CHAPTER 21

Color Phenomena

The light receptors in cameras and in the eye respond more or less strongly to different wavelengths of light. Most cameras and most eyes have several different types of receptor, whose sensitivity to different wavelengths varies. Comparing the response of several types of sensor yields information about the distribution of energy with wavelength for the incoming light; this is color information. Color information can be used to remove shadows. The color of an object seen in an image depends on how the object was lit, but there are algorithms that can correct for this effect.

21.1 HUMAN COLOR PERCEPTION

The light coming out of sources or reflected from surfaces has more or less energy at different wavelengths, depending on the processes that produced the light. This distribution of energy with wavelength is sometimes called a *spectral energy density*; Figure 21.1 shows spectral energy densities for sunlight measured under a variety of different conditions. The visual system responds to light in a range of wavelengths from approximately 400nm to approximately 700nm. Light containing energy at just one wavelength looks deeply colored (these colors are known as *spectral colors*). The colors seen at different wavelengths have a set of conventional names, which originate with Isaac Newton (the sequence from 700nm to 400nm goes Red Orange Yellow Green Blue Indigo Violet, or **R**ichard **o**f **Y**ork **g**ot **b**listers **i**n **V**enice, although indigo is now frowned upon as a name because people typically cannot distinguish indigo from blue or violet). If the intensity is relatively uniform across the wavelengths, the light will look white.

Different kinds of color receptor in the human eye respond more or less strongly to light at different wavelengths, producing a signal that is interpreted as color by the human vision system. The precise interpretation of a particular light is a complex function of context; illumination, memory, object identity, and emotion can all play a part. The simplest question is to understand which spectral energy densities produce the same response from people under simple viewing conditions (Section 21.1.1). This yields a simple, linear theory of color matching that is accurate and extremely useful for describing colors. We sketch the mechanisms underlying the transduction of color in Section 21.1.2.

21.1.1 Color Matching

The simplest case of color perception is obtained when only two colors are in view on a black background. In a typical experiment, a subject sees a colored light—the *test light*—in one half of a split field (Figure 21.2). The subject can then adjust a mixture of lights in the other half to get it to match. The adjustments involve changing the intensity of some fixed number of *primaries* in the mixture.

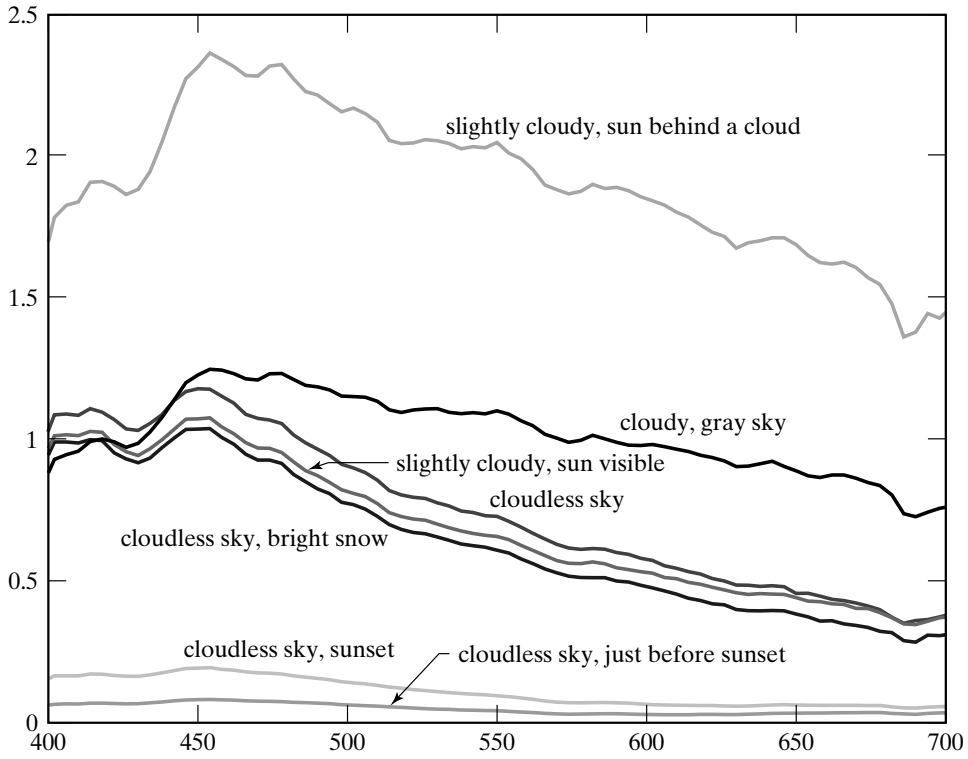

FIGURE 21.1: *Daylight has different amounts of power at different wavelengths. These plots show the spectral energy density of daylight measured at different times of day and under different conditions. The figure plots relative power against wavelength for wavelengths from 400 nm to 700 nm for a series of seven different daylight measurements, made by Jussi Parkkinen and Pertti Silfsten, of daylight illuminating a sample of barium sulphate (which gives a high reflectance white surface). At the foot of the plot, we show the names used for spectral colors of the relevant wavelengths. Plot from data obtainable at* `http://www.it.lut.fi/ip/research/color/database/database.html`.

Write T for the test light, an equals sign for a match, the weights—which are non-negative—as w_i, and the primaries P_i. A match can then be written in an algebraic form as

$$T = w_1 P_1 + w_2 P_2 + \ldots,$$

meaning that test light T matches the particular mixture of primaries given by $(w_1, w_2, \ldots)$. The situation is simplified if *subtractive matching* is allowed. In subtractive matching, the viewer can add some amount of some primaries to the *test light* instead of to the match. This can be written in algebraic form by allowing the weights in the expression above to be negative.

Under these conditions, most observers require only three primaries to match a test light. This phenomenon is known as the principle of *trichromacy*. However,

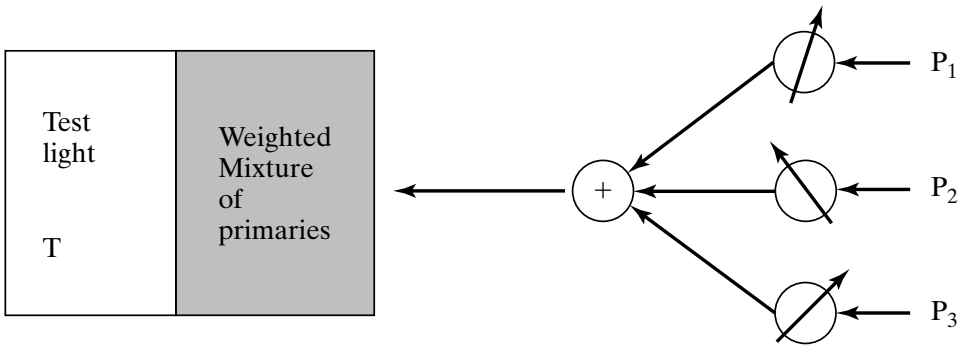

FIGURE 21.2: *Human perception of color can be studied by asking observers to mix colored lights to match a test light shown in a split field. The drawing shows the outline of such an experiment. The observer sees a test light T and can adjust the amount of each of three primaries in a mixture displayed next to the test light. The observer is asked to adjust the amounts so that the mixture looks the same as the test light. The mixture of primaries can be written as $w_1P_1 + w_2P_2 + w_3P_3$; if the mixture matches the test light, then we write $T = w_1P_1 + w_2P_2 + w_3P_3$. It is a remarkable fact that for most people three primaries are sufficient to achieve a match for many colors, and three primaries are sufficient for all colors if we allow subtractive matching (i.e., some amount of some of the primaries is mixed with the test light to achieve a match). Some people require fewer primaries. Furthermore, most people choose the same mixture weights to match a given test light.*

there are some caveats. First, subtractive matching must be allowed; second, the primaries must be independent, meaning that no mixture of two of the primaries may match a third. There is now clear evidence that trichromacy occurs because there are three distinct types of color transducer in the eye [**?**, **?**]. Given the same primaries and test light, most observers select the *same* mixture of primaries to match that test light, because most people have the same types of color receptor.

Matching is (to an accurate approximation) linear. This yields *Grassman's laws*. First, if we mix two test lights, then mixing the matches will match the result. Second, if two test lights can be matched with the same set of weights, then they will match each other. Finally, matching is linear: a test light with doubled intensity is matched by doubling the weights.

Given the same test light and set of primaries, most people use the same set of weights to match the test light. This, trichromacy, and Grassman's laws are about as true as any law covering biological systems can be. The exceptions include the following:

- people with too few kinds of color receptor as a result of genetic ill fortune (who may be able to match everything with fewer primaries);
- people with neural problems (who may display all sorts of effects, including a complete absence of the sensation of color);
- some elderly people (whose choice of weights differ from the norm because of

the development of macular pigment in the eye);

- very bright lights (whose hue and saturation look different from less bright versions of the same light);

- and very dark conditions (where the mechanism of color transduction is somewhat different than in brighter conditions).

21.1.2 Color Receptors

Human retinas contain two types of cell that are sensitive to light, differentiated by their shape. The light-sensitive region of a *cone* has a roughly conical shape, whereas that in a *rod* is roughly cylindrical. Cones largely dominate color vision. Cones are somewhat less sensitive to light than rods are, meaning that in low light, color vision is poor.

Trichromacy occurs because there are (usually!) three distinct types of cone in the eye that mediate color perception. Each of these types turns incident light into neural signals. The *principle of univariance* states that the activity of these cones is of one kind (i.e., they respond strongly or weakly, but do not signal the wavelength of the light falling on them). Univariance is a powerful idea because it gives us a good and simple model of human reaction to colored light: two lights will match if they produce the same receptor responses, *whatever their spectral energy densities*.

