looks black. We now shine red light on it. Since the wall is white (i.e., it reflects all colors), it will reflect the red light and will look red. The same is true for green light. If we now shine both red and green lights on the wall, it will reflect both, which our brain interprets as yellow. We say that in this case the colors are added.

Notice that the wall does not reflect yellow light. It reflects red and green lights, and this mixture creates the sensation of yellow in our brain.

To understand the concept of subtracting colors, imagine a white sheet of paper in a bright environment. The paper reflects all colors, so it looks white. If we want to paint a certain spot red, we have to cover it with a chemical (red paint) that absorbs all colors except red. We say that the red paint *subtracts* the green and blue from the original white reflection, so the spot now reflects just red light. Similarly, if we want a yellow spot, we have to use yellow paint, which is a substance that absorbs blue and reflects red and green.

We conclude that in the case where we shine white light on a reflecting surface, we have to subtract colors in order to get the precise color we want. In the case where we shine light of several colors on such a surface, we have to add colors to get any desired mixture (Figure 21.7).

> The various colours that may be obtained by the mixture of other colours, are innumerable. I only propose here to give the best and simplest modes of preparing those which are required for use. Compound colours, formed by the union of only two colours, are called by painters virgin tints. The smaller the number of colours of which any compound colour is composed, the purer and the richer it will be. They are prepared as follows: ...
>
> —Daniel Young, *Young's Translation of Scientific Secrets* (1861).

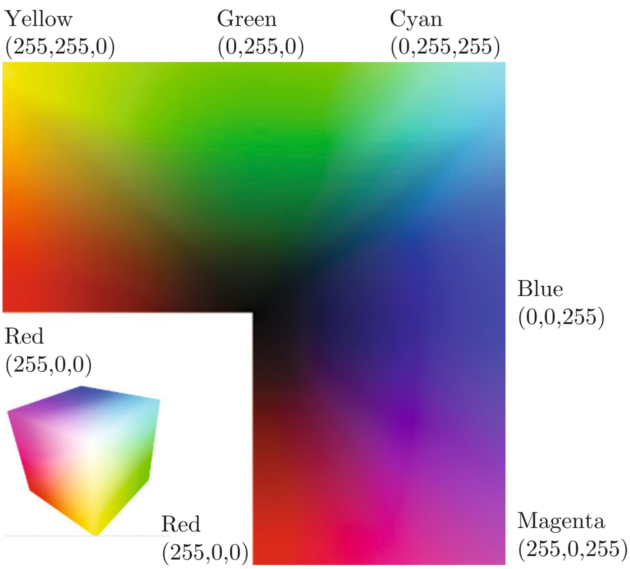

Figure 21.10: The RGB Cube.

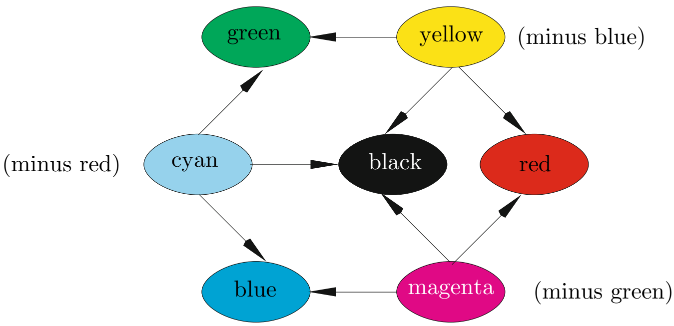

Figure 21.11: RGB and CMYK Relationships.

> For example, the human eye and its controlling software implicitly embody the false theory that yellow light consists of a mixture of red and green light (in the sense that yellow light gives us the same sensation as a mixture of red light and green light does). In reality, all three types of light have different frequencies and cannot be created by mixing light of other frequencies. The fact that a mixture of red and green light appears to us to be yellow light has nothing whatever to do with the properties of light, but is a property of our eyes. It is a result of a design compromise that occurred at some time during our distant ancestors' evolution.
>
> —David Deutsch, *The Fabric Of Reality*.

21.7.1 Subtractive Color Models

There are two subtractive color models, painter's pigments and printing pigments.

Painter's Pigments. The primary colors of this color model are red, yellow, and blue. They were chosen because artists in the past believed that they were pure colors, containing no traces of any other colors. These three primaries can be mixed, two at a time (see Section 21.9), to produce the secondaries purple (R + B), green (B + Y), and orange (Y + R). Mixing equal amounts of all three primaries subtracts all colors and hence produces black.

Printing Pigments. This color model is also known as *process color* and is the result of development in color ink and printing processes. The three primaries are magenta, yellow, and cyan. The three secondaries are blue (M + C), red (M + Y), and green (C + Y). Mixing equal amounts of all three primaries should yield black, but because of the properties of real inks this black is normally not dark enough. In practice, true black is included in this model as an artificial fourth primary (also because black ink is cheaper). It is used when grayscale or black printing is required. The model is therefore called CMYK (K for black, to avoid confusion with blue) and color printing is known as the four-color process. Figure 21.11 shows the relationships between the three CMY primaries and their secondaries. Figure 21.14 shows examples of CMYK colors and Figure 21.15 shows the CMY components of a complex pattern (the black component, K, is not shown because it is very weak).

Because of the particular primaries and secondaries of the CMY model there is a simple relationship between it and the RGB model. The relation is

$$(r, g, b) = (1, 1, 1) - (c, m, y).$$

This relationship shows that, for example, increasing the amount of cyan in a color, reduces the amount of red.

