


# 10 Light and Color

*I cannot pretend to feel impartial about colours.  
I rejoice with the brilliant ones  
and am genuinely sorry for the poor browns.*  
Winston Churchill



In ancient times it was believed that the eye radiated a cone of visual flux which mixed with visible objects in the world to create a sensation in the observer, like the sense of touch, the extromission theory. Today we consider that light from an illuminant falls on the scene, some of which is reflected into the eye of the observer to create a perception about that scene. The light that reaches the eye, or the camera, is a function of the illumination impinging on the scene and the material property known as reflectivity.

This chapter is about light itself and our perception of light in terms of brightness and color. Section 10.1 describes light in terms of electro-magnetic radiation and mixtures of light as continuous spectra. Section 10.2 provides a brief introduction to colorimetry, the science of color perception, human trichromatic color perception and how colors can be represented in various color spaces. Section 10.3 covers a number of advanced topics such as color constancy, gamma correction, and an example concerned with distinguishing different colored objects in an image.

## 10.1 Spectral Representation of Light

Around 1670 Sir Isaac Newton discovered that white light was a mixture of different colors. We now know that each of these colors is a single frequency or wavelength of electro-magnetic radiation. We perceive the wavelengths between 400 and 700 nm as colors as shown in Fig. 10.1.


In general the light that we observe is a mixture of many wavelengths and can be represented as a function  $E(\lambda)$  that describes intensity as a function of wavelength  $\lambda$ . Monochromatic light from a laser that emits light at a single wavelength in which case  $E$  is an impulse.

The most common source of light is incandescence which is the emission of light from a hot body such as the Sun or the filament of a light bulb. In physics this is modeled as a blackbody radiator or Planckian source. The emitted power is a function of wavelength  $\lambda$  and given by Planck's radiation formula


$$E(\lambda) = \frac{2hc^2}{\lambda^5 (e^{hc/k\lambda T} - 1)} \text{ W m}^{-2} \text{ m}^{-1} \quad (10.1)$$

where  $T$  is the absolute temperature (K) of the source,  $h$  is Planck's constant,  $k$  is Boltzmann's constant, and  $c$  the speed of light. ◀ This is the power emitted per unit area per unit wavelength.


$$\begin{aligned} c &= 2.998 \times 10^8 \text{ m s}^{-1}, \\ h &= 6.626 \times 10^{-34} \text{ Js}, \\ k &= 1.381 \times 10^{-23} \text{ J K}^{-1}. \end{aligned}$$



**Spectrum of light.** During the plague years of 1665–1666 Isaac Newton developed his theory of light and color. He demonstrated that a prism could decompose white light into a spectrum of colours, and that a lens and a second prism could recombine the multicoloured spectrum into white light. Importantly he showed that the color of the light did not change when it was reflected from different objects, from which he concluded that color is an intrinsic property of light not the object. (Newton's sketch to the left)



**Fig. 10.1.** The spectrum of visible colors as a function of wavelength in nanometres. The visible range depends on viewing conditions and the individual but is generally accepted as being the range 400–700 nm. Wavelengths greater than 700 nm are termed infra-red and those below 400 nm are ultra-violet



**Fig. 10.2.** Blackbody spectra. **a** Blackbody emission spectra for temperatures from 1000–6000 K. **b** Blackbody emissions for the Sun (6500 K), a tungsten lamp (2600 K) and the response of the human eye – all normalized to unity for readability

We can plot the emission spectra for a blackbody at different temperatures. First we define a range of wavelengths

```
>> lambda = [300:10:1000]*1e-9;
```

in this case from 300 to 1000 nm, and then compute the blackbody spectra

```
>> for T=1000:1000:6000
>> plot( lambda*1e9, blackbody(lambda, T)); hold all
>> end
```

as shown in Fig. 10.2a. We can see that as temperature increases the maximum amount of power increases and the wavelength at which the peak occurs decreases. The total amount of power radiated (per unit area) is the area under the blackbody curve and is given by the Stefan-Boltzman law

$$\frac{2\pi^5 k^4}{15c^2 h^3} T^4 \text{ W m}^{-2}$$

and the wavelength corresponding to the peak of the blackbody curve is given by Wien’s displacement law

$$\lambda_{\max} = \frac{2.8978 \times 10^{-3}}{T} \text{ m}$$

The wavelength of the peak decreases with increasing temperature and in familiar terms this is what we observe when we heat an object. It starts to glow faintly red at around 800 K and moves through orange and yellow toward white as temperature increases. ▶

The filament of tungsten lamp has a temperature of 2600 K and glows *white hot*. The Sun has a surface temperature of 6500 K. The spectra of these sources

```
>> lamp = blackbody(lambda, 2600);
>> sun = blackbody(lambda, 6500);
>> plot(lambda*1e9, [lamp/max(lamp) sun/max(sun)])
```

Incipient red heat	770– 820 K,
dark red heat	920–1020 K,
bright red heat	1120–1220 K,
yellowish red heat	1320–1420 K,
incipient white heat	1520–1620 K,
white heat	1720–1820 K.



Sir Humphry Davy demonstrated the first electrical incandescent lamp using a platinum filament in 1802. Sir Joseph Swan demonstrated his first light bulbs in 1850 using carbonized paper filaments. However it was not until advances in vacuum pumps in 1865 that such lamps could achieve a useful lifetime. Swan patented a carbonized cotton filament in 1878 and a carbonized cellulose filament in 1881. His lamps came into use after 1880 and the Savoy Theatre in London was completely lit by electricity in 1881. In the USA Thomas Edison did not start research into incandescent lamps until 1878 but he patented a long-lasting carbonized bamboo filament the next year and was able to mass produce them. The Swan and Edison companies merged in 1883.

