1 Foundations

Although this is a very practical book, you must have some understanding of the foundations of light and color, as they relate to photography, in order to use its practical information effectively. This chapter provides a basic look at color and color management. It won't make you a color expert, but it will help you understand how color works and how a color management system can help you.

Chapter Contents

The Nature of Light

Although the science of color is incredibly complicated, the subject of color management doesn't need to include a steep learning curve. The first topic on that learning path is how light carries and contains color (Figure 1.1).

Figure 1.1
Dewitt Jones has made a career of transforming the light reflected off a scene into dramatic images that truly captivate the viewer. By understanding how light works, you are better equipped to produce the very best results possible. In this image Dewitt has started with a digital infrared photographic image, and then hand-tinted the color in Photoshop to produce the final image. (Photograph by Dewitt Jones, www.dewittjones.com.)

Note: Tempted though you may be to skip ahead in this book and get your hands dirty (figuratively speaking), I encourage you to read this chapter first. It will give you a stronger foundation to build on as you implement color management in your own workflow. Learning the concepts presented in this chapter will make the rest of the topics in this book much easier to understand.

Light is a fascinating thing. The light from the sun travels over 93 million miles—in about eight minutes—through the void of space to meet the earth. It provides the heat that makes life possible and also allows us to see and helps make life more enjoyable. Light allows us to enjoy the golden rays of a sunrise, the cool blue of deep shade, the vibrant green of a lush hillside, the dramatic yellows, oranges, and reds of a sunset (see Figure 1.2), and the pink haze of twilight. Even after the sun has disappeared below the horizon, its light can be reflected from the moon to illuminate our world in magical ways. Our fascination with color begins at an early age, and for photographers it certainly continues to be a fascination. Although photographers who focus on black-and-white photography may not place as much emphasis on color, they are very much in tune with the subtleties of light in general.

Figure 1.2 Light provides the source of color that we enjoy capturing in photographs. (Photograph by Ira Meyer, www.irameyer.com.)

Light is energy. Energy travels as waves, with a wavelength that defines the energy level of a particular source. These various ranges of wavelengths are categorized in sections, from gamma rays and x-rays, through ultraviolet, to the visible spectrum, and infrared, and on to microwaves, radar waves, and radio waves. Color is observable if it falls within the "visible spectrum," which is a range of wavelengths from approximately 380 to 780 nanometers (see Figure 1.3). One nanometer is one billionth of a meter, so we're talking about pretty small waves of energy.

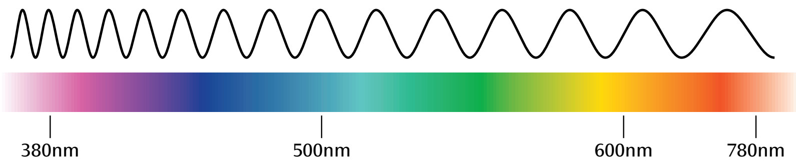

Figure 1.3 The visible spectrum includes light that ranges from wavelengths of about 380 nanometers to 780 nanometers.

In a sense, you could say that color doesn't exist. Light energy does not actually have color as a tangible property. Rather, various wavelengths of energy cause our optical system to perceive particular colors. For example, red is around 650 nanometers, green is around 525 nanometers, and blue is around 450 nanometers. When energy at these wavelengths reaches our eye, it causes us to perceive red, green, or blue. Although color can be defined based on specific wavelengths of light, it isn't a physical attribute of light by itself, but rather the way we perceive the energy at particular wavelengths. The very real property we perceive as color is created in our minds as a response to the specific wavelengths of light observed.

Quite literally, all the colors of the rainbow—red, orange, yellow, green, blue, indigo, violet—are found on the visible spectrum, with red at the longer wavelength end (around 650 nanometers) and violet at the shorter wavelength end (around 380 nanometers). This range of colors represents the full spectrum of visible light.

Note: You may have learned the name "Roy G. Biv" in science class while you were growing up. This serves as an acronym (Red, Orange, Yellow, Green, Blue, Indigo, Violet) for the colors found in the rainbow, extending from the longest wavelengths (red) to the shortest (violet).

Our eyes contain photoreceptors that enable us to convert light into signals in the brain, providing the experience of vision. The photoreceptors are of two types: rods and cones. The rods number around 120 million, and are very sensitive to light but aren't able to perceive color. There are around 6 to 7 million cones, and they provide color sensitivity. The human visual system allows us to perceive a huge range of colors, and all at the same time. When we look at a landscape before us, as shown in Figure 1.4, we are seeing many colors at once, and we are able to perceive all of them concurrently.

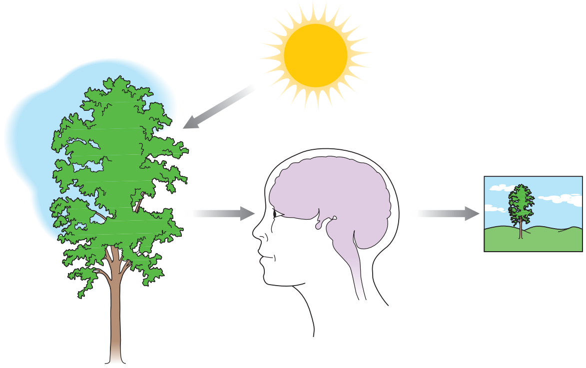

Figure 1.4
The human optical system translates light energy of various wavelengths into the perception of color.

In the real world, of course, color is a bit more complicated than I've made it sound so far. The colors we see are created by the reflection of light observed by our eyes. Although objects do not actually possess color from a technical standpoint, they have properties that cause them to absorb and reflect particular wavelengths of light. If an object absorbed all light, it would be perceived as completely black. Of course, the only object that can achieve this is a black hole. All other objects reflect some degree of light, but perhaps not enough for us to distinguish it. If an object reflected all light, it would be perceived as being pure white. Again, real world objects aren't going to have this purity, but the basic concept holds true. Most objects, of course, absorb and reflect various wavelengths of light in various combinations.