Traditional color printing uses color separation. The first step is to photograph the original image through different color filters. Each filter separates a primary color from the multicolored original. A blue filter separates the yellow parts of the original and creates a transparency with those parts printed in grayscale. A red filter separates the cyan parts and a green filter separates the magenta parts. Another transparency is prepared, with the black parts. Each of the four transparencies is then converted to a halftone image (Section 2.27) and the four images become masters for the final printing (Figure 21.12). They are placed in different stages of the printing machine and, as the paper moves through the machine, each stage adds halftone dots of colored ink to the paper. The result is a picture made of four halftone grids, each in one of the CMYK colors. The grids are not superimposed on the paper but are printed offset. The eye sees dots colored in the four primaries, and the brain creates a mixed color that depends on the number of halftone dots of each primary.

When such a color print is held close to the eye, the individual dots in the four colors can be seen (Figure 21.13). There are dye sublimation printers that mix wax of different dyes inside the printer to create a drop of wax of the right color which is then deposited on the paper. No halftoning is used and the printing is continuous, in contrast with inkjet printers, which spray small dots of ink. The result is a picture in vivid colors. However, current inkjet printers have become so good, that they have replaced most dye

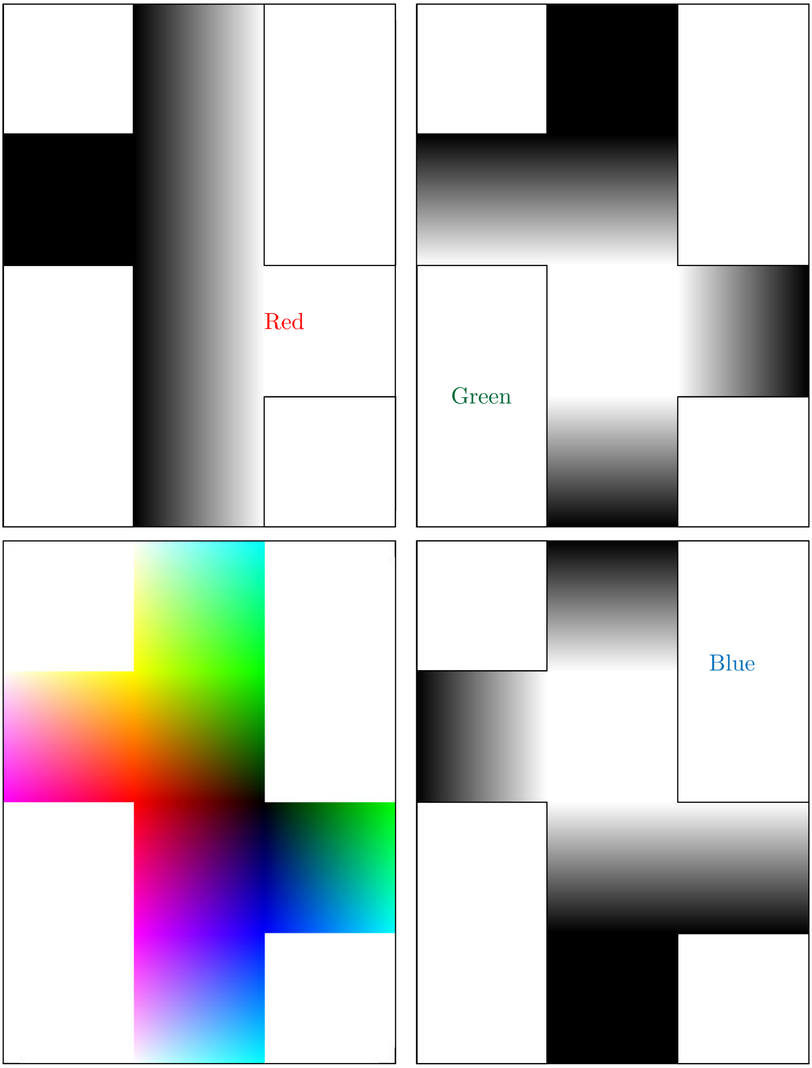

Figure 21.12: Color Separation of the RGB Cube.

In Tati's view, the varied answers proved that color was not part of what people see unless it has some function or meaning. We recall significant colors, but for the rest, our memories are mostly monochrome.

—David Bellos, *Jacques Tati*, (1999).

sublimation printers. For more information on color printing and the CMYK model, see [Stone et al. 88] and [wiki-color-print 10].

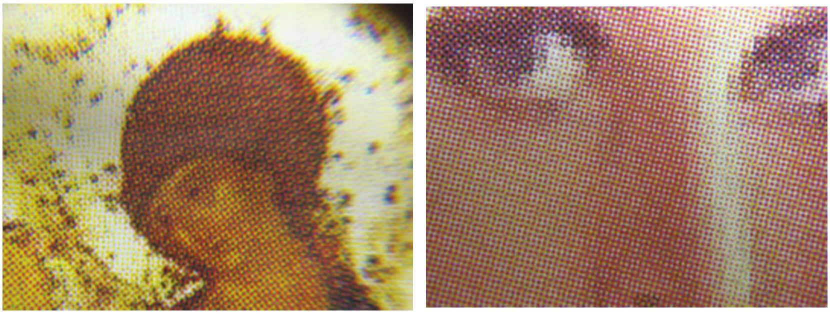

Figure 21.13: Halftone Color Dots.