The light bulb subsequently became the dominant source of light on the planet but is now being phased out due to its poor energy efficiency. (Photo by Douglas Brackett, Inv., Edisonian.com)

are compared in Fig. 10.2b. The tungsten lamp curve is much lower in magnitude, but has been scaled up for readability. The peak of the Sun's emission is around 450 nm and it emits a significant amount of power in the visible part of the spectrum. The peak for the tungsten lamp is at a much longer wavelength and perversely very little of its power falls within the human visible region. The bulk of the power is infra-red which we perceive as heat not light.

### 10.1.1 Absorption

The Sun's spectrum at ground level on the Earth has been measured and tabulated

```
>> sun_ground = loadspectrum(lambda, 'solar.dat');
>> plot(lambda*1e9, sun_ground)
```

and is shown in Fig. 10.3a. It differs markedly from that of a blackbody since some wavelengths have been absorbed more than others by the atmosphere. Our eye's peak sensitivity has evolved to be closely aligned to the peak of the spectrum of atmospherically filtered sunlight.



Transmittance  $T$  is the inverse of absorptance and is the fraction of light passed as a function of wavelength. It is described by Beer's law

$$T = 10^{-Ad} \quad (10.2)$$

where  $A$  is the absorption coefficient in units of  $\text{m}^{-1}$  and  $d$  is the path length. The absorption spectrum  $A(\lambda)$  for water is loaded from tabulated data

```
>> [A, lambda] = loadspectrum([400:10:700]*1e-9, 'water.dat');
```

**Fig. 10.3.** **a** Modified solar spectrum at ground level (blue). The dips in the solar spectrum correspond to various water absorption bands.  $\text{CO}_2$  absorbs radiation in the infra-red region, and ozone  $\text{O}_3$  absorbs strongly in the ultra-violet region. The Sun's blackbody spectrum (normalized) is shown in red and the response of the human eye is shown dashed. **b** Transmission through 5 m of water. The longer wavelengths, reds, have been strongly attenuated



and the transmission through 5 m of water is

```
>> d = 5;
>> T = 10.^(-A*d);
>> plot(lambda*1e9, T);
```

which is plotted in Fig. 10.3b. Differential absorption of wavelengths is a significant concern when imaging underwater and we revisit this topic in Sect. 10.3.1.

### 10.1.2 Reflection

The light reflected from a surface, its luminance, has a spectrum given by

$$L(\lambda) = E(\lambda)R(\lambda) \text{ W m}^{-2} \quad (10.3)$$

where  $E$  is the incident illumination and  $R \in [0, 1]$  is the reflectivity or reflectance of the surface and is a function of wavelength. White paper for example has a reflectance of around 70%. The reflectance spectra of many materials have been measured and tabulated. ▶ Consider for example the reflectivity of a red house brick

```
>> [R, lambda] = loadspectrum([100:10:10000]*1e-9, 'redbrick.dat');
>> plot(lambda*1e6, R);
```

which is plotted in Fig. 10.4. We see that it reflects red colors more than blue.

The illuminance of the Sun in the visible region

```
>> lambda = [400:10:700]*1e-9;           % visible spectrum
```

is

```
>> E = loadspectrum(lambda, 'solar.dat');
```

at ground level. The reflectivity of the brick is


```
>> R = loadspectrum(lambda, 'redbrick.dat');
```

and the light reflected from the brick is


```
>> L = E .* R;
>> plot(lambda*1e9, L);
```

which is shown in Fig. 10.5. It is this spectrum that is interpreted by our eyes as the color red.

From [http://speclib.jpl.nasa.gov/weathered red brick \(0412UUUBRK\)](http://speclib.jpl.nasa.gov/weathered_red_brick_(0412UUUBRK)).



**Fig. 10.4.** Reflectance of a weathered red house brick (data from ASTER, Baldrige et al. 2009). **a** Full range measured from 300 nm visible to 10 000 nm (infrared); **b** closeup of visible region



**Fig. 10.5.**  
Luminance of the weathered red house brick under illumination from the Sun at ground level, based on data from Fig. 10.3a and 10.4b

## 10.2 Color

*Color is the general name for all sensations arising from the activity of the retina of the eye and its attached nervous mechanisms, this activity being, in nearly every case in the normal individual, a specific response to radiant energy of certain wavelengths and intensities.*

T. L. Troland,  
Report of Optical Society of America  
Committee on Colorimetry 1920–1921

We have described the spectra of light in terms of power as a function of wavelength, but our own perception of light is in terms of subjective quantities such as brightness and color. Light that is visible to humans lies in the range of wavelengths from 400 nm (violet) to 700 nm (red) with the colors blue, green, yellow and orange in between, as shown in Fig. 10.1.

Our eyes contain two types of light sensitive cells as shown in Fig. 10.6. Cone cells respond to particular colors and provide us with our normal daytime vision. Rod cells are much more sensitive than cone cells but respond to intensity only and are used at night. ◀

Therefore at night you have no color vision.

The brightness we associate with a particular wavelength is known as luminosity and is measured in units of lumens per watt. For our daylight cone-cell vision the luminosity as a function of wavelength has been experimentally determined, tabulated and forms the basis of the 1931 CIE standard that represents the average human observer. ◀ The luminosity function is provided by the Toolbox

This is the photopic response for a light-adapted eye using the cone photoreceptor cells. The dark adapted, or scotopic response, using the eye's monochromatic rod photoreceptor cells is different, and peaks at around 510 nm.

```
>> human = luminos(lambda);
>> plot(lambda*1e9, human)
```

and is shown in Fig. 10.7a. Consider two lights emitting the same power (in watts) but one has a wavelength of 550 nm (green) and the other has a wavelength of 450 nm (blue). The perceived brightness of these two lights is quite different, in fact the blue light appears only

```
>> luminos(450e-9) / luminos(550e-9)
ans =
    0.0382
```

or 3.8% as bright as the green one.