Write p_k for the response of the kth type of receptor, $\sigma_k(\lambda)$ for its sensitivity, $E(\lambda)$ for the light arriving at the receptor, and Λ for the range of visible wavelengths. We can obtain the overall response of a receptor by adding up the response to each separate wavelength in the incoming spectrum so that

$$p_k = \int_{\Lambda} \sigma_k(\lambda) E(\lambda) d\lambda.$$

Comparing color matching data for normal observers and those lacking one cone type yields the sensitivities of the three different kinds of cone to different wavelengths (Figure 21.3). The three types of cone are properly called *S cones*, *M cones*, and *L cones* (for their peak sensitivity being to short-, medium-, and long-wavelength light, respectively).

21.2 THE PHYSICS OF COLOR

Light sources can produce different amounts of light at different wavelengths, so incandescent lights look orange or yellow, and fluorescent lights look bluish. For most diffuse surfaces, albedo depends on wavelength, so that some wavelengths may be largely absorbed and others largely reflected. This means that most surfaces will look colored when lit by a white light. The light reflected from a colored surface is affected by both the color of the light falling on the surface, and by the surface. For example, a white surface lit by red light will reflect red light, and a red surface lit by white light will also reflect red light.

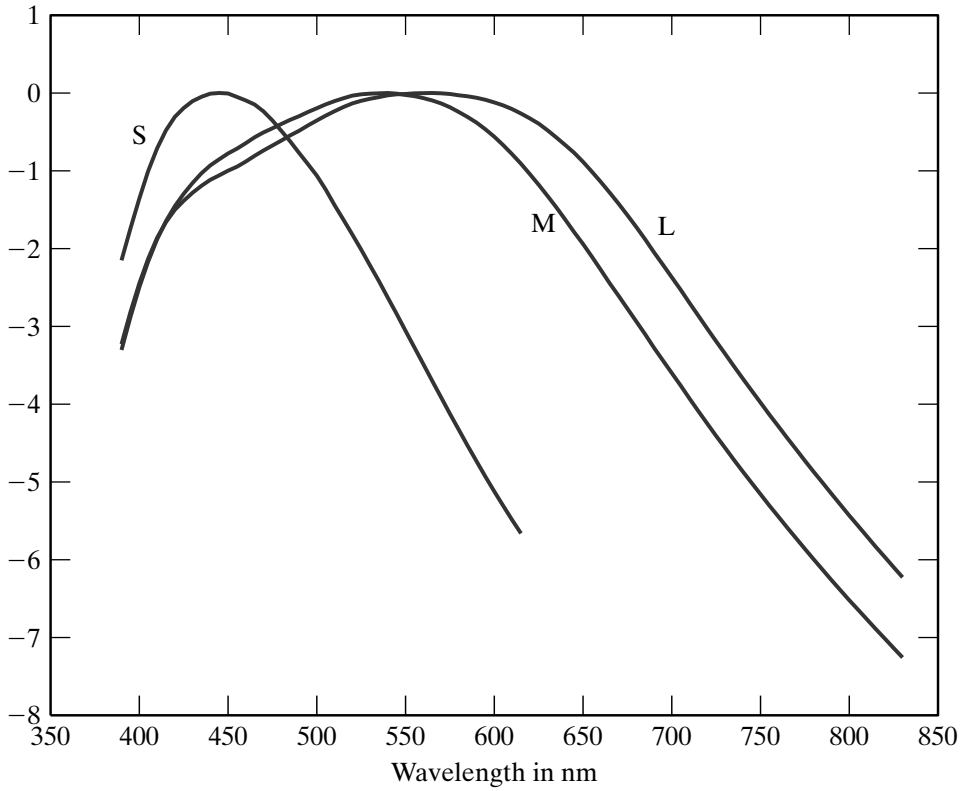

FIGURE 21.3: *There are three types of color receptor in the human eye, usually called* cones. *These receptors respond to all photons in the same way, but in different amounts. The figure shows the log of the relative spectral sensitivities of the three kinds of color receptor in the human eye, plotted against wavelength. On the wavelength axis, we have shown the color name usually associated with lights which contain energy only at that wavelength. The first two receptors—properly named the* long- *and* medium-wavelength *receptors—have peak sensitivities at quite similar wavelengths. The third receptor (*short-wavelength *receptor) has a different peak sensitivity. The response of a receptor to incoming light can be obtained by summing the product of the sensitivity and the spectral energy density of the light over all wavelengths. Notice that each receptor responds to quite a broad range of wavelengths. This means that human observers must perceive color by comparing the response of the receptors to one another, and that there must be many spectral energy densities that cannot be distinguished by humans.* Figures plotted from data disseminated by the Color and Vision Research Laboratories database, compiled by Andrew Stockman and Lindsey Sharpe, and available at `http://www.cvrl.org/`.

21.2.1 The Color of Light Sources

A patch of surface outdoors during the day is illuminated both by light that comes directly from the sun—usually called *daylight*—and by light from the sun that has

been scattered by the air (sometimes called *skylight* or *airlight*; the presence of clouds or snow can add other, important, phenomena). The color of daylight varies with time of day (Figure 21.1) and time of year.

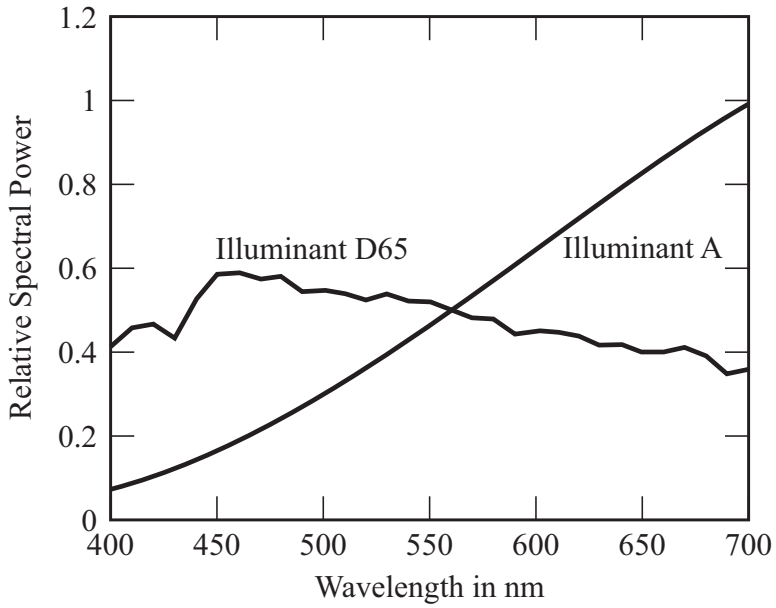

FIGURE 21.4: *There is a variety of illuminant models; the graph shows the relative spectral power distribution of two standard CIE models, illuminant A—which models the light from a 100W Tungsten filament light bulb, with color temperature 2800K—and illuminant D-65—which models daylight.* Figure plotted from data available at `http://www.cvrl.org/`.

For clear air, the intensity of radiation scattered by a unit volume depends on the fourth power of the frequency; this means that light of a long wavelength can travel much farther before being scattered than light of a short wavelength (this is known as *Rayleigh scattering*). This means that, when the sun is high in the sky, blue light is scattered out of the ray from the sun to the earth—meaning that the sun looks yellow—and can scatter from the sky into the eye—meaning that the sky looks blue. There are standard models of the spectral energy density of the sky at different times of day and latitude. Surprising effects occur when there are fine particles of dust in the sky (the larger particles cause much more complex scattering effects, usually modeled rather roughly by the *Mie scattering* model, described in **?** or in **?**).

Artificial Illumination

Typical artificial light sources are commonly of a small number of types:

- An *incandescent light* contains a metal filament that is heated to a high temperature. The spectrum roughly follows the black-body law (Section 21.2.1),

but the melting temperature of the element limits the color temperature of the light source, so the light has a reddish tinge.

- A *fluorescent light* works by generating high-speed electrons that strike gas within the bulb. The gas releases ultraviolet radiation, which causes phosphors coating the inside of the bulb to fluoresce. Typically the coating consists of three or four phosphors, which fluoresce in quite narrow ranges of wavelengths. Most fluorescent bulbs generate light with a bluish tinge, but some bulbs mimic natural daylight (Figure 21.5).

- In some bulbs, an arc is struck in an atmosphere consisting of gaseous metals and inert gases. Light is produced by electrons in metal atoms dropping from an excited state to a lower energy state. Typical of such lamps is strong radiation at a small number of wavelengths, which correspond to particular state transitions. The most common cases are *sodium arc lamps* and *mercury arc lamps*. Sodium arc lamps produce a yellow-orange light extremely efficiently and are quite commonly used for freeway lighting. Mercury arc lamps produce a blue-white light and are often used for security lighting.

- **TODO:** LED lights

Figure 21.5 shows a sample of spectra from different light bulbs.

Black Body Radiators

One useful abstraction is the *black body*, a body that reflects no light. A heated black body emits electromagnetic radiation. The spectral power distribution of this radiation depends only on the temperature of the body. If we write T for the temperature of the body in Kelvins, h for Planck's constant, k for Boltzmann's constant, c for the speed of light, and λ for the wavelength, we have

$$E(\lambda) \propto \frac{1}{\lambda^5} \frac{1}{(\exp(hc/k\lambda T) - 1)}.$$

This means that there is one parameter family of light colors corresponding to black body radiators—the parameter being the temperature—and so we can talk about the *color temperature* of a light source. This is the temperature of the black body that looks most similar. At relatively low temperatures, black bodies are red, passing through orange to a pale yellow-white to white as the temperature increases (Figure 21.10 shows this locus). When $hc \gg k\lambda T$, we have $1/(\exp(hc/k\lambda T) - 1) \approx \exp(-hc/k\lambda T)$, so

$$E(\lambda; T) = C\frac{\exp(-hc/k\lambda T)}{\lambda^5}$$

where C is the constant of proportionality; this model is somewhat easier to use than the exact model (Section 22.2.1).