In software for drawing and illustration, each of the four CMYK components is normally specified as an integer in the interval $[0, 100]$. In principle, any combination of four such integers is a valid color, but experts claim that it is important to keep the sum of the four integers (which can be up to 400) below about 280, because higher values require so much ink that the paper cannot fully absorb it and the print may later smear on the page. This point is especially important when printing text and figures in rich black. This color, which in principle is obtained by setting CMYK $= (100, 100, 100, 100)$, is better specified as $(50, 50, 50, 100)$.

Several color standards have been developed in order to simplify the task of editors and graphics designers. Instead of figuring out the ratios of CMYK, the graphics designer browses a table that has many color samples, selects one, and uses its name to specify it to the printer. One such standard in common use today is the PANTONE matching system. It is described in [Pantone 91].

◇ **Exercise 21.3:** A surface has a certain color because of its ability to absorb and reflect light. A surface that absorbs most of the light frequencies appears dark; a surface that reflects most frequencies appears bright. What colors are absorbed and what are reflected by a yellow surface?

Blueness doth express trueness.
—Ben Jonson.

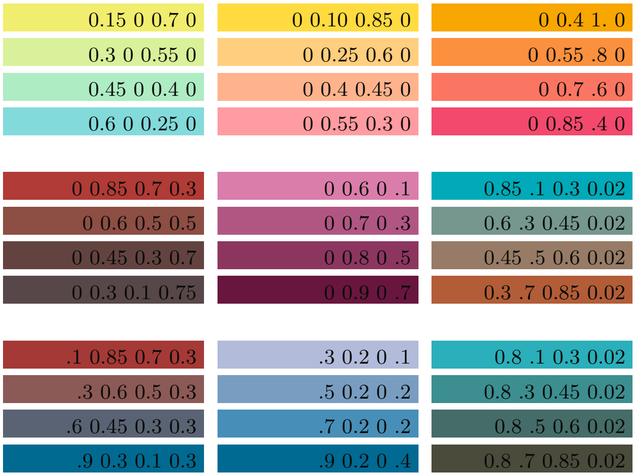

Figure 21.14: Examples of CMYK Colors.

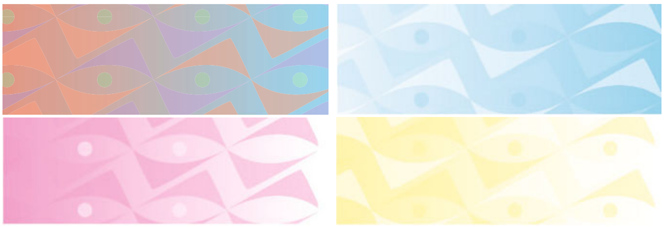

Figure 21.15: CMY Components of a Pattern.

21.8 Complementary Colors

The concept of complementary colors is based on the idea that two colors appear psychologically harmonious if their mixture produces white. Imagine the entire color spectrum. The sum of all the colors produces white. If we subtract one color, say, blue, the sum of the remaining colors produces the complementary color, yellow. Blue and yellow are thus complementary colors (a dyad) in an additive color model. Other dyads are green and magenta, red and cyan, yellow-orange and cyan-blue, cyan-green and red-magenta, and yellow-green and blue-violet.

Subtractive complementary colors are based on the idea that two colors look harmonious if their mixture yields a shade of gray. The subtractive dyads are yellow and violet, red and green, blue and orange, yellow-orange and blue-violet, blue-green and red-orange, and yellow-green and red-violet.

Complementary colors produce strong visual contrast, which creates a feeling of color vibrations or activity.

◇ **Exercise 21.4:** Is there such a thing as additive color triads?

When a color is represented in HSV or HLS (or in any other color space where hue is one of the components), then computing the complementary of a given color is easy; simply complement the hue. If the hue is in the interval $[0, 360]$, then compute $\text{hue} = 360 - \text{hue}$. If a color is given in another color space, then the complementary is computed by converting the color to HSV, complementing H, and converting back.

The well-known Adobe Illustrator program employs a shortcut to compute the complementary of an RGB color. It adds the lowest and highest RGB values of a given color, and then subtracts each component from the sum to obtain the new RGB components. Thus, given $\text{RGB} = (50, 75, 125)$, first compute the sum $50 + 125 = 175$, and then subtract $175 - 50 = 125$, $175 - 75 = 100$, and $175 - 125 = 50$.

Color

Colors, like features, follow the changes of the emotions.
—Pablo Picasso.

There is no blue without yellow and without orange.
—Vincent Van Gogh.

21.9 The Color Wheel

A color wheel is a simple model designed to help in selecting and matching colors. Figure 21.16 shows how to construct the most common color wheel. It consists of twelve hues that are classified into three basic colors and colors derived from those.

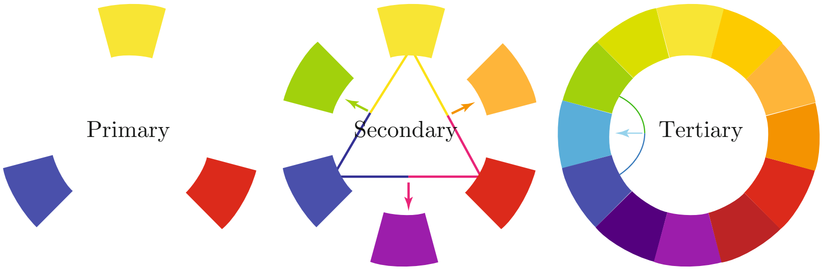

Figure 21.16: Color Wheel Construction.