**Radiometric and photometric quantities.** Two quite different sets of units are used when discussing light: radiometric and photometric. Radiometric units are used in Sect. 10.1 and are based on quantities such as power and are expressed in familiar SI units such as watts.

Photometric units are analogs of radiometric units but take into account the *visual sensation* in the observer. Luminous power or luminous flux is the *perceived* power of a light source and is measured in *lumens* (abbreviated to lm) rather than *watts*. A 1 W light source at 555 nm, the peak response, *by definition* emits a luminous flux of 683 lm. By contrast a 1 W light source at 800 nm emits a luminous flux of 0 lm – it causes no visual sensation at all.


A 1 W incandescent lightbulb however produces a perceived visual sensation of less than 15 lm or a luminous efficiency of 15 lm W<sup>-1</sup>. Fluorescent lamps achieve efficiencies up to 100 lm W<sup>-1</sup> and white LEDs up to 150 lm W<sup>-1</sup>.

The eyes of different species have different spectral responses. Many insects are able to see well into the ultra-violet region of the spectrum. The silicon sensors used in digital cameras have strong sensitivity in the red and infra-red part of the spectrum which we can also plot



```
>> camera = ccdresponse(lambda);
>> hold on
>> plot(lambda*1e9, camera*max(human), '--')
```

and is shown superimposed in Fig. 10.7.

The LED on an infra-red remote control can be seen as a bright light in most digital cameras – try this with your mobile phone camera and TV remote. Some security cameras provide infra-red scene illumination for covert night time monitoring. Note that some cameras are fitted with infra-red filters to prevent the sensor becoming saturated by ambient infra-red radiation.



**Fig. 10.6.** A coloured scanning electron micrograph of rod cells (white) and cone cells (yellow) in the human eye. The cells diameters are in the range 0.5–4 μm. The cells contain different types of light-sensitive protein called photopsin. The cell bodies (red) of the receptor cells are located in a layer above the rods and cones



**Fig. 10.7.** a Luminosity curve for standard observer human observer. The peak response is 683 lm W<sup>-1</sup> at 555 nm (green). The response of a silicon CCD camera is shown dashed for comparison. b Spectral response of human cones (normalized)

Solid angle is measured in steradians, a full sphere is  $4\pi$  sr.

**Lightmeters, illuminance and luminance.** A photographic lightmeter measures luminous flux which has units of  $\text{lm m}^{-2}$  or lux (lx). The luminous intensity of a point light source is the luminous flux per unit solid angle measured in  $\text{lm sr}^{-1}$  or candelas (cd). For a point source of luminous intensity  $I$  the illuminance  $E$  falling normally onto a surface is

$$E = \frac{I}{d^2} \text{ lx}$$

where  $d$  is the distance between source and the surface. Outdoor illuminance on a bright sunny day is approximately 10 000 lx whereas office lighting levels are typically around 1 000 lx.

The luminance or *brightness* of a surface is

$$L_s = E_i \cos \theta \text{ nt}$$

which has units of  $\text{cd m}^{-2}$  or nit (nt), and where  $E_i$  is the incident illuminance at an angle  $\theta$  to the surface normal.

In normal daylight conditions our cone photoreceptors are active and these are color sensitive. Humans are trichromats and have three types of cones that respond to different parts of the spectrum. They are referred to as long (L), medium (M) and short (S) according to the wavelength of their peak response, or more commonly as red, green and blue. The spectral responses of the cones can be loaded

```
>> cones = loadspectrum(lambda, 'cones.dat');
>> plot(lambda*1e9, cones)
```

where `cones` has three columns for each of the L, M and S cone responses and each row corresponds to the wavelength in `lambda`. The spectral response of the cones  $L(\lambda)$ ,  $M(\lambda)$  and  $S(\lambda)$  are shown in Fig. 10.7b. ◀

Other species have different numbers of cones. Birds, fish and amphibians are tetrachromats, that is, they have four types of cones. Most other mammals, for instance dogs, are dichromats and have only two types of cones. There is speculation that some human females are tetrachromats. ◀

The retina of the human eye has a central or foveal region which is only 0.6 mm in diameter and contains most of the 6 million cone cells: 65% sense red, 33% sense green and only 2% sense blue. We unconsciously scan our high-resolution fovea over the world to build a large-scale mental image of our surrounds. In addition there are 120 million rod cells, which are also motion sensitive, distributed over the retina.

The sensor in a digital camera is analogous to the retina, but instead of rod and cone cells there is a regular array of light sensitive photosites on a silicon chip. Each photosite is of the order 1–10  $\mu\text{m}$  square and outputs a signal proportional to the intensity of the light falling over its area. ◀ For a color camera the photosites are covered by color filters which pass either red, green or blue light to the photosites. The spectral response of the filters is the functional equivalent of the cones' response  $M(\lambda)$  in Fig. 10.7b. A very common arrangement of color filters is the Bayer pattern shown in Fig. 10.8. It uses a regular  $2 \times 2$  photosite pattern comprising two green filters, one red and one blue. ◀

The luminance of an object  $L(\lambda)$  given by Eq. 10.3 results in a particular response from each of the three cones

$$\begin{aligned} \rho &= \int_{\lambda} L(\lambda) M_r(\lambda) d\lambda \\ \gamma &= \int_{\lambda} L(\lambda) M_g(\lambda) d\lambda \\ \beta &= \int_{\lambda} L(\lambda) M_b(\lambda) d\lambda \end{aligned} \tag{10.4}$$

where  $M_r(\lambda)$ ,  $M_g(\lambda)$  and  $M_b(\lambda)$  are the spectral response of the red, green and blue cones respectively as shown in Fig. 10.7b. The response is a 3-vector  $(\rho, \gamma, \beta)$  which is known as a tristimulus.