21.2.2 The Color of Surfaces

The color of surfaces is a result of a large variety of mechanisms, including differential absorbtion at different wavelengths, refraction, diffraction, and bulk scattering

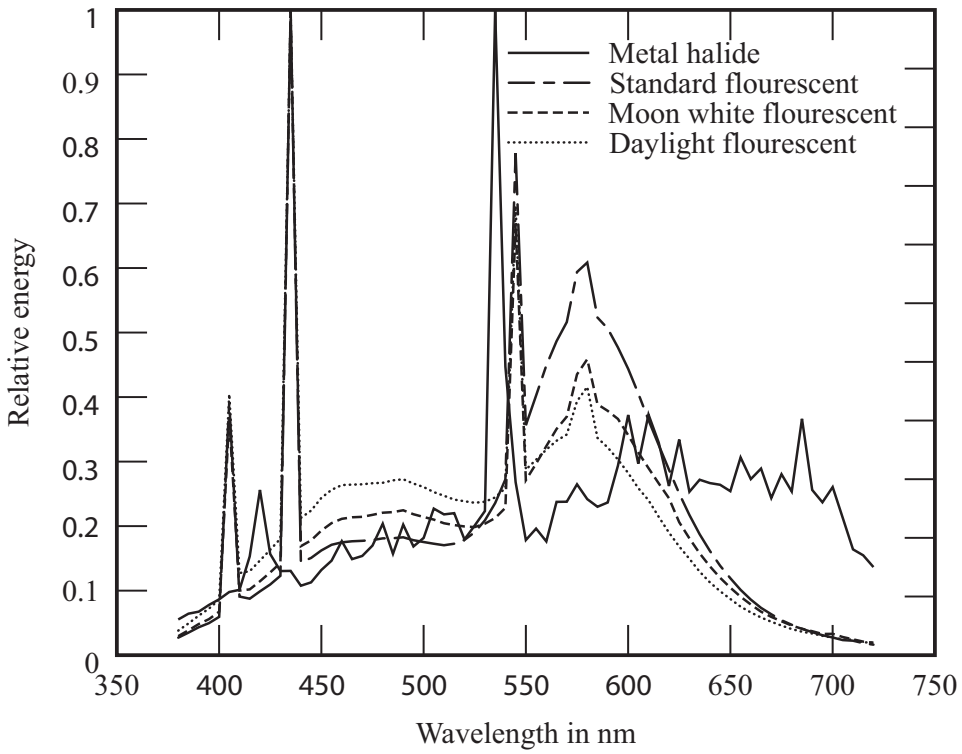

FIGURE 21.5: *The relative spectral power distribution of four different lamps from the Mitsubishi Electric Corporation. Note the bright, narrow bands that come from the flourescing phosphors in the fluorescent lamp.* The figure was plotted from data made available by the Coloring Info Pages at `http://www.colorpro.com/info/data/lamps.html`; the data was measured by Hiroaki Sugiura.

(for more details, see, for example **?**, **?**, **?**, or **?**). We can model surfaces as having a diffuse and a specular component, each of which has a wavelength-dependent albedo. The wavelength-dependent diffuse albedo is sometimes referred to as the *spectral reflectance* (sometimes abbreviated to *reflectance* or, less commonly, *spectral albedo*). Figures 21.6 and 21.7 show examples of spectral reflectances for a number of different natural objects.

There are two color regimes for specular reflection. If the surface is dielectric (i.e., does not conduct electricity), specularly reflected light tends to take the color of the light source. If the surface is a conductor, the specular albedo may depend quite strongly on wavelength, so that white light may result in colored specularities.

21.3 REPRESENTING COLOR

Describing colors accurately is a matter of great commercial importance. Many products are closely associated with specific colors—for example, the golden arches,

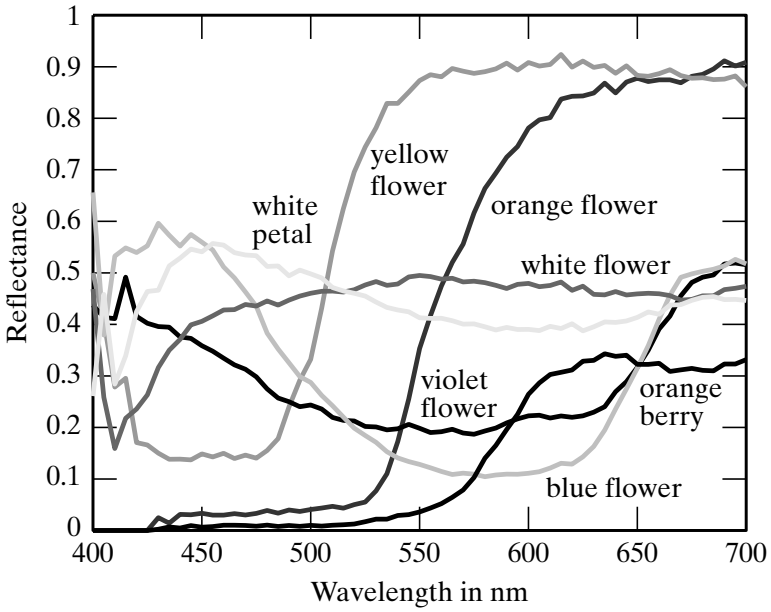

FIGURE 21.6: *Spectral albedoes for a variety of natural surfaces measured by Esa Koivisto, Department of Physics, University of Kuopio, Finland, plotted against wavelength in nanometers.* These figures were plotted from data available at `http://www.it.lut.fi/ip/research/color/database/database.html`.

the color of various popular computers, and the color of photographic film boxes—and manufacturers are willing to go to a great deal of trouble to ensure that different batches have the same color. This requires a standard system for talking about color. Simple color names are insufficient because relatively few people know many color names, and most people are willing to associate a large variety of colors with a given name. There are many linear and non-linear color spaces (**?** is a good reference). Generally, the choice of color space is driven by application. One important consideration is that, some color representations are more redundant than others. For example, the R, G and B layers in an RGB image are typically very similar, but a linear transformation can decorrelate these layers quite well (exercises). Redundancy is obviously a nuisance if one wishes to compress images. It is also a nuisance if one wishes to synthesize images, because the synthesis process must produce layers that are very, but not exactly, like each other. Another important consideration is consistency with perception. In some color spaces a small change in coordinates can result in a large change in perceived color. This is a problem if one wishes to control errors in color, for example, when mapping or synthesizing colors.

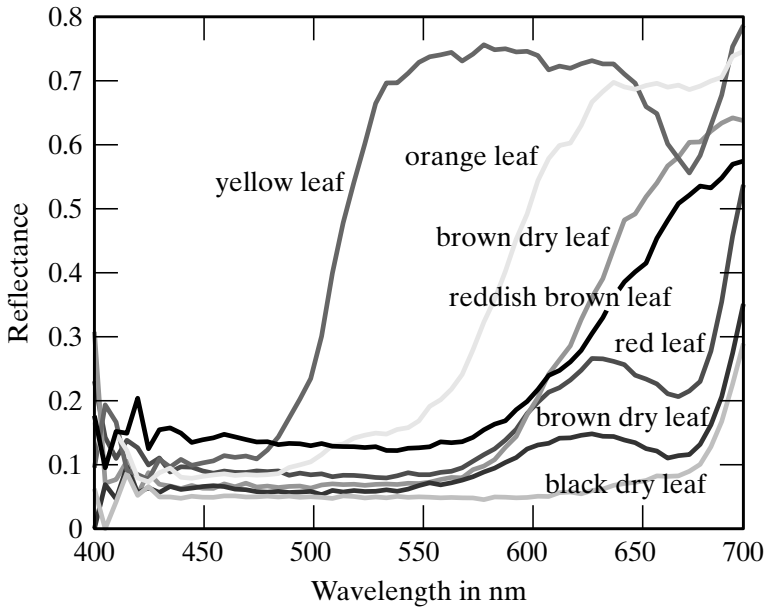

FIGURE 21.7: *Spectral albedoes for a variety of natural surfaces measured by Esa Koivisto, Department of Physics, University of Kuopio, Finland, plotted against wavelength in nanometers.* These figures were plotted from data available at `http://www.it.lut.fi/ip/research/color/database/database.html`.

21.3.1 Additive Linear Color Spaces

There is a natural mechanism for representing color: agree on a standard set of primaries, and then describe any colored light by the three values of weights that people would use to match the light using those primaries. This approach extends to give a representation for surface colors as well if we use a standard light for illuminating the surface (and if the surfaces are equally clean, etc.). Performing a matching experiment each time we wish to describe a color can be practical (paint stores will mix paint to match a flake, for example), but a simpler procedure is available.

A *linear color space* is defined by a choice of primaries P_1, P_2, and P_3. These may not be physically realizable. One then obtains a set of *color matching functions* from the primaries by experiment. The color matching functions $f_1(\lambda)$, $f_2(\lambda)$, and $f_3(\lambda)$ have the property that, if a source $S(\lambda)$ is matched by $w_1P_1 + w_2P_2 + w_3P_3$, then

$$w_i = \int f_i(\lambda)S(\lambda)d\lambda.$$

There is a form of duality between primaries and color matching functions, so one can obtain a linear color space by constructing the color matching functions and then looking for primaries that produce these color matching functions. A variety of different systems have been standardized by the CIE (the *commission international d'éclairage*, which exists to create standards for such things).

The *CIE XYZ color space* is one quite popular standard. The color matching functions were chosen to be everywhere positive, so that the coordinates of any real light are always positive. It is not possible to obtain CIE X, Y, or Z primaries because for some wavelengths the value of their pectral energy density is negative. However, given color matching functions alone, one can specify the XYZ coordinates of a color and hence describe it.