To construct such a wheel, start with three primaries (in the figure, those are the painter's pigments red, blue, and yellow). Mix each of the three pairs of primaries to obtain a secondary color (in the figure, those are green, orange, and purple). Finally, generate the tertiary colors by mixing each primary with its two near neighbors, one at a time. Thus, for example, mixing green (secondary) and yellow (primary) creates a tertiary hue that we can term yellow-green.

Once the wheel is ready, we can use it in several ways for matching colors and hues for a color scheme. Figure 21.17 shows how the 12 hues of the wheel can be partitioned into warm and cool colors. The former set consists of hues from red-purple to yellow, while the latter set is made of yellow-green to indigo. The figure also illustrates how easy it is to determine the complementary of a given hue. It is simply the hue on the other side of the wheel (thus, the complementary of yellow-green is red-purple).

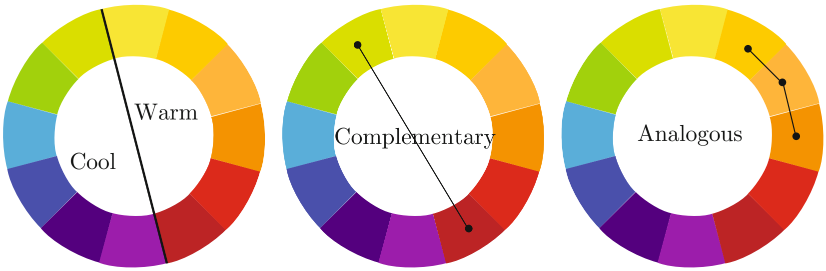

Figure 21.17: Color Wheel Properties.

Finally, Figure 21.17 demonstrates how the color wheel makes it easy to determine the analogous to any given hue. Analogous hues are those that look good to us when they are used together, in the same environment (this is an example of extrinsic attribute). Thus, a designer having to decide on a color scheme for, say, a room, may decide on orange as a dominant hue, and may also include orange-yellow and red-orange because they are analogous.

In addition to analogous hues, the neutral colors (white, black, and shades of gray), blend nicely with many color schemes. Architects, interior designers, and clothes designers generally recommend using only warm or only cool hues in a color scheme, and to limit the number of hues.

Some color experts also feel that complementary colors enhance each other because of the big difference in the color sensation that they cause, and often look good together for this reason. If this belief is true, then it may be the reason why green and red are the dominant hues in many national flags. A color scheme based on two sets of complementary hues is known as a tetradic.

> They always told us, didn't they, the teachers and grans, orange and pink, they make you blink, blue and green should not be seen, mauve and red cannot be wed, but I say, there're all there, the colours, God made them all, and mixes them all in His creatures, what exists goes together somehow or other, don't you think, Mrs Dennison?
>
> —A. S. Byatt, *The Matisse Stories*, (1993).

Figure 21.18 illustrates two more features, namely split complementary and triadic, that can easily be determined from the color wheel. A split complementary color scheme uses a base hue plus the two hues that are the immediate neighbors of its complementary (yellow and red-purple–blue-purple in the figure). Such a color scheme is supposed to offer the contrast of complementary colors while also avoiding the intensity of the difference between them.

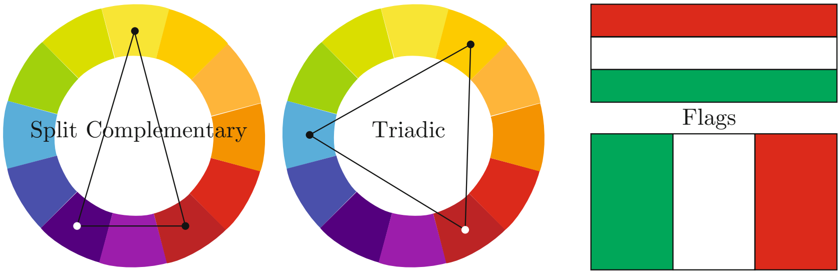

Figure 21.18: Color Wheel Properties and Two Flags.

A triadic color scheme is based on three equally-spaced hues in the color wheel (cyan, orange-yellow, and red-purple in the figure). Such a color scheme produces attractive results especially if it uses one hue as dominant and the other two as supplementary hues, to accent differences between objects.

There are many tools on the Internet for experimenting with hues, colors, and color schemes. One that is especially easy to use is [Kuler 11].

21.10 Spectral Density

A laser is capable of emitting "pure" light, light that consists of a single wavelength. Most light sources, however, emit "dirty" light that's a mixture of many wavelengths, normally with one dominating. For each light source, the graph of light intensity as a function of the wavelength λ is called the *spectral density* of the light source. Figure 21.19 shows the spectral densities of several typical light sources.

These simple diagrams illustrate one problem in attempting to specify color systematically and unambiguously. Several different spectral densities may be perceived by us as identical. When the colors created by these spectral densities are placed side by side, we find it impossible to distinguish between them. The first step in solving the problem is color matching. Suppose that we use a color model defined by the three primaries $A(\lambda)$, $B(\lambda)$, and $C(\lambda)$ and we have a color described by the spectral density $S(\lambda)$. How can we express $S(\lambda)$ in terms of our three primaries? One way is to shine a spot of $S(\lambda)$ on a white screen and, right next to it, a spot of light $P(\lambda) = \alpha A(\lambda) + \beta B(\lambda) + \gamma C(\lambda)$ created by mixing the three primaries (where $0 \leq \alpha, \beta, \gamma \leq 1$). Now vary the amounts of α, β, and γ until a trained observer judges the spots to be indistinguishable. We can now say that, in some sense, $S(\lambda)$ and $P(\lambda)$ are identical, and write $S = P$.