The spectral characteristics are due to the different photopsins in the cone cell. George Wald won the 1967 Nobel Prize in Medicine for his experiments in the 1950s that showed the absorbance of retinal photopsins.

They have an extra variant version of the long-wave (L) cone type which would lead to greater ability in color discrimination.

More correctly the output is proportional to the total number of photons captured by the photosite since the last time it was read. See page 260.

Each pixel therefore cannot provide independent measurements of red, green and blue but it can be estimated. For example, the amount of red at a blue sensitive pixel is obtained by interpolation from its red filtered neighbours. More expensive "3 CCD" cameras can make independent measurements at each pixel since the light is split by a set of prisms, filtered and presented to one CCD array for each primary color. Digital camera raw image files contain the actual outputs of the Bayer-filtered photosites.

**Color blindness**, or color deficiency, is the inability to perceive differences between some of the colors that others can distinguish. Protanopia, deuteranopia, tritanopia refer to the absence of the L, M and S cones respectively. More common conditions are protanomaly, deuteranomaly and tritanomaly where the cone pigments are mutated and the peak response frequency changed. It is most commonly a genetic condition since the red and green photopsins are coded in the X chromosome. The most common form (occurring in 6% of males including the author) is deuteranomaly where the M-cone's response is shifted toward the red end of the spectrum resulting in reduced sensitivity to greens and poor discrimination of hues in the red, orange, yellow and green region of the spectrum.



	red	green	blue
$\lambda$ (nm)	700.0	546.1	435.8

For our red brick example the tristimulus can be computed by approximating the integrals of Eq. 10.4 as a summation

```
>> sum( (L*ones(1,3)) .* cones )
ans =
    16.3578    10.0702     2.8219
```

The dominant response is from the L cone, which is unsurprising since we know that the brick is red.

An arbitrary continuous spectrum is an infinite-dimensional vector and cannot be uniquely represented by just 3 parameters. A consequence of this is that many *different* spectral power distributions will produce the *same* visual stimulus and these are referred to as metamers. More important is the corollary – an arbitrary visual stimulus can be generated by a mixture of just three monochromatic stimuli. These are the three primary colors we learnt about as children. There is no unique set of primaries – any three will do so long as none of them can be matched by a combination of the others. The wavelength of the CIE 1976 standard primaries are given in Table 10.1.

**Fig. 10.8.** Bayer filtering. The grey blocks represents the array of light-sensitive silicon photosites over which is an array of red, green and blue filters. Invented by Bryce E. Bayer of Eastman Kodak, U.S. Patent 3,971,065.

**Table 10.1.** The CIE 1976 primaries (Commission Internationale de L'Éclairage 1987) are spectral colors corresponding to the emission lines in a mercury vapor lamp

Primary colors are not a fundamental property of light – they are a fundamental property of the observer. There are three primary colors only because we, as trichromats, have three types of cones. Birds would have four primary colors and dogs would have two.

### 10.2.1 Reproducing Colors

A computer or television display is able to produce a variable amount of each of three primaries at every pixel. The primaries for a cathode ray tube (CRT) are created by exciting phosphors on the back of the screen. For a liquid crystal display (LCD) the colors are obtained by filtering white light emitted by the backlight. The important problem is to determine how much of each primary is required to match a given tristimulus.

We start by considering a monochromatic stimulus of wavelength  $\lambda_s$  which is defined as

$$L(\lambda) = \begin{cases} L_{\lambda} & \text{if } \lambda = \lambda_s \\ 0 & \text{otherwise} \end{cases}$$

The response of the cones to this stimulus is given by Eq. 10.4 but because  $L(\cdot)$  is an impulse we can drop the integral to obtain the tristimulus

$$\begin{aligned}\rho &= L_\lambda M_r(\lambda_s) \\ \gamma &= L_\lambda M_g(\lambda_s) \\ \beta &= L_\lambda M_b(\lambda_s)\end{aligned}\quad (10.5)$$

The units are chosen such that equal quantities of the primaries are required to match the equal-energy white stimulus.

Consider next three primary light sources denoted **R**, **G** and **B** with wavelengths  $\lambda_r$ ,  $\lambda_g$  and  $\lambda_b$  and intensities  $R$ ,  $G$  and  $B$  respectively. The tristimulus from these light sources is

$$\begin{aligned}\rho &= RM_r(\lambda_r) + GM_r(\lambda_g) + BM_r(\lambda_b) \\ \gamma &= RM_g(\lambda_r) + GM_g(\lambda_g) + BM_g(\lambda_b) \\ \beta &= RM_b(\lambda_r) + GM_b(\lambda_g) + BM_b(\lambda_b)\end{aligned}\quad (10.6)$$

For the perceived color of these three light sources to match that of the monochromatic stimulus the two tristimuli must be equal. We equate Eq. 10.5 and Eq. 10.6 and write compactly in matrix form as

$$L_\lambda \begin{pmatrix} M_r(\lambda_s) \\ M_g(\lambda_s) \\ M_b(\lambda_s) \end{pmatrix} = \begin{pmatrix} M_r(\lambda_r) & M_r(\lambda_g) & M_r(\lambda_b) \\ M_g(\lambda_r) & M_g(\lambda_g) & M_g(\lambda_b) \\ M_b(\lambda_r) & M_b(\lambda_g) & M_b(\lambda_b) \end{pmatrix} \begin{pmatrix} R \\ G \\ B \end{pmatrix}$$

which we can invert to give the required amounts of primary colors

$$\begin{pmatrix} R \\ G \\ B \end{pmatrix} = L_\lambda \begin{pmatrix} M_r(\lambda_r) & M_r(\lambda_g) & M_r(\lambda_b) \\ M_g(\lambda_r) & M_g(\lambda_g) & M_g(\lambda_b) \\ M_b(\lambda_r) & M_b(\lambda_g) & M_b(\lambda_b) \end{pmatrix}^{-1} \begin{pmatrix} M_r(\lambda_s) \\ M_g(\lambda_s) \\ M_b(\lambda_s) \end{pmatrix}$$

The required tristimulus values are simply a linear transformation of the cone's response to the monochromatic excitation. The transformation matrix is constant, but depends upon the spectral response of the cones to the chosen primaries ( $\lambda_r$ ,  $\lambda_g$ ,  $\lambda_b$ ). This is the basis of trichromatic matching.