Linear color spaces allow a number of useful graphical constructions that are more difficult to draw in three dimensions than in two, so it is common to intersect the XYZ space with the plane $X + Y + Z = 1$ (as shown in Figure 21.8) and draw the resulting figure using coordinates

$$(x, y) = \left(\frac{X}{X + Y + Z}, \frac{Y}{X + Y + Z} \right).$$

This space, which is often referred to as the *CIE xy color space* is shown in Figure 21.10. CIE xy is widely used in vision and graphics textbooks and in some applications, but is usually regarded by professional colorimetrists as out of date.

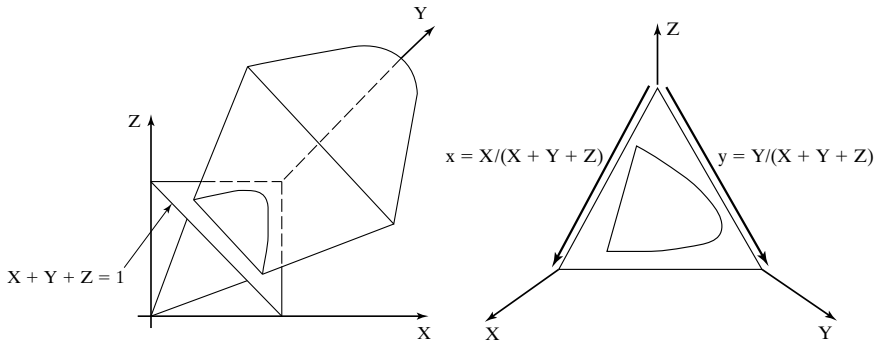

FIGURE 21.8: *The volume of all visible colors in the CIE XYZ coordinate space is a cone whose vertex is at the origin. Usually it is easier to suppress the brightness of a color, which we can do because, to a good approximation, perception of color is linear, and we do this by intersecting the cone with the plane* $X + Y + Z = 1$ *to get the CIE xy space shown in Figure 21.10.*

The *RGB color space* is a linear color space that formally uses single wavelength primaries (645.16 nm for R, 526.32 nm for G, and 444.44 nm for B; see Figure **??**). Informally, RGB uses whatever phosphors a monitor has as primaries. Available colors are usually represented as a unit cube—usually called the *RGB cube*—whose edges represent the R, G, and B weights. The cube is drawn in Figure 21.11.

The *opponent color space* is a linear color space derived from RGB. There is evidence that there are three kinds of color system in primates (e.g., see **?**; **?**). The oldest responds to intensity (i.e., light-dark comparisons). A more recent, but still old, color system compares blue with yellow. The most recent color system

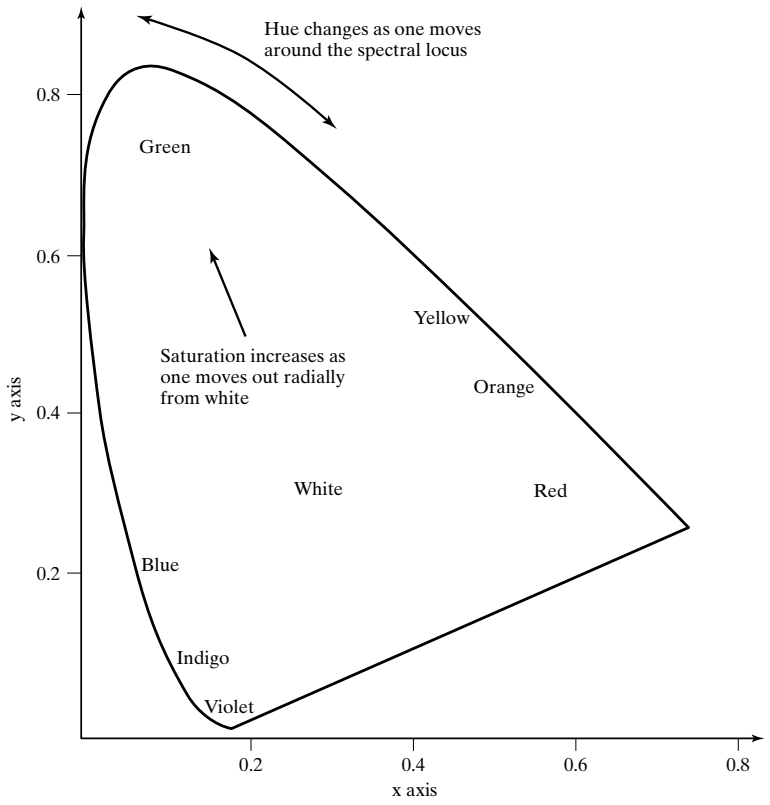

FIGURE 21.9: *The figure shows a constant brightness section of the standard 1931 standard CIE xy color space, with color names marked on the diagram. Generally, colors that lie farther away from the neutral point are more saturated—the difference between deep red and pale pink—and hue—the difference between green and red—as one moves around the neutral point.*

compares red with green. In some applications, it is useful to use a comparable representation. This can be obtained from RGB coordinates using $I = (R+G+B)/3$ for intensity, $(B - (R + G)/2)/I$ for the blue-yellow comparison (sometimes called B-Y), and $(R - G)/I$ for the red-green comparison (sometimes called R-G). Notice that B-Y (resp. R-G) is positive for strongly blue (resp. red) colors and negative for strongly yellow (resp. green) colors, and is intensity independent.

There are two useful constructions that work in linear color spaces, but are most commonly applied in CIE xy. First, because the color spaces are linear, and color matching is linear, all colors that can be obtained by mixing two primaries A and B lie on the line segment joining them plotted on the color space. Second, all colors that can be obtained by mixing three primaries A, B, and C lie in the triangle formed by the three primaries plotted on the color space. Typically, we use this construction to determine the set of colors (or *gamut*) that a set of monitor phosphors can display.

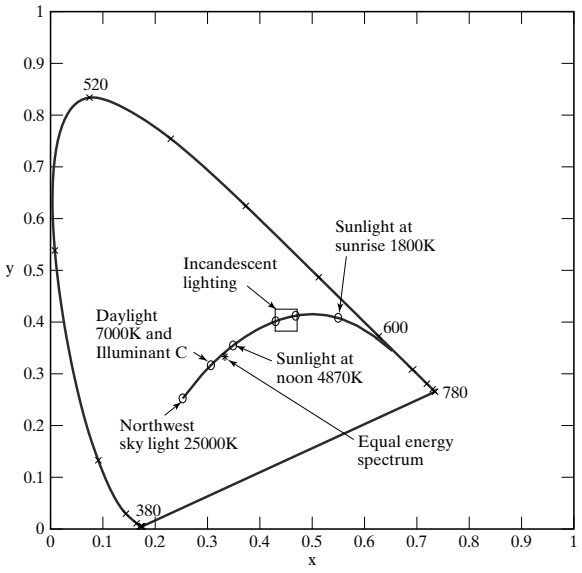

FIGURE 21.10: *The figure shows a constant brightness section of the standard 1931 standard CIE xy color space. This space has two coordinate axes. The curved boundary of the figure is often known as the* spectral locus*; it represents the colors experienced when lights of a single wavelength are viewed. The figure shows a locus of colors due to black-body radiators at different temperatures and a locus of different sky colors. Near the center of the diagram is the neutral point, the color whose weights are equal for all three primaries. CIE selected the primaries so that this light appears achromatic. Generally, colors that lie farther away from the neutral point are more saturated—the difference between deep red and pale pink—and hue—the difference between green and red—as one moves around the neutral point.*

21.3.2 Subtractive Mixing and Inks

Intuition from one's finger-painting days suggests that the primary colors should be red, yellow, and blue, and that yellow and blue mix to make green. The reason this intuition doesn't apply to monitors is that paints involve pigments—which mix subtractively—rather than lights. Pigments can behave in quite complex ways, but the simplest model is that pigments remove color from incident light, which is reflected from paper. Thus, red ink is really a dye that absorbs green and blue light—incident red light passes through this dye and is reflected from the paper. This is *subtractive color mixing*.

Color spaces for this kind of mixing can be quite complicated. In the simplest case, mixing is linear (or reasonably close to linear), and the *CMY space* applies. In this space, there are three primaries: *cyan* (a blue-green color), *magenta* (a purplish color), and *yellow*. These primaries should be thought of as subtracting a light primary from white light; cyan is $W - R$ (white − red); magenta is $W - G$ (white − green), and yellow is $W - B$ (white − blue). Now the appearance of mixtures can be evaluated by reference to the RGB color space. For example, cyan

and magenta mixed give

$$(W - R) + (W - G) = R + G + B - R - G = B,$$

that is, blue. Notice that $W + W = W$ because we assume that ink cannot cause paper to reflect more light than it does when uninked. Practical printing devices use at least four inks (cyan, magenta, yellow, and black) because mixing color inks leads to a poor black, it is difficult to ensure good enough registration between the three color inks to avoid colored haloes around text, and color inks tend to be more expensive than black inks. One reason that fingerpainting is hard is that the color resulting from mixing paints can be quite hard to predict. This is because the outcome depends very strongly on details such as the specific pigment in the paint, the size of pigment particles, the medium in which the pigment is suspended, the care put into stirring the mixture, and similar parameters; usually, we do not have enough detailed information to use a full physical model of these effects. A useful study of this difficult topic is [**?**].