In what sense is the preceding true? It turns out that the above statement is meaningful because of a remarkable property of colors. Suppose that two spectral densities

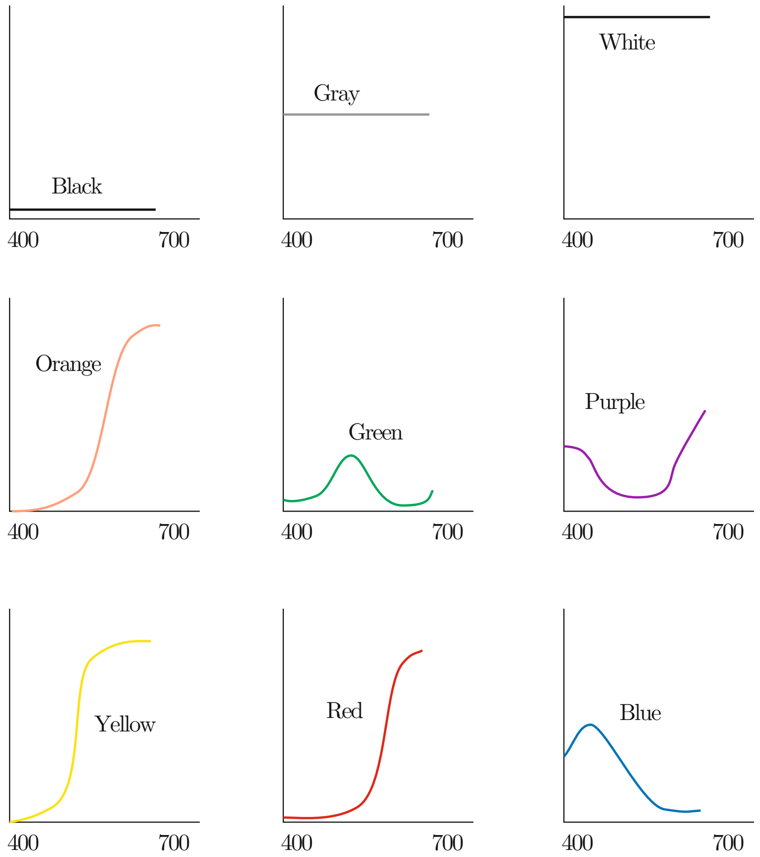

Figure 21.19: Some Spectral Densities.

$S(\lambda)$ and $P(\lambda)$ have the same perceived color, so we write $S = P$. We now select another color Q and shine it on both spots S and P. We know from experience that the two new spots would also be indistinguishable. This means that we can use the symbol $+$ for adding lights and we can describe the two spots by $S(\lambda) + Q(\lambda)$ and $P(\lambda) + Q(\lambda)$. In short, we can say "if $S = P$, then $S + Q = P + Q$." The same is true for changing intensities. If $S = P$, then $\alpha S = \alpha P$ for any intensity α. We therefore end up with a vector algebra for colors, where a color can be treated as a three-dimensional vector, with the usual vector operations (Section 8.1 and Appendix A).

Given the above, we can select a color model based on three primaries, A, B, and C, and can represent any color S as a linear combination of the primaries. Thus, $S = \alpha A + \beta B + \gamma C$. We can say that the vector (α, β, γ) is the representation of S in the basis (A, B, C). Equivalently, we can say that S is represented as the point (α, β, γ)

in the three-dimensional space defined by the vectors $A = (1, 0, 0)$, $B = (0, 1, 0)$, and $C = (0, 0, 1)$.

Three-dimensional graphs are difficult to draw on paper, so we would like to artificially reduce the representation from three dimensions to two. This is done by realizing that the vector $(2\alpha, 2\beta, 2\gamma)$ represents the same color as (α, β, γ) but is twice as bright. We therefore restrict ourselves to vectors (α, β, γ), where $\alpha + \beta + \gamma = 1$. These are vectors normalized to unit brightness. All vectors of unit brightness lie in the $\alpha + \beta + \gamma = 1$ plane (see Section 9.2.2 for the equation of a plane) and it is this two-dimensional plane that we plot on paper. Any point on this plane, i.e., any color, can be specified with three numbers α, β, and $\gamma = 1 - \alpha - \beta$, only two of which are independent.

For the RGB color model, we now select the pure spectral colors using trained human experts. The idea is to shine a spot of a pure color, say, 500 nm, on a screen and, right next to it, a spot that's a mixture $(r, g, b = 1 - r - g)$ of the three primaries of the RGB model. The values of r and g are varied until the observer judges the two spots to be indistinguishable. The point (r, g, b) is then plotted in the three-dimensional RGB color space. When all the pure colors have been plotted in this way, the points are connected to form a smooth curve, $\mathbf{P}(\lambda) = (r(\lambda), g(\lambda), b(\lambda))$. This is the *pure spectral color curve* of the RGB model (Figure 21.20).