While we can define a specific color based on a wavelength of light, a single observed color is actually composed of light at a variety of wavelengths, which means that a single color is built from many other colors. The amount of light at different wavelengths used to form a specific color can be described using *spectral curves* like the one in Figure 1.5. These are charts that illustrate the amount of light energy at various wavelengths is contained within a specific color being measured.

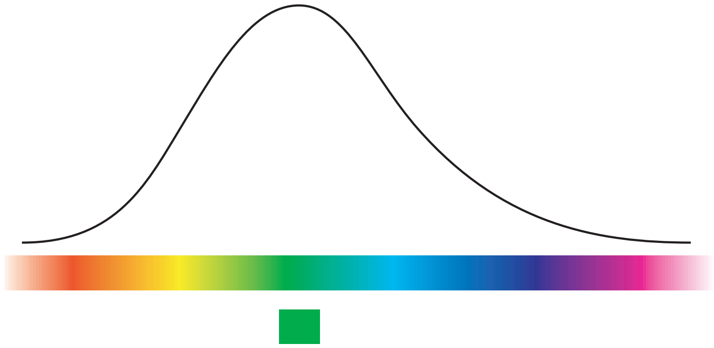

Figure 1.5
A spectral curve provides a visual indication of the amounts of light at various wavelengths that produce a particular color.

Light and color are intertwined, and they represent the foundation of photography and the genesis of color management. While color management is a science unto itself, involving an effort to control color in one form or another, understanding the underlying light and color you are trying to tame with a color management system is helpful for photographers at all levels.

Light in Photography

Photography has everything to do with light. In fact, you could say that photography is nothing more than controlling—to the extent possible—and recording light. When you capture a photographic image, whether on film or with a digital camera, you are storing the visible effect of light so it can be presented again for the enjoyment of others, as depicted in Figure 1.6. Film and digital image files effectively store light so it can be reproduced to provide a similar sensation to observing the light reflected off the scene in the first place.

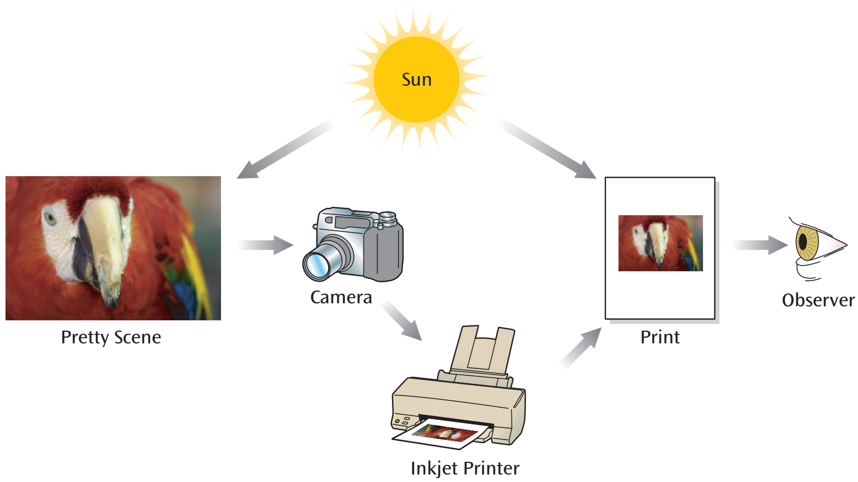

Figure 1.6 Light from an illumination source is reflected from objects and then recorded by the camera. The image can then be printed to simulate what was originally observed, again depending on reflected light to be seen.

Light is the key ingredient that makes photography possible. Light is emitted from the sun or artificial light sources, bombarding the objects in the world around us. When light energy comes in contact with an object, it passes through, is absorbed by, or is reflected by the object (in various degrees and combinations depending on characteristics of the object such as translucency). The reflected light moves outward from the object, so it can continue on in all directions. This light can then be observed by your eyes and recorded by your camera. The scene you are able to see is, in fact, created in your mind, as a response to the light energy stimulating your optic nerves.

The light reflected from a scene does much more than allow you to see what is there. Light isn't just a source of illumination. In a sense, light *is* the scene before you, because that light is what your brain is able to process as the image you perceive.

Each object has its own reflective properties, causing light of certain wavelengths to be absorbed while the rest is reflected. Capturing an image is a matter of recording that reflected light as it streams through your lens. Optimizing that image in the digital darkroom involves making judgments about the light emitted from your monitor to produce the perfect image. The final print closes the circle, reflecting light off the paper to re-create the image you envisioned at capture.

The Nature of Color

We become familiar with color at a very early age, and naming colors is one of the first skills many children learn after they are able to talk. We live with color and experience it constantly. Despite our familiarity with color, many are not familiar with the terminology used to describe color. Understanding these terms so that you have a better understanding of color, and are better able to deal with variations in color once you establish your color-managed workflow, is important. The following terms describe the various attributes of color:

Hue is what we usually refer to as "color." Of course, color consists of properties other than hue, but the hue is what we consider the color to be. The names we use to describe colors, such as "red" or "blue," represent the hue of that particular color. If you prefer a more technical definition of hue, it is the dominant wavelength of light in a particular color. For example, red consists of mostly red light, but unless it is absolutely pure it also contains some component of light at other wavelengths.

Saturation refers to the purity of the color, and it is what we would normally consider the vibrancy of a particular color. High saturation represents a very vibrant color, while low saturation would represent less vibrant colors, or a shade of gray. Saturation can also be referred to as chroma.