**Color matching experiments** are performed using a light source comprising three adjustable lamps that correspond to the primary colors and whose intensity can be individually adjusted. The lights are mixed and diffused and compared to some test color. In color matching notation the primaries, the lamps, are denoted by **R**, **G** and **B**, and their intensities are  $R$ ,  $G$  and  $B$  respectively. The three lamp intensities are adjusted by a human subject until they appear to match the test color. This is denoted

$$\mathbf{C} \equiv R\mathbf{R} + G\mathbf{G} + B\mathbf{B}$$


which is read as the visual stimulus **C** (the test color) is matched by, or looks the same as, a mixture of the three primaries with brightness  $R$ ,  $G$  and  $B$ . The notation  $R\mathbf{R}$  can be considered as the lamp **R** at intensity  $R$ .

Experiments show that color matching obeys the algebraic rules of additivity and linearity which is known as Grassmann's laws. For example two light stimuli  $C_1$  and  $C_2$

$$\begin{aligned}C_1 &\equiv R_1\mathbf{R} + G_1\mathbf{G} + B_1\mathbf{B} \\ C_2 &\equiv R_2\mathbf{R} + G_2\mathbf{G} + B_2\mathbf{B}\end{aligned}$$

when mixed will match

$$C_1 + C_2 \equiv (R_1 + R_2)\mathbf{R} + (G_1 + G_2)\mathbf{G} + (B_1 + B_2)\mathbf{B}$$



**Fig. 10.9.** The 1931 color matching functions for the standard CIE observer, based on the CIE standard primaries

We can write this in an even more compact form

$$\begin{pmatrix} R \\ G \\ B \end{pmatrix} = \begin{pmatrix} \bar{r}(\lambda_s) \\ \bar{g}(\lambda_s) \\ \bar{b}(\lambda_s) \end{pmatrix} \quad (10.7)$$

where  $\bar{r}(\lambda)$ ,  $\bar{g}(\lambda)$ ,  $\bar{b}(\lambda)$  are known as color matching functions. These functions have been tabulated for the standard CIE primaries listed in Table 10.1 and are returned by the function `cmfrgb`

```
>> lambda = [400:10:700]*1e-9;
>> cmf = cmfrgb(lambda);
>> plot(lambda*1e9, cmf);
```

and shown graphically in Fig. 10.9. Each curve shows how much of the corresponding primary is required to match the monochromatic light of wavelength  $\lambda$ .

For example to create the sensation of light at 600 nm (orange) we would need

```
>> orange = cmfrgb(600e-9)
orange =
    2.8717    0.3007   -0.0043
```

Surprisingly this requires a very small *negative* amount of the blue primary. To create 500 nm green we would need

```
>> green = cmfrgb(500e-9)
green =
   -0.2950    0.4906    0.1075
```

and this requires a significant *negative* amount of the red primary. This is problematic since a light source cannot have a negative luminance.

We reconcile this by adding some white light ( $R = G = B = w$ , see Sect. 10.2.6) so that the tristimulus values are all positive. For instance

```
>> w = -green(1);
>> white = [w w w];
>> feasible_green = green + white
feasible_green =
    0    0.7856    0.4025
```

If we looked at this color side-by-side with the desired 500 nm green we would say that the generated color had the correct hue but was not as *saturated*.

Saturation refers to the purity of the color. Spectral colors are *fully saturated* but become less saturated (more pastel) as increasing amounts of white is added. In this case we have mixed a stimulus of

```
>> white
white =
    0.2950    0.2950    0.2950
```

which is a light grey.

This leads to a very important point about color reproduction – it is *not* possible to reproduce every possible color using just three primaries. This makes intuitive sense since a color is properly represented as an infinite-dimensional spectral function and a 3-vector can only approximate it. To understand this more fully we need to consider chromaticity spaces.

The Toolbox function `cmfrgb` can also compute the CIE tristimulus for an arbitrary spectrum. The luminance spectrum of the redbrick illuminated by sunlight at ground level was computed earlier and its tristimulus is

```
>> RGB_brick = cmfrgb(lambda, L)
RGB_brick =
    0.6137    0.1416    0.0374
```

These are the respective amounts of the three CIE primaries that are perceived as having the same color as the brick.

### 10.2.2 Chromaticity Space

The tristimulus values describe color as well as brightness. Relative tristimulus values are obtained by normalizing the tristimulus values

$$r = \frac{R}{R+G+B}, g = \frac{G}{R+G+B}, b = \frac{B}{R+G+B} \quad (10.8)$$

which results in chromaticity coordinates  $r$ ,  $g$  and  $b$  that are invariant to overall brightness. By definition  $r + g + b = 1$  so one coordinate is redundant and typically only  $r$  and  $g$  are considered. Since the effect of intensity has been eliminated the 2-dimensional quantity  $(r, g)$  represents *color*.

We can plot the locus of spectral colors, the colors of the rainbow, on the chromaticity diagram using a variant of the color-matching functions


```
>> [r,g] = lambda2rg( [400:700]*1e-9 );
>> plot(r, g)
>> rg_addticks
```

which results in the horseshoe-shaped curve shown in Fig. 10.10. The Toolbox function `lambda2rg` computes the color matching function Eq. 10.7 for the specified wavelength and then converts the tristimulus value to chromaticity coordinates using Eq. 10.8.