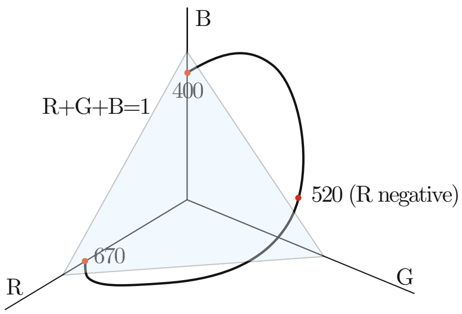

Figure 21.20: Pure RGB Spectral Color Curve.

An important property of the curve is that some of r, g, and b may sometimes have to be negative. An example is $\lambda \approx 520\,\text{nm}$, where r turns out to be negative. What is the meaning of adding a negative quantity of green in a color defined by, for example, $S = 0.8R - 0.1G + 0.3B$? To understand this we rewrite the equation in the form $S + 0.1G = 0.8R + 0.3B$. Written in this form, the equation implies that color S cannot be constructed from the RGB primaries, but color $S + 0.1G$ can. The important conclusion is that not every color can be created in the RGB model! This is illustrated in Figure 21.21. For some colors, we can create only an approximation. This fact applies to all color models that can be created in practice.

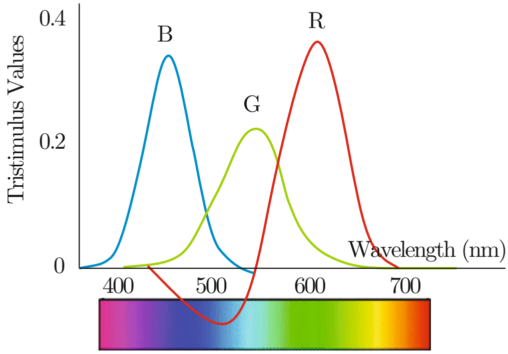

Figure 21.21: RGB Color Combinations.

> Certain combinations have certain effects. For instance the opposition of yellow and violet, blue and orange, that can appear *natural* in a way, because natural shadows are blue or violet. Light and shade, you see? Whereas red and green, if you put them next to each other—sometimes you can see a kind of dancing yellow line where they meet . . .
>
> —A. S. Byatt, *The Matisse Stories*, (1993).

21.11 The CIE Standard

The CIE standard was developed in 1931 by the International Committee on Illumination (Commission Internationale de l'Éclairage). It is based on three carefully chosen artificial color primaries X, Y, and Z. They don't correspond to any real colors, but they have the important property that any real color can be represented as a linear combination $xX + yY + zZ$, where $x + y + z = 1$ and none of x, y, and z are negative (Figure 21.22).

The plane $x + y + z = 1$ in the XYZ space is called the CIE chromaticity diagram (Figure 21.23). The curve of pure spectral color in the CIE diagram covers all the pure colors, from 420 nm to 660 nm. It is shaped like a horseshoe.

Point $w = (0.310, 0.316)$ in the CIE diagram is special and is called "illuminant white." It is assumed to be the fully unsaturated white and is used in practice to match to colors that should be pure white.

◇ **Exercise 21.5:** If illuminant white is pure white, why isn't it on the curve of pure spectral color in the CIE diagram?

The CIE diagram provides a standard for describing colors. There are instruments that can determine the (x, y) coordinates of a given color. Also, given the CIE coordinates of a color, those instruments can generate a sample of the color. The diagram can also be used for useful color calculations. Here are some examples:

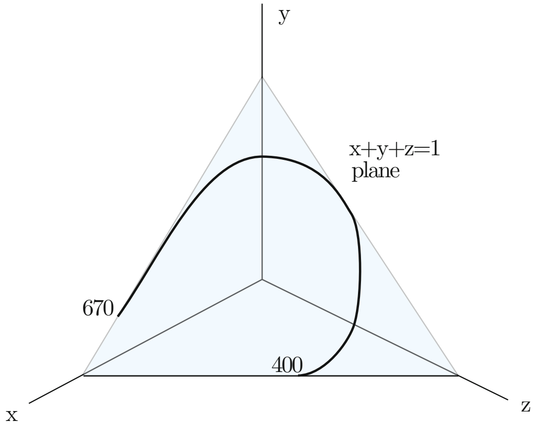

Figure 21.22: Pure Spectral Color Curve.

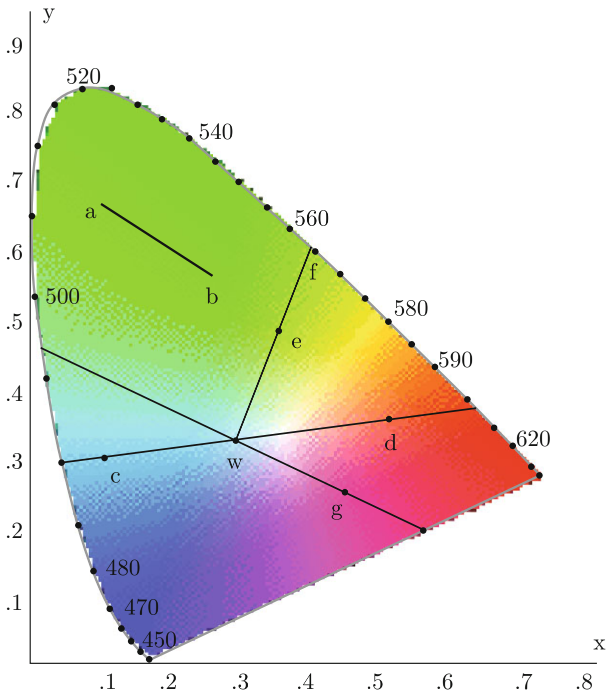

Figure 21.23: The CIE Chromaticity Diagram.