Brightness, also referred to as luminance, is a measure of the amount of light emitted or reflected from a subject. When referring to a color, it refers to the brightness or lightness of the color.

Together, these three primary components allow you to precisely describe a color in general terms. They form the basis of the HSB (hue, saturation, brightness) system for defining specific colors. Even with other color models (which will be discussed later in this chapter), these terms are used to describe the attributes of color.

Perceived Color

So far the presentation of color makes it sound rather absolute. An object absorbs and reflects light energy of different wavelengths in a particular combination, and the combination of reflected wavelengths determines what color is seen. Of course, it isn't quite that simple.

Color, as we perceive it, is the product of three individual factors working together:

The light source, also referred to as the illuminant. When talking about basic color theory, we usually assume a light source that is pure white, meaning that it contains all wavelengths of the visible spectrum. As a photographer, you already know that the sources of light available to you for your photography are anything but white. The actual color of the light illuminating your subject has a significant effect on the final color appearance of that object. The standard illuminant for color management is 5000 degrees Kelvin, often referred to as D50.

The object, which reflects the light from the light source. Each object has its own properties that affect the wavelengths of light that get reflected or absorbed, and to what degree. These properties affect the appearance of the light reflected off the object and, therefore, the color that is observed for an object under particular lighting conditions.

The observer of the reflected light. The ultimate observer is, of course, a human observer. But when it comes to the effect on a photographic image, there are other observers to consider. For example, the camera is an observer of reflected light, and it reacts differently than the human optic system does. Furthermore, filters and other devices can be used to modify how the light reflected from an object is observed. Other observers include the devices used for calibrating and profiling various devices, which will be discussed in later chapters.

Each of these three factors plays its own role in the creation of the perceived color (see Figure 1.7), and so changes to any of them affects the final perceived color.

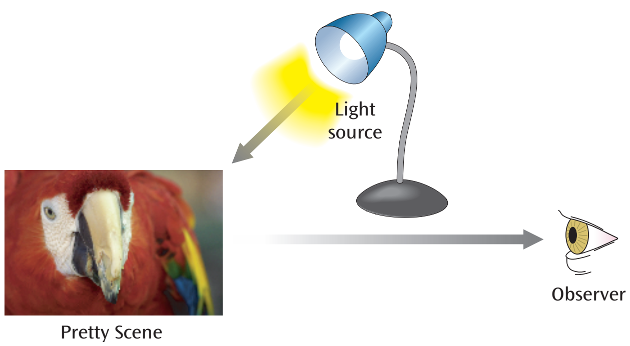

Figure 1.7
The light source, object, and observer each influence the observed color. If any of the three changes, the perceived color changes.

The Color Wheel

A color wheel (such as the one in Figure 1.8) is a very effective tool for understanding the relationships between colors. Earlier in this chapter, I referred to the full visible spectrum, which includes all the colors of the rainbow. The color wheel arranges the range of colors in the visible spectrum into a circle. Instead of a linear representation based on wavelength of light, hue is defined based on the position on the wheel, measured in degrees around the circle.

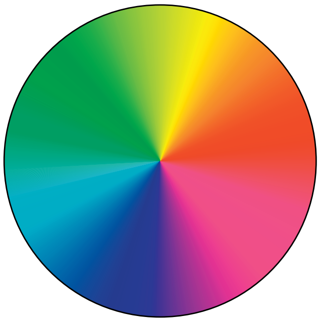

Figure 1.8
The color wheel provides a way to visualize the color spectrum in a way that illustrates the relationships between different colors. This basic color wheel shows only actual hues at pure saturation.

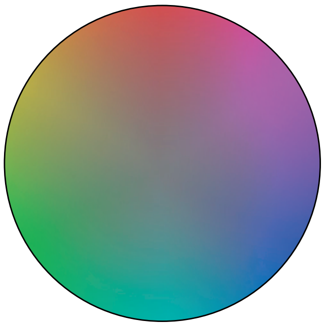

Figure 1.9
This color wheel shows both hue and saturation in one display.

Saturation is often shown on a color wheel with highly saturated colors at the edge and unsaturated colors in the middle (see Figure 1.9). This allows more information about specific colors to be represented on a single graphic.

A color wheel generally doesn't show lightness (brightness) information, simply because it would require a three-dimensional construct to show all three values of hue, saturation, and brightness. With software that uses some form of color wheel to select a color, hue and saturation can often be viewed at one time, with a separate control for brightness (see Photoshop's version of this in Figure 1.10). With some programs, you can change the representation of the color wheel being displayed so you can view any two properties of color at once, with a separate control for the third.

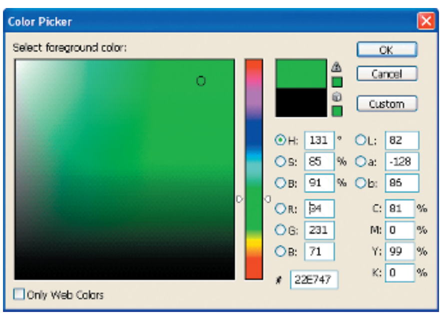

Figure 1.10
The Photoshop Color Picker provides a visualization of the relationships of the three color attributes of hue, saturation, and brightness. The selected attribute is shown on the smaller bar, while the other two attributes are displayed on the larger box along the horizontal and vertical axes.

Colors on opposite sides of the wheel are opposite—or complementary—to each other. Adjusting color balance shifts the colors in your image from one side of the wheel to the other on a specific axis. For example, as you're adjusting the color balance for the green/magenta axis, you are shifting the pixel values in your image between more magenta on one side of the color wheel and more green on the other side of the color wheel. I'll talk more about color balance adjustments in Chapter 6, "Optimization."