21.3.3 Non-linear Color Spaces

The coordinates of a color in a linear space may not necessarily encode properties that are common in language or are important in applications. Useful color terms include: *hue*, the property of a color that varies in passing from red to green; *saturation*, the property of a color that varies in passing from red to pink; and *brightness* (sometimes called *lightness* or *value*, the property that varies in passing from black to white. Another difficulty with linear color spaces is that the individual coordinates do not capture human intuitions about the topology of colors; it is a common intuition that hues form a circle, in the sense that hue changes from red through orange to yellow, and then green, and from there to cyan, blue, purple, and then red again. Another way to think of this is to picture local hue relations: red is next to purple and orange; orange is next to red and yellow; yellow is next to orange and green; green is next to yellow and cyan; cyan is next to green and blue; blue is next to cyan and purple; and purple is next to blue and red. Each of these local relations works, and globally they can be modeled by laying hues out in a circle. This means that no individual coordinate of a linear color space can model hue, because that coordinate has a maximum value that is far away from the minimum value.

Applying a non-linear transformation to the RGB space can produce a color space that respects these relations. The *HSV space* (for hue, saturation, and value), is obtained by looking down the center axis of the RGB cube. Because RGB is a linear space, brightness—called *value* in HSV—varies with scale out from the origin. We can flatten the RGB cube to get a 2D space of constant value and for neatness deform it to be a hexagon. This gets the structure shown in Figure 21.11, where hue is given by an angle that changes as one goes round the neutral point and saturation changes as one moves away from the neutral point.

In some applications, it is important to know whether a color difference would be noticeable to a human viewer. One can determine *just noticeable differences* by modifying a color shown to observers until they can only just tell it has changed in a comparison with the original color. With an appropriate choice of non-linear

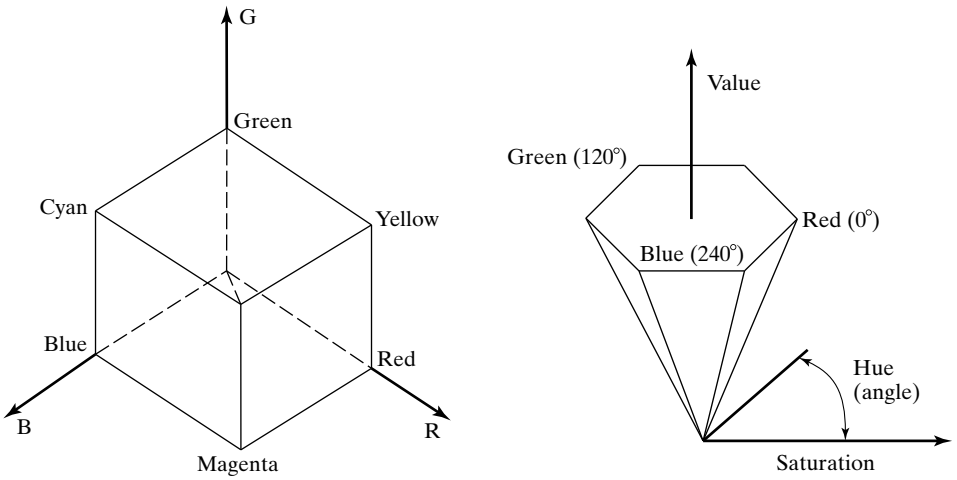

FIGURE 21.11: *On the* **left**, *we see the RGB cube; this is the space of all colors that can be obtained by combining three primaries (R, G, and B—usually defined by the color response of a monitor) with weights between zero and one. It is common to view this cube along its neutral axis—the axis from the origin to the point (1, 1, 1)—to see a hexagon. This hexagon codes hue (the property that changes as a color is changed from green to red) as an angle, which is intuitively satisfying. On the* **right**, *we see a cone obtained from this cross-section, where the distance along a generator of the cone gives the value (or brightness) of the color, the angle around the cone gives the hue, and the distance out gives the saturation of the color.*

transformation applied to linear color coordinates, one can find a *uniform color space*. In such a space, if the distance in coordinate space is below some threshold, a human observer would not be able to tell the colors apart.

A uniform space can be obtained from CIE XYZ using a projective transformation to obtain the *CIE u′v′ space CIE u'v' space*. The coordinates are:

$$(u', v') = \left(\frac{4X}{X + 15Y + 3Z}, \frac{9Y}{X + 15Y + 3Z} \right).$$

Generally, the distance between coordinates in $u\prime$, $v\prime$ space is a fair indicator of the significance of the difference between two colors. Of course, this omits differences in brightness.

CIE LAB is now almost universally the most popular uniform color space. Coordinates of a color in LAB are obtained as a non-linear mapping of the XYZ

coordinates:

$$L^* = 116\left(\frac{Y}{Y_n}\right)^{\frac{1}{3}} - 16$$
$$a^* = 500\left[\left(\frac{X}{X_n}\right)^{\frac{1}{3}} - \left(\frac{Y}{Y_n}\right)^{\frac{1}{3}}\right]$$
$$b^* = 200\left[\left(\frac{Y}{Y_n}\right)^{\frac{1}{3}} - \left(\frac{Z}{Z_n}\right)^{\frac{1}{3}}\right]$$

Here X_n, Y_n, and Z_n are the X, Y, and Z coordinates of a reference white patch. The reason to care about the LAB space is that it is substantially uniform. In some problems, it is important to understand how different two colors will look *to a human observer*, and differences in LAB coordinates give a good guide.

CHAPTER 22

Using Color Models

22.1 A MODEL OF IMAGE COLOR

Assume that an image pixel is the image of some surface patch. Many phenomena affect the color of this pixel. The main effects are: the camera response to illumination; the choice of camera receptors; the amount of light that arrives at the surface; the color of light arriving at the surface; the dependence of the diffuse albedo on wavelength; and specular components. We have already dealt with the camera response (Section 20.2.1) and we will assume that the camera is linear, or has been radiometrically calibrated. A quite simple model can be used to separate the other effects.

Assume that the surfaces that we are dealing with can be described by the diffuse+specular model. Write $\mathbf{x}$ for a point, λ for wavelength, $E(\mathbf{x}, \lambda)$ for the spectral energy density of the light leaving a surface, $\rho(\mathbf{x}, \lambda)$ for the albedo of a surface as a function of wavelength and position, $S_d(\mathbf{x}, \lambda)$ for the spectral energy density of the light source (which may vary with position; for example, the intensity might change), and $S_i(\mathbf{x}, \lambda)$ for the spectral energy density of interreflected light. Then we have that:

$$\begin{aligned} E(\mathbf{x}, \lambda) &= [\text{diffuse term}] + (\text{specular term}) \\ &= [(\text{direct term}) + (\text{interreflected term})] + (\text{specular term}) \\ &= (\rho(\mathbf{x}, \lambda)(\text{geometric term}))[(S_d(\mathbf{x}, \lambda) + S_i(\mathbf{x}, \lambda))] + (\text{specular term}). \end{aligned}$$

The geometric terms represent how intensity is affected by surface normal. Notice that the diffuse term is affected both by the color of the surface and by the color of the light (examples in Figures 22.1 and 22.2).

Because the camera is linear, the pixel value at $\mathbf{x}$ is a sum of terms corresponding to each of the terms in $E(\vec{x}, \lambda)$. Write $\mathbf{d}(\mathbf{x})$ for the color taken by a flat patch facing the light source at $\mathbf{x}$ with the same albedo as the actual patch there, $g(\mathbf{x})$ for a geometric term (explained below), $\mathbf{i}(\mathbf{x})$ for the contribution of the interreflected term, $\mathbf{s}(\mathbf{x})$ for the unit intensity color of the specular term, and $g_s(\mathbf{x})$ for a geometric term (explained below). Then we have:

$$\begin{aligned} \mathbf{C}(\mathbf{x}) &= [(\text{direct term}) + (\text{interreflected term})] + (\text{specular term}) \\ &= g_d(\mathbf{x})\mathbf{d}(\mathbf{x}) + \mathbf{i}(\mathbf{x}) + g_s(\mathbf{x})\mathbf{s}(\mathbf{x}). \end{aligned}$$

Generally, to work with this model, we ignore $\mathbf{i}(\mathbf{x})$; we identify and remove specularities, using the methods of Section **??**, and so assume that $\mathbf{C}(\mathbf{x}) = g_d(\mathbf{x})\mathbf{d}(\mathbf{x})$.

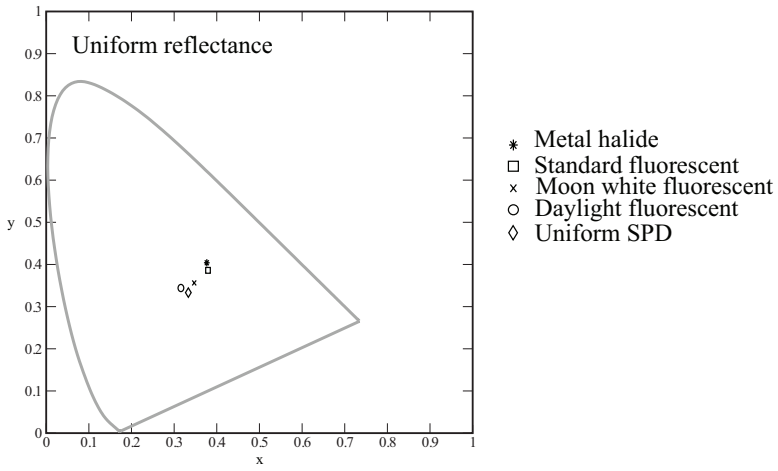

FIGURE 22.1: *Light sources can have quite widely varying colors. This figure shows the color of the four light sources of Figure 21.5, compared with the color of a uniform spectral power distribution, plotted in CIE x, y coordinates.*

22.1.1 The Diffuse Term

There are two diffuse components. One, $\mathbf{i}(\mathbf{x})$, is due to interreflections. Interreflections can be a significant source of colored light. If a large colored surface reflects light onto another surface, that surface's color can change quite substantially. This is an effect that people find hard to see, but which is usually fairly easy to spot in photographs. There are no successful models for removing these color shifts, most likely because they can be very hard to predict. This is because many different surface reflectances can have the same color, so that two surfaces with the same color (but different reflectances) can have quite differently colored interreflections. The interreflection term is often small, and usually is simply ignored.