The CIE primaries listed in Table 10.1 can be plotted as well

```
>> primaries = cmfrgb( [700, 546.1, 435.8]*1e-9 );
>> plot(primaries(:,1), primaries(:,2), 'd')
```

and are shown as diamonds in Fig. 10.10.




**Fig. 10.10.**

The spectral locus on the  $r$ - $g$  chromaticity plane. The CIE standard primary colors are marked by diamonds. Spectral wavelengths (in nm) are marked. The straight line joining the extremities is the purple boundary and is the locus of saturated purples

**Colorimetric standards.** Colorimetry is a complex topic and standards are very important. Two organizations, CIE and ITU, play a leading role in this area.

The Commission Internationale de l'Éclairage (CIE) or International Commission on Illumination was founded in 1913 and is an independent non-profit organisation that is devoted to worldwide cooperation and the exchange of information on all matters relating to the science and art of light and lighting, colour and vision, and image technology. The CIE's eighth session was held at Cambridge, UK, in 1931 and established international agreement on colorimetric specifications and formalized the XYZ color space. The CIE is recognized by ISO as an international standardization body. See <http://www.cie.co.at> for more information and CIE datasets.

The International Telecommunication Union (ITU) is an agency of the United Nations and was established to standardize and regulate international radio and telecommunications. It was founded as the International Telegraph Union in Paris on 17 May 1865. The International Radio Consultative Committee or CCIR (Comité Consultatif International des Radiocommunications) became, in 1992, the Radiocommunication Bureau of ITU or ITU-R. It publishes standards and recommendations relevant to colorimetry in its broadcasting service (television) or BT series. See <http://www.itu.int> for more detail.



**Fig. 10.11.** Chromaticity diagram showing 500 nm green (square), equal-energy white (asterisk), a feasible green (star) and a displayable green (circle). The locus of different saturated greens is shown as a green line


The centre of gravity law states that a mixture of two colors lies along a line between those two colors on the chromaticity plane. A mixture of  $N$  colors lies within a region bounded by those colors. Considered with respect to Fig. 10.10 this has significant implications. Firstly, since all color stimuli are combinations of spectral stimuli all real color stimuli must lie on or inside the spectral locus. Secondly, any colors we create from mixing the primaries can only lie *within* the triangle bounded by the primaries – the color gamut. It is clear from Fig. 10.10 that the CIE primaries define only a small subset of all possible colors – shown as a dashed triangle. Very many real colors *cannot* be created using these primaries, in particular the colors of the rainbow which lie on the spectral locus from 460–545 nm. In fact no matter where the primaries are located, not all possible colors can be produced. In geometric terms there are no three points within the gamut that form a triangle that includes the entire gamut. Thirdly, we observe that much of the locus requires a negative amount of the red primary and cannot be represented.

The problem on page 232 with displaying 500 nm green is explained by it lying outside the gamut of the CIE primaries and this is shown in Fig. 10.11. We plot the chromaticity of the spectral green color

```
>> green_cc = lambda2rg(500e-9);
green_cc =
    -0.9733    1.6187
>> plot2(green_cc, 's')
```

as a square marker. White is by definition  $R = G = B = 1$  and its chromaticity

We could increase the gamut by choosing different primaries, perhaps using a different green primary would make the gamut larger, but there is the practical constraint of finding a light source (LED or phosphor) that can efficiently produce that color.



**Fig. 10.12.** a The color matching functions for the standard observer, based on the imaginary primaries X, Y (intensity) and Z are tabulated by the CIE (Commission Internationale de l'Éclairage 1987). b Colors on the  $xy$ -chromaticity plane

The units are chosen such that equal quantities of the primaries are required to match the equal-energy white stimulus.

```
>> white_cc = tristim2cc([1 1 1])
white_cc =
    0.3333    0.3333
>> plot2(white_cc, '*')
```

is plotted as an asterisk. According to the centre of gravity law the mixture of our desired green and white must lie along the indicated line. The chromaticity of the least saturated displayable green lies at the intersection of this line and the gamut boundary and is indicated by a circle.

Earlier we said that there are no three points within the gamut that form a triangle that includes the entire gamut. The CIE therefore proposed, in 1931, a system of *imaginary non-physical primaries* known as X, Y and Z that totally enclose the spectral locus of Fig. 10.10. X and Z have zero luminance – the luminance is contributed entirely by Y. All real colors can thus be matched by positive amounts of these three primaries. ◀ The corresponding tristimulus values are denoted (X, Y, Z).

The XYZ color matching functions defined by the CIE

```
>> cmf = cmfxyz(lambda);
>> plot(lambda*1e-9, cmf);
```

are shown graphically in Fig. 10.12a. This shows the amount of each CIE XYZ primary required to match a spectral color and we note that these curves are never negative. The corresponding chromaticity coordinates are

$$x = \frac{X}{X+Y+Z}, y = \frac{Y}{X+Y+Z}, z = \frac{Z}{X+Y+Z} \quad (10.9)$$

and once again  $x + y + z = 1$  so only two parameters are required – by convention  $y$  is plotted against  $x$  in a chromaticity diagram. The spectral locus can be plotted in a similar way as before

```
>> [x,y] = lambda2xy(lambda);
>> plot(x, y);
```

A more sophisticated plot, showing the colors within the spectral locus, can be created

```
>> xycolorspace
```

and is shown ◀ in Fig. 10.12b. These coordinates are a *standard* way to represent color for graphics, printing and other purposes. For example the chromaticity coordinates of peak green (550 nm) is

```
>> lambda2xy(550e-9)
ans =
    0.3016    0.6924
```

Note that the colors depicted are only approximation of the actual color at that point due to the gamut limitation of the printed colors. No display device has a gamut large enough to present an accurate representation of the chromaticity at every point.