Understanding the relationships of adjacent colors on the color wheel is also important. On each side of a color on the color wheel are the colors that produce that color. For example, cyan is found between blue and green on the color wheel, and cyan is composed of blue and green. Understanding the intricacies of these color relationships can be incredibly helpful as you are adjusting your images, and also as you are evaluating your color management workflow to confirm accurate results.

Color Models

Color in the real world is all about wavelengths and intensity of light. When we work with digital images, we need a way to translate those values so color can be accurately described. Color models provide an organizational system for color in digital imaging. These color models can either be device-independent or device-dependent.

A device-independent color model does not depend on any device to describe color. LAB and CIE XYZ are examples of device-independent color models. They are modeled on the way the human visual system translates color, but they can be challenging to work in for many photographers who do not have experience working within this color model. The benefit of a device-independent color model is that it can precisely describe a color without referencing a specific device.

Device-dependent color models, on the other hand, revolve around providing instructions to devices in an effort to match a device-independent color value. The fact that we work with color using device-dependent color models does a lot to explain why profiles are critical to get predictable results in the digital darkroom. Because these color models don't describe a true color, but rather provide color values to a device that is expected to produce the intended color from those values, a method to translate color between devices is required.

The following color models are those used most often in photo-editing applications or with color management tools in general:

The CIE XYZ color model is a device-independent color model developed in 1931 by the International Commission on Illumination (called CIE after its French name). This color model is device-independent, and it is designed to represent color based on the way it is actually perceived under very specific viewing conditions (defined as the "CIE 1931 Standard Observer"), so that it describes actual color as it is perceived by an individual with normal color vision. Each of the values it defines (termed *X*, *Y*, and *Z*) represent imaginary primary values based on a mathematical model rather than real color values.

Note: Although not used as a color model for any photo-editing software, CIE XYZ is the basis of most device-independent color models in use today.

The LAB color model (CIE LAB) is the most commonly used variant of the device-independent color models that have been developed by CIE. The LAB color model consists of three channels to describe a color. The Lightness channel, as the name implies, defines luminosity. The "A" channel describes color on an axis between blue and yellow, and the "B" channel describes color on an axis between red and green. For those already familiar with the other color models, this may seem like an odd arrangement. However, it is based on the way the human vision system perceives color, and it provides an effective way to describe color in a device-independent way. The LAB color model provides the reference color space for most color management systems.

The RGB color model is device-dependent, and it is used when describing color as emitted light, such as that on a monitor. Colors are described by specifying a value for red, green, and blue. In other words, RGB values specify how much light of each of the three constituent colors should be blended to produce the intended result. Mixing all three colors at their maximum value will (in theory at least) produce pure white, and omitting all three colors completely will (again, in theory) produce pure black. Most images you will work on will be edited using an RGB color model, as that is the way monitors, digital projectors, and inkjet printers deal with color.

The CMYK color model is the other commonly used device-dependent color model. It describes colors based on the subtractive primaries of cyan, magenta, and yellow. The obvious example of this is for inks, which absorb and reflect light to produce the colors we see. When none of the colors are included, the result is white (in theory; the color of the paper to which the ink is printed is a significant source of variation here). When all of the colors are mixed at their full values, the result is black (in theory). In reality, because the inks used are not perfectly pure, mixing all the cyan, magenta, and yellow inks won't produce black (it usually creates a muddy brown), so black ink (the Key color, providing the "K" in CMYK) is added to the mix. Although most images are adjusted using an RGB color model, you may have occasion to use the CMYK color model, as I'll discuss in later chapters.

Primary Colors

Within any particular color model, specific attributes are used to describe a color. These attributes are referred to as "primaries," and the values of the primaries describe a specific color. In the case of CIE XYZ or LAB, these primaries are not based on how we actually create color, but rather are based on how we perceive color. For this reason, they tend to be a bit complicated, and not an easy way to adjust color in your photographic images.

Therefore, most digital photo editing is performed using the RGB or CMYK color models. Each of these color models uses specific colors for their primaries, thereby making use of "primary colors." Understanding the concept of primary colors is important both for adjusting your images and for understanding color management.

An understanding of primary colors will help you understand what changes are required to produce the result you are trying to achieve in your images when you are adjusting them.

Each color that we produce on a monitor or printer is composed of some combination of primary colors. For emitted light, the three primary colors are red, green, and blue. They are called *additive primary colors* (see Figure 1.11) because you create white by adding all three colors together, and each of the individual components is adding color to the mix. The absence of all three primary colors represents black because it represents a complete absence of light. The presence of all three primary colors at their maximum value represents white, as shown in Figure 1.12.

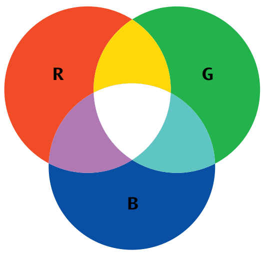

Figure 1.11
The additive primary colors of red, green, and blue can be blended to produce all possible colors with light. Combining all three at maximum value produces white.

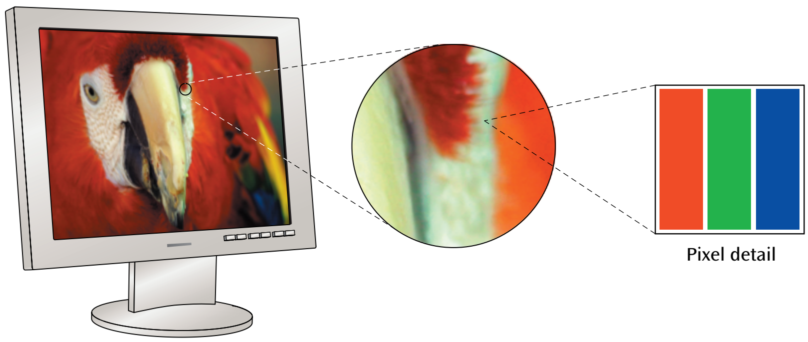

Figure 1.12 The display on a monitor is composed of the three additive primary colors of red, green, and blue. Each pixel is formed by these three colors, with the intensity of each modified to produce the appropriate color.