Ignoring the interreflected component, the diffuse term is

$$g_d(\mathbf{x})\mathbf{d}(\mathbf{x}).$$

Here $\mathbf{d}(\mathbf{x})$ is the *image* color of an equivalent *flat* surface facing the light source and viewed under the same light. The geometric term, $g_d(\mathbf{x})$, varies relatively slowly over space and accounts for the change in brightness due to the orientation of the surface.

We can model the dependence of $\mathbf{d}(\mathbf{x})$ on the light and on the surface by assuming we are viewing flat, diffuse surfaces, illuminated from infinitely far behind the camera. In this case, there will be no effects due to specularities or to surface orientation. The color of light arriving at the camera will be determined by two factors: first, the wavelength-dependent albedo of the surface that the light is leaving; and second, the wavelength-dependent intensity of the light falling on that surface. If a patch of perfectly diffuse surface with diffuse albedo $\rho(\lambda)$ is illuminated by a light whose spectrum is $S(\lambda)$, the spectrum of the reflected light is $\rho(\lambda)S(\lambda)$. Assume the camera has linear photoreceptors, and the k'th type of photoreceptor

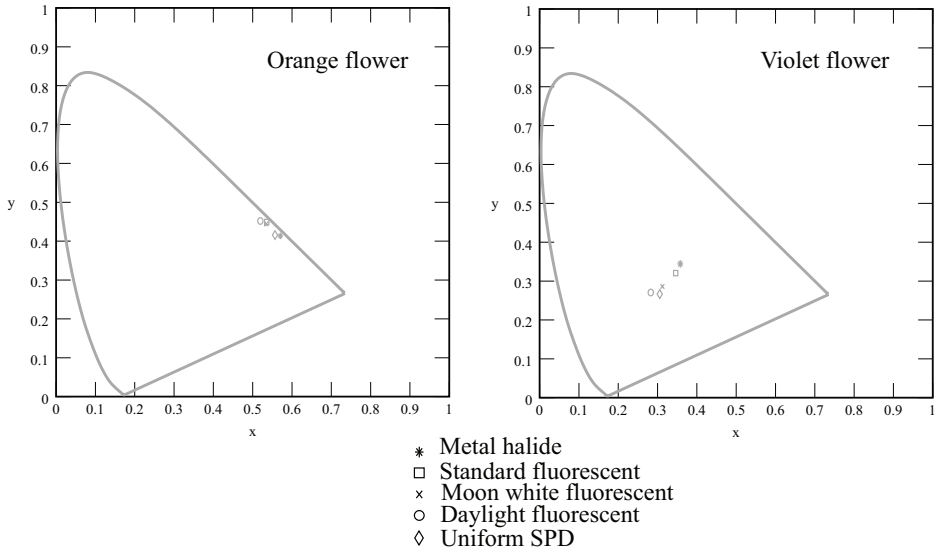

FIGURE 22.2: *The color of a light source affects the color of surfaces lit by the source. The different colors obtained by lighting the violet flower of Figure 21.6 (***left***) and the orange flower of Figure 21.6 (***right***) with the four light sources of Figure 21.5.*

has sensitivity $\sigma_k(\lambda)$. If a linear photoreceptor of the kth type sees this surface patch, its response is:

$$p_k = \int_{\Lambda} \sigma_k(\lambda)\rho(\lambda)S(\lambda)d\lambda,$$

where Λ is the range of all relevant wavelengths.

The main engineering parameter here is the photoreceptor sensitivities $\sigma_k(\lambda)$. For some applications such as shadow removal (Section 22.2.1), it can be quite helpful to have photoreceptor sensitivities that are "narrow-band" (i.e., the photoreceptors respond to only one wavelength). Usually, the only practical methods to change the photoreceptor sensitivities are to either put colored filters in front of the camera or to use a different camera. Using a different camera doesn't work particularly well, because manufacturers try to have sensitivities that are reasonably compatible with human receptor sensitivities. They do this so that cameras give about the same responses to colored lights that people do; as a result, cameras tend to have quite similar receptor sensitivities. There are three ways to proceed: install narrow-band filters in front of the lens (difficult to do and seldom justified); apply a transformation to the receptor outputs that makes them behave more like narrow-band receptors (often helpful, if the necessary data are available, **?**;**?**); or assume that they are narrow-band receptors and tolerate any errors that result (generally quite successful).

22.1.2 The Specular Term

The specular component will have a characteristic color, and its intensity will change with position. We can model the specular component as

$$g_s(\mathbf{x})\mathbf{s}(\mathbf{x}),$$

where $\mathbf{s}(\mathbf{x})$ is the unit intensity *image* color of the specular reflection at that pixel, and $g_s(\mathbf{x})$ is a term that varies from pixel to pixel, and models the amount of energy specularly reflected. We expect $g_s(\mathbf{x})$ to be zero at most points, and large at some points.

The color $\mathbf{s}(\mathbf{x})$ of the specular component depends on the material. Generally, metal surfaces have a specular component that is wavelength dependent and so takes on a characteristic color that depends on the metal (gold is yellow, copper is orange, platinum is white, and osmium is blue or purple). Surfaces that do not conduct—*dielectric surfaces*— have a specular component that is independent of wavelength (e.g., the specularities on a shiny plastic object are the color of the light). Section **??** describes how these properties can be used to find specularities, and to find image regions corresponding to metal or plastic objects.

22.2 INFERENCE FROM COLOR

Our color model supports a variety of inferences. Here we show methods to remove shadows (Section 22.2.1) and to infer surface color (Section 22.2.2).

22.2.1 Shadow Removal Using Color

Lightness methods make the assumption that "fast" edges in images are due to changes in albedo (Section 20.2.2). This assumption is usable, but fails badly at shadows, particularly shadows in sunlight outdoors (Figure 22.4), where there can be a large and fast change of image brightness. People usually are not fooled into believing that a shadow is a patch of dark surface, so must have some method to identify shadow edges. Home users often like editing and improving photographs, and programs that could remove shadows from images would be valuable. A shadow removal program would work something like a lightness method: find all edges, identify the shadow edges, remove those, and then integrate to get the picture back.

There are some cues for finding shadow edges that seem natural, but don't work well. One might assume that shadow edges have very large dynamic range (which albedo edges can't have; see Section 20.1.1), but this is not always the case. One might assume that, at a shadow edge, there was a change in brightness but not in color. It turns out that this is not the case for outdoor shadows, because the lit region is illuminated by yellowish sunlight, and the shadowed region is illuminated by bluish light from the sky, or sometimes by interreflected light from buildings, and so on. However, a really useful cue can be obtained by modelling the different light sources.

We assume that light sources are black bodies, so that their spectral energy density is a function of temperature. We assume that surfaces are diffuse. We use the simplified black-body model of Section 21.2.1, where, writing T for the

temperature of the body in Kelvins, h for Planck's constant, k for Boltzmann's constant, c for the speed of light, and λ for the wavelength, we have

$$E(\lambda; T) = C\frac{\exp(-hc/k\lambda T)}{\lambda^5}$$

(C is some constant of proportionality). Now assume that the color receptors each respond only at one wavelength, which we write λ_k for the k'th receptor, so that $\sigma_k(\lambda) = \delta(\lambda - \lambda_k)$. If we view a surface with spectral albedo $\rho(\lambda)$ illuminated by one of these sources at temperature T, the response of the j'th receptor will be

$$r_j = \int \sigma_j(\lambda)\rho(\lambda)K\frac{\exp(-hc/k\lambda T)}{\lambda^5}d\lambda = K\rho(\lambda_j)\frac{\exp(-hc/k\lambda_j T)}{\lambda_j^5}.$$

We can form a color space that is very well behaved by taking $c_1 = \log(r_1/r_3)$, $c_2 = \log(r_2/r_3)$, because

$$\begin{pmatrix} c_1 \\ c_2 \end{pmatrix} = \begin{pmatrix} a_1 \\ a_2 \end{pmatrix} + \frac{1}{T}\begin{pmatrix} b_1 \\ b_2 \end{pmatrix}$$

where $a_1 = \log\rho(\lambda_1) - \log\rho(\lambda_3) + 5\log\lambda_3 - 5\log\lambda_1$ and $b_1 = (hc/k)(1/\lambda_3 - 1/\lambda_1)$ (and a_2, b_2 follow). Notice that, when one changes the color temperature of the source, the (c_1, c_2) coordinates move along a straight line. The direction of the line depends on the sensor, *but not on the surface*. Call this direction the **color temperature direction**. The intercept of the line depends on the surface.

Now consider a world of colored surfaces, and map the image colors to this space. There is a family of parallel lines in this space, whose direction is the color temperature direction. Different surfaces may map to different lines. If we change the color temperature of the illuminant, then each color in this space will move along the color temperature direction, but colors will not move from line to line. We now represent a surface color by its line. For example, we could construct a line through the origin that is perpendicular to color temperature direction, then represent a surface color by distance along this line (Figure 22.3). We can represent each pixel in the image in this space, and in this representation the color image becomes a gray-level image, *where the gray level does not change inside shadows* (because a shadow region just has a different color temperature to the non-shadowed region). **?** calls this the *invariant image*. Any edge that appears in the image but not in the invariant image is a shadow edge, so we can now apply our original formula: find all edges, identify the shadow edges, remove those, and then integrate to get the picture back.