and the chromaticity coordinates of a standard tungsten illuminant at 2600 K is

```
>> lamp = blackbody(lambda, 2600);
>> lambda2xy(lambda, lamp)
ans =
    0.4679    0.4126
```

### 10.2.3 Color Names

Chromaticity coordinates provide a quantitative way to describe and compare colors, however humans refer to colors by name. Many computer operating systems contain a database or file that maps human understood names of colors to their corresponding ( $R, G, B$ ) tristimulus values. A typical database contains more than 800 uniquely named colors which says something about the importance of color to humans. The Toolbox provides a copy of a such a file and an interface function `colorname`. For example we can query a color name that includes a particular substring

```
>> colorname('?burnt')
ans =
    'burntsienna'    'burntumber'
```

The RGB tristimulus value of `burntsienna` is

```
>> colorname('burntsienna')
ans =
    0.5412    0.2118    0.0588
```

with the values normalized to the interval  $[0, 1]$ . We could also request  $xy$ -chromaticity coordinates

```
>> bs = colorname('burntsienna', 'xy')
bs =
    0.5258    0.3840
```

With reference to Fig. 10.12 we see that this point is in the red-brown part of the colorspace and not too far from the color of chocolate

```
>> colorname('chocolate', 'xy')
ans =
    0.5092    0.4026
```

We can also solve the inverse problem. For example consider a tristimulus value close to, but not exactly the same as, burnt Sienna

```
>> colorname([0.54 0.20 0.06])
ans =
    burntsienna
```

and the name of the closest color, in Euclidean terms, is returned. We can repeat this with the color specified in  $xy$ -chromaticity coordinates

```
>> colorname(bs, 'xy')
ans =
    'burntsienna'
```

### 10.2.4 Other Color Spaces

A color space is a 3-dimensional space that contains all possible tristimulus values – all colors and all levels of brightness. If we think of this in terms of coordinate frames as discussed in Sect. 2.2 then there are an infinite number of choices of Cartesian frame with which to define colors. We have already discussed two different Cartesian color spaces: RGB and XYZ. However we could also use polar, spherical or hybrid coordinate systems.

The chromaticity spaces  $r-g$  or  $x-y$  do not account for brightness – we normalized it out in Eq. 10.8 and Eq. 10.9. Brightness is more precisely called luminance and is typically denoted by  $Y$ . The definition from ITU Recommendation 709

The file is named `/etc/rgb.txt` on most Unix-based systems.

$$Y^{709} = 0.2126R + 0.7152G + 0.0722B \quad (10.10)$$

is a weighted sum of the RGB-tristimulus values and reflects the eye's high sensitivity to green and low sensitivity to blue. Chromaticity plus luminance leads to 3-dimensional color spaces such as  $r$ - $g$ - $Y$  or  $x$ - $y$ - $Y$ .

Humans seem to more naturally consider chromaticity in terms of two characteristics: hue and saturation. Hue is the dominant color, the closest spectral color, and saturation refers to the purity, or absence of mixed white. Stimuli on the spectral locus are completely saturated while those closer to its centroid are less saturated. The color spaces that we have discussed lack easy interpretation in terms of hue and saturation so alternative color spaces have been proposed. The two most commonly known are HSV and CIE  $L^*C^*h$ . In color-space notation H is hue, S is saturation which is also known as C or chroma. The intensity dimension is named either V for value or L for lightness but they are computed quite differently. The concepts of hue and saturation is illustrated in geometric terms in Fig. 10.13.

The function `colorspace` can be used to perform conversions between different color spaces. For example the hue, saturation and intensity for each of pure red, green and blue RGB tristimulus value is

```
>> colorspace('RGB->HSV', [1, 0, 0])
ans =
    0    1    1
>> colorspace('RGB->HSV', [0, 1, 0])
ans =
   120    1    1
>> colorspace('RGB->HSV', [0, 0, 1])
ans =
   240    1    1
```

In each case the saturation is 1, the colors are pure, and the intensity is 1. As shown in Fig. 10.13 hue is represented as an angle in the range  $[0, 360)^\circ$  with red at  $0^\circ$  increasing through the spectral colors associated with decreasing wavelength (orange, yellow, green, blue, violet). If we reduce the amount of the green primary

```
>> colorspace('RGB->HSV', [0, 0.5, 0])
ans =
  120.0000    1.0000    0.5000
```

we see that intensity drops but hue and saturation are unchanged. For a medium grey

```
>> colorspace('RGB->HSV', [0.4, 0.4, 0.4])
ans =
  240.0000    0    0.4000
```

the saturation is zero, it is only a mixture of white, and the hue has no meaning since there is no color. If we add the green to the grey




Fig. 10.13.

Hue and saturation. A line is extended from the white point through the chromaticity in question to the spectral locus. The angle of this line is hue, and saturation is the length of the vector normalized with respect to distance to the locus

$L^*$  is a non-linear function of relative luminance and approximates the non-linear response of the human eye. Value is given by  $V = \frac{1}{2}(\min R, G, B + \max R, G, B)$ .

For very dark colors numerical problems lead to imprecise hue and saturation coordinates.



```
>> colorspace('RGB->HSV', [0, 0.5, 0] + [0.4, 0.4, 0.4])
ans =
    120.0000    0.5556    0.9000
```

Fig. 10.14. Flower scene. **a** Original color image; **b** hue image; **c** saturation image. Note that the white flowers have low saturation (they appear dark)

we have the green hue and a medium saturation value.

The `colorspace` function can convert between thirteen different color spaces including YUV, YCbCr,  $L^*a^*b^*$  and  $L^*u^*v^*$ . A limitation of many color spaces is that the perceived color difference between two closely spaced points depends on the position of those points in the space. This has led to the development of perceptually uniform color spaces such as the CIE  $L^*u^*v^*$  (CIELUV) and  $L^*a^*b^*$  spaces.