For reflected light—such as with the inks used in a print—the three primary colors are cyan, magenta, and yellow. These colors are called the *subtractive primaries* (see Figure 1.13) because each of the primary colors is absorbing light (subtracting light) so that only light of certain wavelengths is reflected back. The absence of all three primary colors represents white because no light is being absorbed—although in fact, "white" would be the color of the paper because the CMYK color model is used primarily for printed output. The CMYK color model represents the colors of inks (cyan, magenta, yellow, and black) used to produce printed output for virtually all printers. The presence of all colors represents black (at least in theory) because all light is being absorbed.

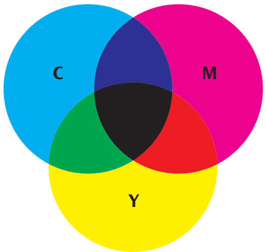

Figure 1.13
The subtractive primary colors of cyan, magenta, and yellow can be blended to produce all possible colors with reflected light. Combining all three at maximum value produces (in theory) black.

The additive and subtractive primary color groups are related to each other in terms of the relationships between colors, as diagrammed in Figure 1.14. Each primary in one group has an opposite color in the other group. These relationships are important to understand when working with your own images in the digital darkroom and making prints of those images.

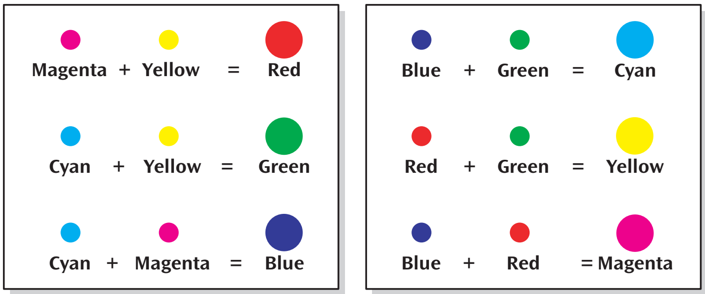

Figure 1.14 The additive and subtractive primary colors are opposites of each other. Each additive primary color can be formed by combining the two subtractive primary colors that are not its opposite, and each subtractive primary color can be formed by combining the two additive primary colors that are not its opposite.

In the color wheel, the opposite color for a given primary is exactly opposite it. For example, red is a primary color for the RGB color model. On a monitor, red is produced by emitting red light at maximum intensity with no green or blue light included. The opposite of red is cyan, which is a primary color in the CMYK color model. Cyan absorbs red light, reflecting back green and blue. Red and cyan are opposites.

You can also refer to the color balance control in most color management software, which usually provides sliders that allow you to shift the balance between each of the primaries. An easy way to remember the colors that are opposite each other in the RGB and CMYK color models is to think of them in order. The first colors in each (red and cyan) are opposite each other, as are the second (green and magenta) and third (blue and yellow). Understanding how these colors interact makes color adjustments much easier, and it will help you better understand what is going on behind the scenes in your color-managed workflow.

Metamerism

If you ever need to impress someone at a cocktail party, try to work the term "metamerism" into the conversation. Microsoft Bookshelf defines metamerism as "The condition of having the body divided into metameres, exhibited in most animals only in the early embryonic stages of development." If you're having a difficult time understanding what that has to do with color management, you're not alone. When it comes to color management, metamerism seems to be a word that is rarely understood, and usually scorned.

Metamerism, despite its bad reputation among those who have actually heard of it, is actually a very good thing. Metamerism is a phenomenon where two colors match under one lighting condition, but don't match under different lighting (see Figure 1.15). That may sound like a bad thing, and in many cases it is, but it is also the underpinning of a very good thing.

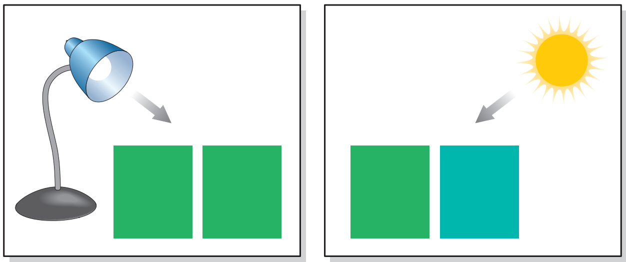

Figure 1.15
Metamerism is a phenomenon where two colors match under one lighting condition, but don't match under different lighting.

Each color has specific spectral properties. In fact, every object that reflects light has specific spectral properties. These properties are the properties of the reflected light, and they are generally defined as how much light is emitted or reflected at each specific wavelength. Because of metamerism, objects with very different spectral responses can actually appear as the same color to the human eye. This is very significant.

Metamerism makes it possible to create the same color using a variety of inksets, for example. Without metamerism, we would have to use exactly the same pigments to produce a particular color. It would also require that to produce a specific color, you would need an ink of that specific color. That would not be very practical. You can imagine how many different inks we would need in our inkjet printers if every color we wanted to produce needed an individual ink. With metamerism, we can produce a variety of colors that will produce a specific color experience despite the spectral properties of the colors used to produce the colors.

Metamerism has a pitfall, of course, or it wouldn't have such a bad reputation. For one thing, it is possible to produce two colors that look like a perfect match under one lighting condition, but don't appear to match in a different lighting condition. This could be a huge problem if you are trying to match two different printed items with the same colors. They could appear to match until you deliver them to the client.

For photographic images, it can also be a problem, though usually not in the color matching sense. Instead, it can create a problem where the relationship between colors appears different under certain lighting conditions. Colors that have an appropriate relationship under certain lighting may not match under other lighting. In this type of situation, different lighting would effectively cause different changes in different areas of the print, which could obviously be a problem. However, it is not a common problem.