Of course, under practical circumstances, usually we do not know enough about the sensors to evaluate the as and bs that define this family of lines, so we cannot get the invariant image directly. However, we can infer a direction in (c_1, c_2) space that is a good estimate by a form of entropy reasoning. We must choose a color temperature direction. Assume the world is rich in differently colored surfaces. Now consider two surfaces S_1 and S_2. If $\mathbf{c}_1$ (the (c_1, c_2) values for S_1) and $\mathbf{c}_2$ are such that $\mathbf{c}_1 - \mathbf{c}_2$ is parallel to the color temperature direction, we can choose T_1 and T_2 so that S_1 viewed under light with color temperature T_1 will look the same as S_2 viewed under light with color temperature T_2. We expect this to be uncommon,

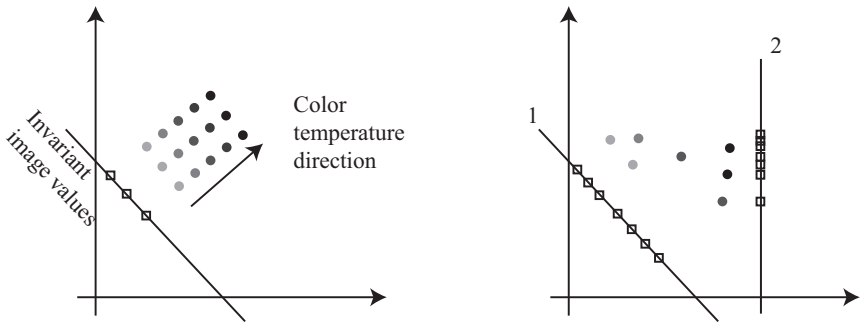

FIGURE 22.3: *Changing the color temperature of the light under which a surface is viewed moves the* (c_1, c_2) *coordinates of that surface along the color temperature direction (***left***; the different gray patches represent the same surface under different lights). If we now project the coordinates along the* (c_1, c_2) *direction onto some line, we obtain a value that doesn't change when the illuminant color temperature changes. This is the invariant value for that pixel. Generally, we do not know enough about the imaging system to estimate the color temperature direction. However, we expect to see many different surfaces in each scene; this suggests that the right choice of color temperature direction on the* **right** *is 1 (where there are many different types of surface) rather than 2 (where the range of invariant values is small).*

because surfaces tend not to mimic one another in this way. This means we expect that colors will tend to spread out when we project along a good estimate of the color temperature direction. A reasonable measure of this spreading out is the *entropy* of the histogram of projected colors. We can now estimate the invariant image, without knowing anything about the sensor. We search directions in (c_1, c_2) space, projecting all the image colors along that direction; our estimate of the color temperature direction is the one where this projection yields the largest entropy. From this we can compute the invariant image, and so apply our shadow removal strategy above. In practice, the method works well, though great care is required with the integration procedure to get the best results (Figure 22.4).

22.2.2 Color Constancy: Surface Color from Image Color

In our model, the image color depends on both light color and on surface color. If we light a green surface with white light, we get a green image; if we light a white surface with a green light, we also get a green image. This makes it difficult to name surface colors from pictures. We would like to have an algorithm that can take an image, discount the effect of the light, and report the actual color of the surface being viewed.

This process is called *color constancy*. Humans have some form of color constancy algorithm. People are often unaware of this, and inexperienced photographers are sometimes surprised that a scene photographed indoors under fluorescent lights has a blue cast, whereas the same scene photographed outdoors may have a warm orange cast. The simple linear models of Section 21.3 can predict the color an

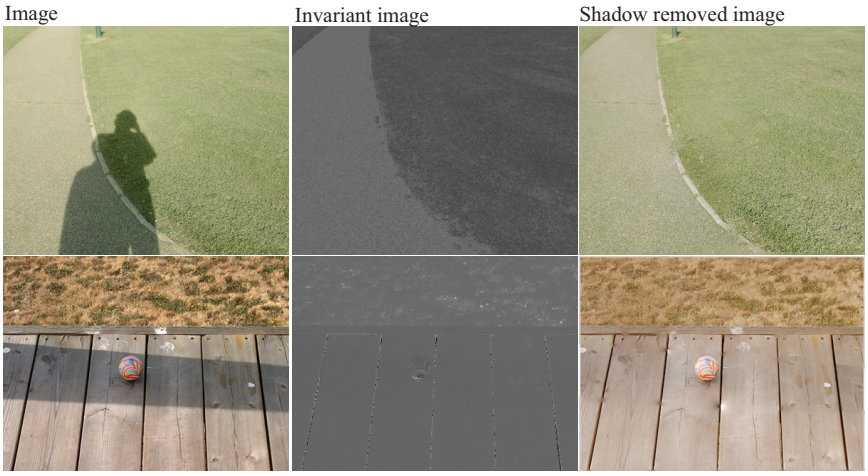

FIGURE 22.4: *The invariant of the text and of Figure 22.3 does not change value when a surface is shadowed. Finlayson* et al. *use this to build a shadow removal system that works by (a) taking image edges; (b) forming an invariant image; then (c) using that invariant image to identify shadow edges; and finally (d) integrating only non-shadow edges to form the result. The results are quite convincing.*

observer will perceive when shown an isolated spot of light of a given power spectral distribution. But if this spot is part of a larger, more complex scene, these models can give wildly inaccurate predictions. This is because the human color constancy algorithm uses various forms of scene information to decide what color to report. Demonstrations by **?**, which are illustrated in Figure 22.5, give convincing examples of this effect. It is surprisingly difficult to predict what colors a human will see in a complex scene (**?**; **?**; [**?**]; [**?**]; [**?**]). This is one of the many difficulties that make it hard to produce really good color reproduction systems.

Human color constancy is not perfectly accurate, and people can choose to disregard information from their color constancy system. As a result, people can often report:

- the color a surface would have in white light (often called *surface color*);
- the color of the light arriving at the eye (a useful skill that allows artists to paint surfaces illuminated by colored lighting); and
- the color of the light falling on the surface.

The model of image color in Section 22.1 is

$$\mathbf{C}(\mathbf{x}) = g_d(\mathbf{x})\mathbf{d}(\mathbf{x}) + g_s(\mathbf{x})\mathbf{s}(\mathbf{x}) + \mathbf{i}(\mathbf{x}).$$

We decided to ignore the interreflection term $\mathbf{i}(\mathbf{x})$. In principle, we could use the methods of Section **??** to generate new images without specularities. This brings us to the term $g_d(\mathbf{x})\mathbf{d}(\mathbf{x})$. Assume that $g_d(\mathbf{x})$ is a constant, so we are

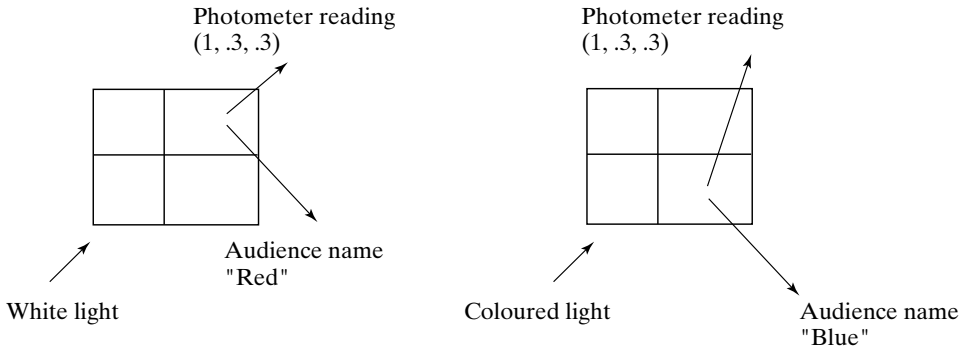

FIGURE 22.5: *Land showed an audience a quilt of rectangles of flat colored papers—since known as a Mondrian for a purported resemblance to the work of that artist—illuminated using three slide projectors, casting red, green and blue light respectively. He used a photometer to measure the energy leaving a particular spot in three different channels, corresponding to the three classes of receptor in the eye. He recorded the measurement, and asked the audience to name the patch. Assume the answer was "red" (on the* **left***). Land then adjusted the slide projectors so that some other patch reflected light that gave the same photometer measurements, and asked the audience to name that patch. The reply would describe the patch's color in white light—if the patch looked blue in white light, the answer would be "blue" (on the* **right***). In later versions of this demonstration, Land put wedge-shaped neutral density filters into the slide projectors so that the color of the light illuminating the quilt of papers would vary slowly across the quilt. Again, although the photometer readings vary significantly from one end of a patch to another, the audience sees the patch as having a constant color.*

viewing a flat, frontal surface. The resulting term, $\mathbf{d}(\mathbf{x})$, models the world as a collage of flat, frontal, diffuse colored surfaces. Such worlds are sometimes called *Mondrian worlds*, after the painter. Notice that, under our assumptions, $\mathbf{d}(\mathbf{x})$ consists of a set of patches of fixed color. We assume that there is a single illuminant that has a constant color over the whole image. This term is a conglomeration of illuminant, receptor, and reflectance information. It is impossible to disentangle these completely in a realistic world. However, current algorithms can make quite usable estimates of surface color from image colors given a well-populated world of colored surfaces and a reasonable illuminant.

Recall from Section 22.1 that if a patch of perfectly diffuse surface with diffuse spectral reflectance $\rho(\lambda)$ is illuminated by a light whose spectrum is $E(\lambda)$, the spectrum of the reflected light is $\rho(\lambda)E(\lambda)$ (multiplied by some constant to do with surface orientation, which we have already decided to ignore). If a linear photoreceptor of the kth type sees this surface patch, its response is:

$$p_k = \int_{\Lambda} \sigma_k(\lambda)\rho(\lambda)E(\lambda)d\lambda,$$

where Λ is the range of all relevant wavelengths, and $\sigma_k(\lambda)$ is the sensitivity of the

kth photoreceptor.