1. Given two points a and b in the diagram (Figure 21.23), the line connecting them has the form $(1-\alpha)a+\alpha b$ for $0 \leq \alpha \leq 1$. This line shows all the colors that can be created by adding varying amounts of colors a and b.

2. Imagine two points, such as c and d in Figure 21.23. They are on the opposite sides of illuminant white and, therefore, correspond to complementary colors.

◇ **Exercise 21.6:** Why is that true?

3. The dominant wavelength of any color, such as e in Figure 21.23 can be measured in the diagram. Just draw a straight line from illuminant white w to e and continue it until it intercepts the curve of pure spectral color. Then read the wavelength at the interception point f (564 nm in our example).

◇ **Exercise 21.7:** How can the saturation of color e be calculated from the diagram?

◇ **Exercise 21.8:** What is the dominant wavelength of point g in the CIE diagram?

> Have you lost your mind? What color is this bill?
>
> —Lori Petty (as Georgia "George" Sanders) in *Lush Life* (1996).

4. The *color gamut* of a device is the range of colors that can be displayed by the device. This can also be calculated with the CIE diagram. An RGB monitor, for example, can display combinations of red, green, and blue, but what colors are included in those combinations? To find the color gamut of an RGB monitor, we first have to find the locations of pure red, green, and blue in the diagram (these are points $(0.628, 0.330)$, $(0.285, 0.590)$, and $(0.1507, 0.060)$), then connect them with straight lines. The color gamut consists of all the colors within the resulting triangle. Each can be expressed as a linear combination of red, green, and blue, with non-negative coefficients (a convex combination).

Interestingly, because of the shape of the horseshoe, no three colors on or inside it can serve as ideal primaries. No matter what three points we select, some colors will be outside the triangle defined by them. This means that no set of three primaries can be used to create all the colors. The RGB set has the advantage that the triangle created by it is large, and thus contains many colors. The CMY triangle, for example, is much smaller by comparison. This is another reason for using red, green, and blue as the RGB primaries.

21.12 Luminance

RGB and CMY(K) are color spaces that specify a color in terms of three primary colors. Certain applications that deal with color pixels, most notably image compression methods such as JPEG, can benefit from a color space whose components are other quantities, not necessarily primary colors. The most common of these spaces is designated "Y Cb Cr" and is referred to as the luminance-chrominance color space.

The CIE defines color as the perceptual result of light in the visible region of the spectrum, having wavelengths in the region of 400–700 nm, incident upon the retina (a nanometer, nm, equals 10^{-9} meter). Physical power (or radiance) is expressed in a spectral power distribution (SPD), often in 31 components, each representing a 10-nm band.

The CIE defines brightness as the attribute of a visual sensation according to which an area appears to emit more or less light. The brain's perception of brightness is impossible to define, so the CIE defines a more practical quantity called *luminance*. It is defined as radiant power weighted by a spectral sensitivity function that is characteristic of vision (the eye is very sensitive to green, slightly less sensitive to red, and much less sensitive to blue). Based on tests for the sensitivity of the eye to colors, luminance is defined as the weighted average $\mathrm{Y} = 0.3\mathrm{R}+0.59\mathrm{G}+0.11\mathrm{B}$. The luminous efficiency of the Standard Observer is defined by the CIE as a positive function of the wavelength, which has a maximum at about 555 nm. When a spectral power distribution is integrated using this function as a weighting function, the result is CIE luminance, which is denoted by Y. Luminance is an important quantity in the fields of digital image processing and compression.

Luminance is proportional to the power of the light source. It is similar to light intensity, but the spectral composition of luminance is related to the brightness sensitivity of human vision.

In simple terms, luminance indicates how much luminous power will be perceived by the eye and brain from a given light source. In even simpler terms, luminance indicates how bright a light source appears.

The eye is very sensitive to small changes in luminance, which is why it is useful to have color spaces that employ Y as one of their three components. A simple way to do this is to subtract Y from the blue and red components of RGB, and use the three components Y, $\mathrm{B}-\mathrm{Y}$, and $\mathrm{R}-\mathrm{Y}$ as a new color space. The last two components are called chroma. They represent color in terms of the presence or absence of blue (Cb) and red (Cr) for a given luminance intensity. Chroma is also related to color saturation.

Various number ranges are used in $\mathrm{B}-\mathrm{Y}$ and $\mathrm{R}-\mathrm{Y}$ for different applications. The YPbPr ranges are optimized for component analog video. The YCbCr ranges are appropriate for component digital video such as studio video, JPEG, JPEG 2000, and MPEG.

Many image compression methods, most notably JPEG, are lossy. They achieve high compression rates by losing image information to which the eye is not sensitive. The eye is very sensitive to changes in luminance, which is why a lossy compression algorithm often starts by converting the color representation of the image pixels from the original (RGB or CMYK) to YCbCr. The algorithm proceeds by losing more data from the two chroma components and less data from the Y component.

The YCbCr color space was developed as part of Recommendation ITU-R BT.601 (formerly CCIR 601) during the development of a worldwide digital component video standard. Y is defined to have a range of 16–235; Cb and Cr are defined to have a range of 16–240, with 128 equal to zero. There are several YCbCr sampling formats, such as 4:4:4, 4:2:2, 4:1:1, and 4:2:0, which are also described in the recommendation.