There are several situations where metamerism has gained a bad reputation it doesn't deserve. One of the most common relates to observing a printed image under differing light sources. This goes back to the discussion of perceived color earlier in this chapter. If you change the illumination source, you change the perception of the print. A print that looks perfect in the lighting of your digital darkroom may not look very good under incandescent lighting, for example. This is not metamerism. This is a change in the illumination that causes a change in the appearance of the image because a different light source means different color will be reflected, which in turn means that different colors will be observed. You can't expect a print to look exactly the same under all lighting conditions, any more than you would expect a mountain range to look the same at sunrise as it does at high noon.

Another issue that has given metamerism a bad name should be referred to as *bronzing*. This is a situation where certain inks have different reflective properties, causing them to appear slightly different under certain lighting It is most common with black ink, particularly with pigment-based inks. For example, photographers using the Epson Stylus Photo 2000P printer frequently complain about bronzing.

For most photographers, understanding metamerism is purely academic. It is a major issue underlying the science of color management and worth understanding so that you will have a greater appreciation of what is involved with producing prints with matching colors. You can't do anything to change the fact that metamerism exists, so it is better to understand it, appreciate the benefits it provides, and accept the potential pitfalls it can produce.

Color Profiles

Profiles are a major component of any color management system. They are what actually makes it possible to produce prints that match what you see on your monitor or produce scans that accurately reflect what is contained in the piece of film you are scanning.

A profile is a data file that describes the color behavior of a specific device. The color values stored in an image can be thought of as instructions for the device that will display or reproduce those colors. In order to maintain accurate colors, those values must be related to a device-independent color model. The profile contains a table that lists the specific device-dependent color values and their equivalent values using the device-independent LAB color model.

Most image files that photographers work with are in either the RGB or CMYK color model, where values are assigned to the additive or subtractive primary colors, defining the color of each pixel. The color values provide instructions for how much of each primary color should be used to produce the intended color. The results are dependent upon the specific primary colors used. For example, each printer uses specific inks, and so "cyan" for one printer is different from "cyan" on a different printer, because the cyan ink for each has a slightly different color value. Even the paper used to print can result in different values, because of the way the inks are absorbed by the paper. Therefore, the color values used in the RGB or CMYK color models don't have any meaning by themselves. They only have a specific color meaning when a device is used to produce output based on the color numbers.

A profile translates these values stored in a digital image to a device-independent color model—usually LAB—so that they have a specific color meaning. In this way, profiles allow consistent colors to be produced by a wide range of devices, because each device can have a specific profile that allows it to translate the colors in your image file.

This translation is key to producing accurate results, and so it is important that the ability to interpret those color values is maintained for each image file. They can, therefore, be "tagged" with a profile, which means that a profile is attached to the file so an application used to read the image file can know exactly what color is represented by each of the color values stored within the image file. In effect, the translation of the color values is contained within the image file itself, so that the meaning of each color value can be accurately determined.

Why Not LAB?

With all this talk about using LAB as a reference color space, you may be wondering why all color values aren't simply stored using the LAB color model rather than RGB or CMYK. Actually, doing exactly that is quite possible, by converting your images to LAB mode. However, many people find it very difficult to do color adjustment in LAB because it doesn't relate to any real world color model to which they are accustomed.

Of course, the LAB color model is a real world model, in that it closely mimics the way the human visual system perceives color. However, that doesn't make it a user-friendly way to work with images for those who are not familiar with it.

For RGB and CMYK images, simply storing the LAB color values would be theoretically possible, but that would circumvent the benefit of using these color models. Specifically, they relate to the colors used by actual output devices, and represent a color model with which many of us are familiar.

Of course, provided a profile is embedded in an image file, the RGB and CMYK values are, in a way, stored as LAB, because the profile contains translation tables that allow the LAB values to be interpreted from the RGB or CMYK values.

The profiles used by most applications and operating systems now comply with the standards developed by the International Color Consortium (ICC) and are commonly referred to as "ICC Profiles." The ICC was founded by Adobe, Agfa, Apple, Kodak, and Sun Microsystems to create industry-wide standards for color management. The profile specifications that the ICC created have become the standard used by most operating systems and photo-editing software.

Profiles can be created for input devices such as a digital camera or film scanner, or output devices such as monitors, digital projectors, or printers. Ideally, a profile will be specific to a particular device—not just any scanner of a specific model, but the exact scanner sitting on your desk. This would be a custom profile, and it represents the ideal. However, because manufacturing tolerances are generally very fine, profiles can be made for a specific model of a device with good results. These profiles are referred to as generic or "canned" profiles. For example, each model of inkjet printer includes canned profiles for different types of papers that can be used with that printer.

Profiles can also be used for a more arbitrary purpose when used as a working space. Such a profile does not relate to a specific device, but rather provides a range of colors—a *gamut*—that are available for your images (see Figure 1.16). The available color values are mapped to specific values, typically using the LAB color model, so that each value represents a device-independent color value.

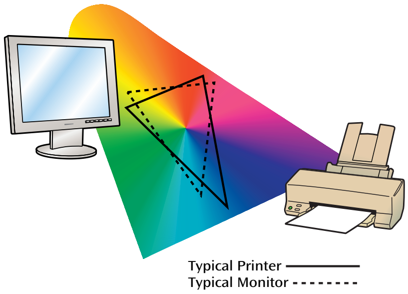

Figure 1.16
Each device has an individual color gamut that defines the range of colors it is able to produce.

Note: In later chapters I'll show you how to build and use custom profiles and how to utilize generic or canned profiles when you aren't able to obtain custom profiles.