Finite-Dimensional Linear Models

This response is linear in the surface reflectance and linear in the illumination, which suggests using linear models for the families of possible surface reflectances and illuminants. A *finite-dimensional linear model* models surface spectral albedoes and illuminant spectral energy density as a weighted sum of a finite number of basis functions. We need not use the same bases for reflectances and for illuminants.

If a finite-dimensional linear model of surface reflectance is a reasonable description of the world, any surface reflectance can be written as

$$\rho(\lambda) = \sum_{j=1}^{n} r_j \phi_j(\lambda),$$

where the $\phi_j(\lambda)$ are the basis functions for the model of reflectance, and the r_j vary from surface to surface. Similarly, if a finite-dimensional linear model of the illuminant is a reasonable model, any illuminant can be written as

$$E(\lambda) = \sum_{i=1}^{m} e_i \psi_i(\lambda),$$

where the $\psi_i(\lambda)$ are the basis functions for the model of illumination.

When both models apply, the response of a receptor of the kth type is

$$\begin{aligned} p_k &= \int \sigma_k(\lambda) \left(\sum_{j=1}^{n} r_j \phi_j(\lambda) \right) \left(\sum_{i=1}^{m} e_i \psi_i(\lambda) \right) d\lambda \\ &= \sum_{i=1,j=1}^{m,n} e_i r_j \left(\int \sigma_k(\lambda) \phi_j(\lambda) \psi_i(\lambda) \right) d\lambda \\ &= \sum_{i=1,j=1}^{m,n} e_i r_j g_{ijk}, \end{aligned}$$

where we expect that the

$$g_{ijk} = \int \sigma_k(\lambda) \phi_j(\lambda) \psi_i(\lambda) d\lambda$$

are known, as they are components of the world model (they can be learned from observations; see the exercises).

Inferring Surface Color

The finite-dimensional linear model describes the interaction between illumination color, surface color, and image color. To infer surface color from image color, we need some sort of assumption. There are several plausible cues that can be used.

Specular reflections at dielectric surfaces have uniform specular albedo. We could find the specularities with the methods of that section, then recover surface

color using this information. At a specularity, we have

$$p_k = \int \sigma_k(\lambda) \sum_{i=1}^{m} e_i \psi_i(\lambda) d\lambda,$$

and so if we knew the spectral sensitivities of the sensor and the basis functions ψ_i, we could solve for e_i by solving a linear system. Now we know all e_i, and all p_k for each pixel. We can solve the linear system

$$p_k = \sum_{i=1,j=1}^{m,n} e_i r_j g_{ijk}$$

in the unknown r_j to recover reflectance coefficients.

Known average reflectance is another plausible cue. In this case, we assume that the spatial average of reflectance in all scenes is constant and known (e.g., we might assume that all scenes have a spatial average of reflectance that is dull gray). In the finite-dimensional basis for reflectance, we can write this average as

$$\sum_{j=1}^{n} \overline{r_j} \phi_j(\lambda).$$

Now if the average reflectance is constant, the average of the receptor responses must be constant too (if the imaging process is linear; see the discussion), and the average of the response of the kth receptor can be written as:

$$\overline{p_k} = \sum_{i=1,j=1}^{m,n} e_i g_{ijk} \overline{r_j}.$$

We know $\overline{p_k}$ and $\overline{r_j}$, and so have a linear system in the unknown light coefficients e_i. We solve this, and then recover reflectance coefficients at each pixel, as for the case of specularities. For reasonable choices of reflectors and dimension of light and surface basis, this linear system will have full rank.

The **gamut** of a color image is revealing. The gamut is the set of different colors that appears in the image. Generally, it is difficult to obtain strongly colored pixels under white light with current imaging systems. Furthermore, if the picture is taken under strongly colored light, that will tend to bias the gamut. One doesn't see bright green pixels in images taken under deep red light, for example. As a result, the image gamut is a source of information about the illumination. If an image gamut contains two pixel values—call them $\mathbf{p}_1$ and $\mathbf{p}_2$—then it must be possible to take an image *under the same illuminant* that contains the value $t\mathbf{p}_1 + (1-t)\mathbf{p}_2$ for $0 \leq t \leq 1$ (because we could mix the colorants on the surfaces). This means that the illuminant information depends on the convex hull of the image gamut. There are now various methods to exploit these observations. There is usually more than one illuminant consistent with a given image gamut, and geometric methods can be used to identify the consistent illuminants. This set can be narrowed down using probabilistic methods (for example, images contain lots of different colors [**?**]) or physical methods (for example, the main sources of illumination are the sun and the sky, well modelled as black bodies [**?**]).

22.3 NOTES

There are a number of important general resources on the use of color. We recommend **?**, **?**, **?**, **?**, **?**, **?**. **?** contains an enormous amount of helpful information. Recent textbooks with an emphasis on color include **?**, **?**, **?**, **?** and **?**.

Trichromacy and Color Spaces

Until quite recently, there was no conclusive explanation of why trichromacy applied, although it was generally believed to be due to the presence of three different types of color receptor in the eye. Work on the genetics of photoreceptors can be interpreted as confirming this hunch (see **?** and **?**), although a full explanation is still far from clear because this work can also be interpreted as suggesting many individuals have more than three types of photoreceptor [**?**].

There is an astonishing number of color spaces and color appearance models available. The important issue is not in what coordinate system one measures color, but how one counts the difference, so color metrics may still bear some thought.

Color metrics are an old topic; usually, one fits a metric tensor to MacAdam ellipses. The difficulty with this approach is that a metric tensor carries the strong implication that you can measure differences over large ranges by integration, whereas it is very hard to see large-range color comparisons as meaningful. Another concern is that the weight observers place on a difference in a Maxwellian view and the semantic significance of a difference in image colors are two very different things.

Specularity Finding

The specularity finding method we describe is due to **?**, with improvements due to **?**, [**?**], and to **?**. Specularities can also be detected because they are small and bright [**?**], because they differ in color and motion from the background [**?**, **?**, **?**], or because they distort patterns [**?**]. Specularities are a prodigious nuisance in reconstruction, because specularities cause matching points in different images to have different colors; various motion-based strategies have been developed to remove them in these applications [**?**, **?**, **?**].

Color Constancy

Land reported a variety of color vision experiments (**?**, [**?**], [**?**], [**?**]). Finite-dimensional linear models for spectral reflectances can be supported by an appeal to surface physics as spectral absorption lines are thickened by solid state effects. The main experimental justifications for finite-dimensional linear models of surface reflectance are measurements, by **?**, of the surface reflectance of a selection of standard reference surfaces known as *Munsell chips*, and measurements of a selection of natural objects by **?**. **?** performed a principal axis decomposition of his data to obtain a set of basis functions, and **?** fitted weighted sums of these functions to Krinov's date to get good fits with patterned deviations. The first three principal axes explained in each case a high percentage of the sample variance (near 99 %), and hence a linear combination of these functions fitted all the sampled functions rather well. More recently, **?** fitted Cohen's [**?**] basis vectors to a large set of data, including Krinov's [**?**] data, and further data on the surface reflectances of Munsell chips,

and concluded that the dimension of an accurate model of surface reflectance was on the order of five or six.

Finite-dimensional linear models are an important tool in color constancy. There is a large collection of algorithms that follow rather naturally from the approach. Some algorithms exploit the properties of the linear spaces involved (**?**; **?**; **?**). Illumination can be inferred from: reference objects [**?**]; specular reflections (Judd [**?**] writing in 1960 about early German work in surface color perception refers to this as "a more usual view"; recent work includes [**?**, **?**, **?**, **?**]); the average color [**?**, **?**, **?**]; and the gamut (**?**, **?**, **?**, [**?**]).

The structure of the family of maps associated with a change in illumination has been studied quite extensively. The first work is due to Von Kries (who didn't think about it quite the way we do). He assumed that color constancy was, in essence, the result of independent lightness calculations in each channel, meaning that one can rectify an image by scaling each channel independently. This practice is known as Von Kries' law. The law boils down to assuming that the family of maps consists of diagonal matrices. Von Kries' law has proved to be a remarkably good law [**?**]. Current best practice involves applying a linear transformation to the channels and then scaling the result using diagonal maps (**?**, [**?**]).

Reference datasets are available for testing methods [**?**]. Color constancy methods seem to work quite well in practice [**?**, **?**]; whether this is good enough is debated [**?**, **?**]. Probabilistic methods can be applied to color constancy [**?**]. Prior models on illumination are a significant cue [**?**].

There is surprisingly little work on color constancy that unifies a study of the spatial variation in illumination with solutions for surface color, which is why we were reduced to ignoring a number of terms in our color model. Ideally, one would work in shadows and surface orientation, too. Again, the whole thing looks like an inference problem to us, but a subtle one. The main papers on this extremely important topic are **?**, **?**. There is substantial room for research here, too.

Interreflections between colored surfaces lead to a phenomenon called *color bleeding*, where each surface reflects colored light onto the other. The phenomenon can be surprisingly large in practice. People seem to be quite good at ignoring it entirely, to the extent that most people don't realize that the phenomenon occurs at all. Discounting color bleeding probably uses spatial cues. Some skill is required to spot really compelling examples. The best known to the authors is occasionally seen in southern California, where there are many large hedges of white oleander by the roadside. White oleander has dark leaves and white flowers. Occasionally, in bright sunlight, one sees a hedge with yellow oleander flowers; a moment's thought attributes the color to the yellow service truck parked by the road reflecting yellow light onto the white flowers. One's ability to discount color bleeding effects seems to have been disrupted by the dark leaves of the plant breaking up the spatial pattern. Color bleeding contains cues to surface color that are quite difficult to disentangle (see **?**, **?**, and **?** for studies).

It is possible to formulate and attack color constancy as an inference problem [**?**, **?**]. The advantage of this approach is that, for given data, the algorithm could report a range of possible surface colors, with posterior weights.