Conversions between RGB with a 16–235 range and YCbCr are linear and therefore simple. Transforming RGB to YCbCr is done by (note the small weight of blue)

$$\begin{aligned}
\mathrm{Y} &= (77/256)\mathrm{R} + (150/256)\mathrm{G} + (29/256)\mathrm{B},\\
\mathrm{C}b &= -(44/256)\mathrm{R} - (87/256)\mathrm{G} + (131/256)\mathrm{B} + 128,\\
\mathrm{C}r &= (131/256)\mathrm{R} - (110/256)\mathrm{G} - (21/256)\mathrm{B} + 128,
\end{aligned}$$

while the opposite transformation is

$$\begin{aligned}
\mathrm{R} &= \mathrm{Y} + 1.371(\mathrm{C}r - 128),\\
\mathrm{G} &= \mathrm{Y} - 0.698(\mathrm{C}r - 128) - 0.336(\mathrm{C}b - 128),\\
\mathrm{B} &= \mathrm{Y} + 1.732(\mathrm{C}b - 128).
\end{aligned}$$

When performing YCbCr to RGB conversion, the resulting RGB values have a nominal range of 16–235, with possible occasional values in 0–15 and 236–255.

There are other color spaces and other quantities associated with color. Here, we'll mention the term luma. Luma (denoted by Y') is gamma-corrected luminance (Section 26.4.4).

21.13 Converting Color to Grayscale

Very few people still use black and white film in their cameras. Similarly, very few still keep their old, black and white television sets. We generally prefer color images to grayscale images, but sometimes a color image has to be converted to grayscale. Perhaps the best example of this is printing. Most books and newspapers are printed in black and white, so a printer driver (software that converts a document from internal representation to commands sent to a printer) has to convert any color images to grayscale.

Part II of this book shows that a three-dimensional object can be projected to two dimensions in several ways because three dimensions are so much more complex and richer than two dimensions. Similarly, color images can be converted to grayscale in a number of ways. Such conversion is not unique and this section outlines three approaches to this problem.

The most common approach uses the concept of luminance. It assigns different weights to the three color components of a pixel such that the luminance of a resulting grayscale pixel matches that of the original pixel.

The conversion is generally done in three steps. The three color components of a pixel are gamma corrected (or gamma expanded), the luminance is computed as a weighted average, and the result is reverse gamma corrected (gamma compressed). The weighted average is of the form Y = 0.3R + 0.59G + 0.11B or similar (point Y in Figure 21.24).

The three weights have been determined experimentally in tests that measured the sensitivity of the eye to the three RGB primary colors. The eye is more sensitive to green than to blue, so when equal intensities of green and blue enter the eye, we perceive the green as brighter. Thus, the weighted average above produces a grayscale pixel brightness that is perceptually equivalent to the brightness of the original color pixel.

Color can also be converted to grayscale by removing its saturation. The next approach is based on desaturating the color. Most colors that we perceive are a mixture of many wavelengths, with perhaps one wavelength dominating; pure colors are rare. Section 21.2 defines color saturation as the percentage of the luminance that resides in the dominant wavelength of the color. Thus, the larger the saturation, the closer the color is to a pure color, while smaller saturation brings any color closer to white.

One way to desaturate a color is illustrated by the RGB color cube. This cube has a diagonal (referred to as the neutral axis) from its white corner $(0, 0, 0)$ to its black corner $(255, 255, 255)$. Given a point A in the RGB color cube, it can be desaturated by finding the point B that's closest to A on the neutral axis and replacing the three RGB values of A with those of B. This process is illustrated in Figure 21.24 for $A = (220, 60, 120)$.

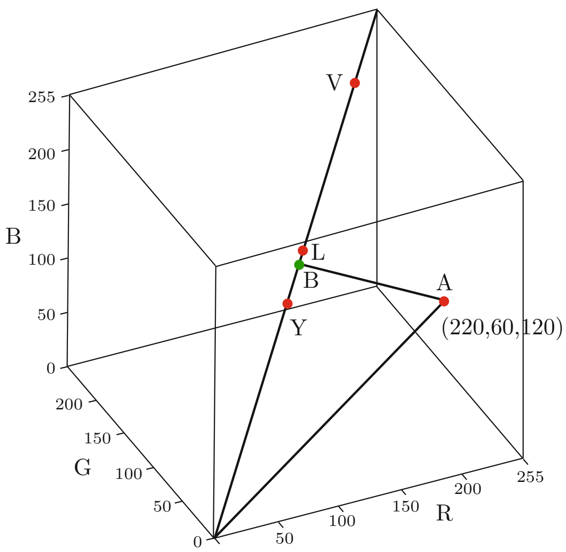

Figure 21.24: RGB Cube and Neutral Axis.

When this approach to color conversion is implemented in practice, speed is important, so instead of finding the exact location of B, graphics software often employs the approximation $L = [\max(R, G, B) - \min(R, G, B)]/2$ and uses point L instead of B.

Point L shown in the figure has coordinates $(220 - 60)/2 = 80$ and is close to the ideal point B. Thus, this approach to color conversion converts color pixel $(220, 60, 120)$ to grayscale 80. In general, converting to grayscale with this approach produces an image that many would judge too dark and flat (i.e., with little contrast).

Yet another approach to color-to-grayscale conversion is based on the HSV color space. When a color is converted to hue, saturation, and value, the "value" component can be used as the grayscale equivalent of the color. In practice, this conversion is done by simply selecting the largest of the RGB color components and using it as the grayscale (point V in Figure 21.24).

Laundry is the only thing that should be separated by color.
—Anonymous