Regardless of their specific implementation, whether to describe the color behavior of a specific device or to define a working space for photo editing, profiles provide a way to relate the color values stored in an image file or interpreted by an input or output device to actual colors based on the way the human visual system perceives color. Profiles, therefore, perform what is arguably the most important task within a color management system.

For color displayed on a monitor, profiles allow a conversion of color values from those stored in your image file to the LAB color model before they are then converted to appropriate RGB values that will cause the monitor to display the same color accurately. For prints, the color values in an image file are similarly translated to LAB before being converted to values that the printer understands to produce the specific colors in the image. Similarly, when a digital camera or film scanner reads color values, a profile can help map the values that are "seen" into correct values in the LAB color model, so that they can then be interpreted for accurate display on any output device.

Any of these translations of color values from one profile to another require conversions that are calculated by a color management engine in the software performing the operations. Because each device has a specific color gamut, which is the range of colors it is able to capture or produce, not all of the colors in the source profile will be available in the destination profile. A method of conversion of image data from one profile to another must include a strategy for dealing with both colors that are within the destination profile's color gamut, and those that are not. That method is implemented by one of four rendering intents.

Rendering Intents

In a color-managed workflow, images will need to be converted from one profile to another for various reasons. During such a conversion, colors in the source profile may not be available in the destination profile. Colors that exist in both profiles can be converted quite easily. The out-of-gamut colors need to be dealt with more carefully. The *rendering intent* determines how the color conversion is performed.

In later chapters, we'll look at the different situations where conversions will be performed and how to choose the right rendering intent, but for now we want to gain an understanding of how the rendering intents behave. Four rendering intents are available: saturation, relative colorimetric, absolute colorimetric, and perceptual. They are demonstrated in Figure 1.17.

Absolute Colorimetric

Perceptual

Relative Colorimetric

Saturation

Figure 1.17
Each rendering intent will have a slightly different effect on an image. These pictures show the result of converting the same image to the same profile using each of the four rendering intents. As you can see, saturation causes a shift in color values while maintaining saturation, while the other three produce good (though slightly different) renderings of the image.

Saturation

With the saturation rendering intent, as the name implies, the most important property of each color is the saturation. This is the purity of a color, and it is what we think of as having vibrant colors.

The saturation rendering intent does not change colors that are within the gamut of the destination profile. For out-of-gamut colors, the color is changed to maintain the saturation of the color without regard for the hue. In other words, a highly saturated green that doesn't exist in the destination profile could be converted to a highly saturated red hue if that is the closest match in terms of saturation. This could be an obvious problem for photographic images. Although the results will not be as garish as changing a color from green to red, the important point is that the actual color appearance can potentially change quite dramatically, causing problems for photographic images.

Because of the way the saturation profile works, it is not recommended for use with photographic images. It is designed for situations that require saturated colors where you don't care about the hue, such as when using charts in a PowerPoint presentation.

Absolute Colorimetric

The absolute colorimetric rendering intent maintains the appearance of colors that are within the gamut of the destination space, but changes the values of colors that are out of gamut to the closest reproducible hue. Saturation and lightness may be sacrificed in order to obtain the closest matching hue. Absolute colorimetric includes white in the conversion, so that the "color" of white in the source profile is reproduced in the destination profile. That means that white in your image may not match paper white, which can cause the image to look like it has a color cast. This is because the human visual system is incredibly capable of adapting to different light sources, always translating white to appear white. With absolute colorimetric, the eye may adapt to paper white, resulting in a color cast if the white of the image doesn't match paper white. Black is not adjusted with the Absolute Colorimetric rendering intent.

This issue may cause you to avoid using absolute colorimetric as a rendering intent, and in general I would say that is a good decision. However, there are some situations that make it useful. For example, if you are using your printer as a proofing device of a different output method, the absolute colorimetric rendering intent can allow you to proof what white will look like in the final print, producing a more accurate proof.

Relative Colorimetric

Relative colorimetric is virtually identical to absolute colorimetric, in that it keeps colors that are within the gamut of the destination space unchanged, but it changes the value of colors that are outside that gamut to the closest matching hue. The difference with relative colorimetric is that it maps white in the source profile to white in the destination profile, so you can achieve an accurate white that is as pure as possible based on the destination profile.

Perceptual

The perceptual rendering intent attempts to retain the perception of colors within the image by maintaining the relationships between those colors. When there are colors outside the gamut of the destination profile, the colors are effectively compressed to fit within that profile. All colors are shifted so that the relationships between colors are maintained, even if the actual colors shift slightly. This can cause a loss of saturation for many colors, which can be a problem for some images. It may be necessary when there are significant colors outside the gamut of the destination profile.

The perceptual rendering intent is best suited for photographic images that contain a significant number of colors that are outside the gamut of the destination profile, because the relationships between all colors in the image will be maintained, producing the best visual result. We'll talk more about how to determine if there are significant colors that are outside the destination color gamut in later chapters.

The Right Rendering Intent

There is no single best rendering intent. As a general rule for photographic images, I would only use the relative colorimetric or perceptual rendering intents. I tend to favor the perceptual intent for many images printed to inkjet printers, but for color space or profile-to-profile conversions, I generally prefer relative colorimetric. There is no right answer for every situation. I'll discuss the selection of rendering intents in more detail in later chapters, but understand that you may want to use a different rendering intent for different images even with the same output process.

Introducing Color Management

Photographers who do their own color image editing in the digital darkroom tend to have a love/hate relationship with color. They use color as their way to express creativity, and they love what they are able to express through color. They capture light on film or with digital sensors, but they do so much more than capture the light. They control that light through various lenses, filters, editing techniques, and printing methods. Nothing provides more satisfaction for the photographer than to use all the tools at their disposal to blend the palette of colors into a beautiful image.

And yet, as much as photographers love color, it seems nothing can frustrate them as easily as when colors don't match, or the color they expected isn't produced. Making the final print match what they envisioned when they clicked the shutter can present a tremendous challenge, particularly when you consider how many steps the image must pass through from capture to final print.

Color management provides the tools to ensure that the photographer is able to produce a print in the digital darkroom with predictable color. The tools to do so are available; however, learning what they are and how to properly use them seems to be one of the most misunderstood issues in digital imaging.

The complexity of color management is partly due to the number of components involved in working with digital images. Capture devices, such as digital cameras and film scanners, read the color in the original image. A monitor displays the color that is stored within the digital image file and provides the photographer with a visual representation on which to base adjustments to the image. Software settings determine the available colors in the image file. Finally, a printer attempts to reproduce the colors stored in the image file. Each of these components must be controlled to ensure proper color throughout the workflow.

As mentioned earlier in this chapter, color management revolves around profiles, and those profiles provide a common translation for color throughout the color-managed workflow (see Figure 1.18). You can build a custom profile for your digital camera or film scanner so that the colors recorded are as true to the scene that was photographed as possible. A custom profile also ensures that the color information in the image file is presented accurately by the monitor, because the values are translated to those that will produce accurate color on the monitor. This ensures that when you adjust the image, you are doing so based on an accurate view of the image and, therefore, the adjustments will truly produce the results you intend. Finally, when it comes time to print the image, a profile translates the numbers in the image file to the appropriate color information for the printer, resulting in a print that accurately reflects the color information in the image file.

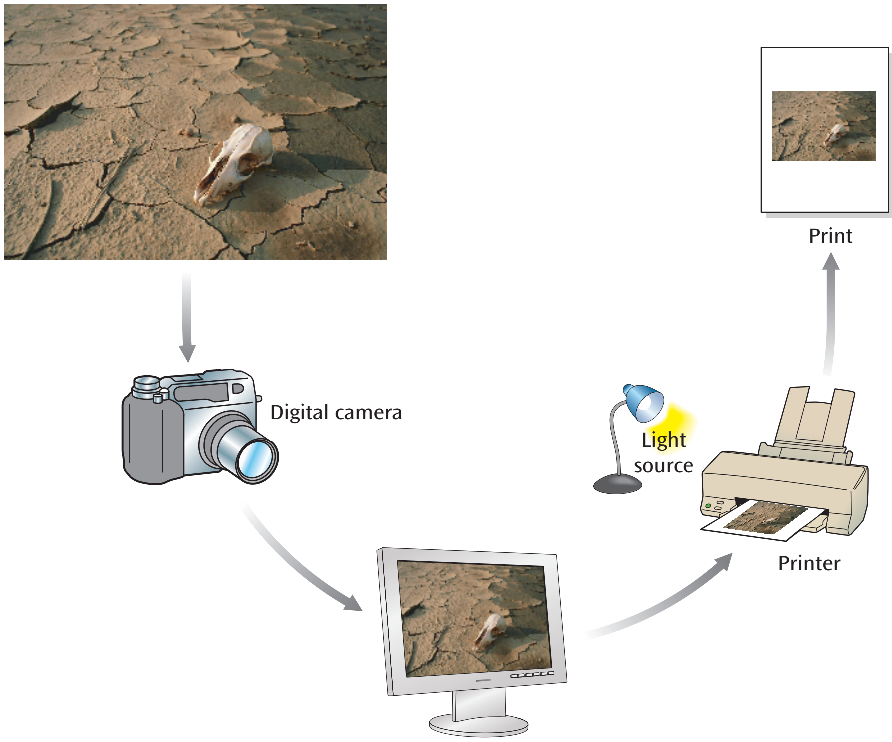

Figure 1.18 Color management provides a way to maintain the appearance of colors as the image moves through each device in the color-managed workflow.

In later chapters, I'll address each of the components of the digital darkroom in turn, providing solutions to achieve accurate and consistent color through the full process of optimizing your digital images.

Limitations of Color Management

I hate to burst your bubble, but color management isn't perfect. If you were hoping to buy this book and suddenly be able to produce a print that is indistinguishable from the display on your monitor, I'm going to have to bring you back down to earth.

Understanding the limitations of color management is just as important as understanding how color management works and what is involved in implementing a good color-managed workflow. We still don't have a complete understanding of all aspects of human vision, and new technology related to color is continually being developed. We have come a long way in understanding color and implementing color management solutions, but it isn't a perfect science.

Although you can achieve excellent results with a color management system, having realistic expectations is important. For example, the holy grail of color management for photographers is a print that matches the monitor. You need to understand that a print will never look exactly like the image on the monitor (see Figure 1.19). For one thing, both use very different mediums to present an image. A monitor emits light to present an image, while a print depends upon reflected light. Although the colors themselves can be very accurate, the experience of viewing each image is very different. The monitor is able to produce an image that is very luminous, by virtue of the fact that the image is actually composed of emitted light. Prints, on the other hand, will never be quite as luminous because they depend upon reflected light.

Figure 1.19 No matter how good your color management system, a print will never perfectly match the display on your monitor due to the differences in how the image is produced.

As I'll discuss in later chapters, evaluating your prints requires a certain amount of interpretation. If you understand the different properties of the different devices used in a color-managed workflow, you'll better understand the limitations of color management.

Of course, that doesn't mean you can't expect excellent results with a proper color-managed workflow. Quite the contrary, you can achieve incredibly accurate prints that exactly reflect what you intended the image to look like. Although it is important to understand that there are limitations in color management, you can get excellent results with an appropriate workflow.

One point I would emphasize is that while you may not get prints that exactly match what you see on the monitor in all respects, you can achieve predictable results. If you accept the inherent limitations involved, you can learn to interpret the results you achieve, learn to get the most accurate results, and ensure that you can get consistent, accurate, and predictable output for your images.