The `colorspace` function can also be applied to a color image as shown in Fig. 10.14. In the hue image dark represents red and bright white represents violet. The red flowers appear as both a very small hue angle (dark) and a very large angle close to 360°. The yellow flowers and the green background can be seen as distinct hue values. The saturation image shows that the red and yellow flowers are highly saturated, while the green leaves and stems are less saturated. The white flowers have very low saturation, since by definition the color white contains a lot of white. This example is explained in more detail, and extended, in Sect. 10.3.5.

### 10.2.5 Transforming between Different Primaries

The CIE standards were defined in 1931 which was well before the introduction of color television in the 1950s. The CIE primaries are based on the emission lines of a mercury lamp which are highly repeatable and suitable for laboratory use. Early television receivers used CRT monitors where the primary colors were generated by phosphors that emit light when bombarded by electrons. The phosphors used, and their colors has varied over the years in pursuit of brighter displays. An international agreement, ITU recommendation 709, defines the primaries for high definition television (HDTV) and these are listed in Table 10.2.

This raises the problem of converting tristimulus values from one sets of primaries to another. Consider for example that we wish to display an image, where the tristimulus values are with respect to CIE primaries, on a screen that uses ITU Rec. 709 primaries. Using the notation we introduced earlier we define two sets of primaries:  $P_1, P_2, P_3$  with tristimulus values  $(S_1, S_2, S_3)$ , and  $P'_1, P'_2, P'_3$  with tristimulus values  $(S'_1, S'_2, S'_3)$ . We can always express one set of primaries as a linear combination of the other

$$\begin{pmatrix} P_1 \\ P_2 \\ P_3 \end{pmatrix} = \begin{pmatrix} a_{11} & a_{21} & a_{31} \\ a_{12} & a_{22} & a_{32} \\ a_{13} & a_{23} & a_{33} \end{pmatrix} \begin{pmatrix} P'_1 \\ P'_2 \\ P'_3 \end{pmatrix} \tag{10.11}$$

and since the two tristimuli match then

$$(S'_1 \ S'_2 \ S'_3) \begin{pmatrix} P'_1 \\ P'_2 \\ P'_3 \end{pmatrix} \equiv (S_1 \ S_2 \ S_3) \begin{pmatrix} P_1 \\ P_2 \\ P_3 \end{pmatrix} \tag{10.12}$$

The coefficients can be negative so the new primaries do not have to lie within the gamut of the old primaries.

**Table 10.2.** xyz-chromaticity of standard primaries and whites. The CIE primaries of Table 10.1 and the more recent ITU recommendation 709 primaries defined for HDTV.  $D_{65}$  is the white of a blackbody radiator at 6500 K, and  $E$  is equal-energy white

	$R_{\text{CIE}}$	$G_{\text{CIE}}$	$B_{\text{CIE}}$	$R_{709}$	$G_{709}$	$B_{709}$	$D_{65}$	$E$
$x$	0.7347	0.2738	0.1666	0.640	0.300	0.150	0.3127	0.3333
$y$	0.2653	0.7174	0.0089	0.330	0.600	0.060	0.3290	0.3333
$z$	0.0000	0.0088	0.8245	0.030	0.100	0.790	0.3582	0.3333

Substituting Eq. 10.11, equating tristimulus values and then transposing we obtain

$$\begin{pmatrix} S'_1 \\ S'_2 \\ S'_3 \end{pmatrix} = \begin{pmatrix} a_{11} & a_{21} & a_{31} \\ a_{12} & a_{22} & a_{32} \\ a_{13} & a_{23} & a_{33} \end{pmatrix}^T \begin{pmatrix} S_1 \\ S_2 \\ S_3 \end{pmatrix} = \mathbf{C} \begin{pmatrix} S_1 \\ S_2 \\ S_3 \end{pmatrix} \quad (10.13)$$

which is simply a linear transformation of tristimulus values.

Consider the concrete problem of transforming from CIE primaries to XYZ tristimulus values. We know from Table 10.2 the CIE primaries in terms of XYZ primaries

```
>> C = [ 0.7347, 0.2653, 0; 0.2738, 0.7174, 0.0088; 0.1666,
0.0089, 0.8245]
C =
    0.7347    0.2738    0.1666
    0.2653    0.7174    0.0089
         0     0.0088    0.8245
```

which is exactly the first three columns of Table 10.2. The transform is therefore

$$\begin{pmatrix} X \\ Y \\ Z \end{pmatrix} = \mathbf{C} \begin{pmatrix} R \\ G \\ B \end{pmatrix}$$

Recall from page 235 that luminance is contributed entirely by the Y primary. It is common to apply the constraint that unity  $R, G, B$  values result in unity luminance  $Y$  and a white with a specified chromaticity. We will choose  $D_{65}$  white whose chromaticity is given in Table 10.2 and which we will denote  $(x_w, y_w, z_w)$ . We can now write

$$\frac{1}{y_w} \begin{pmatrix} x_w \\ y_w \\ z_w \end{pmatrix} = \mathbf{C} \begin{pmatrix} J_R & 0 & 0 \\ 0 & J_G & 0 \\ 0 & 0 & J_B \end{pmatrix} \begin{pmatrix} 1 \\ 1 \\ 1 \end{pmatrix}$$

where the left-hand side has  $Y = 1$  and we have introduced a diagonal matrix  $\mathbf{J}$  which scales the luminance of the primaries. We can solve for the elements of  $\mathbf{J}$

$$\begin{pmatrix} J_R \\ J_G \\ J_B \end{pmatrix} = \mathbf{C}^{-1} \begin{pmatrix} x_w \\ y_w \\ z_w \end{pmatrix} \frac{1}{y_w}$$

Substituting real values we obtain

```
>> J = inv(C) * [0.3127 0.3290 0.3582]' * (1/0.3290)
J =
    0.5609
    1.1703
    1.3080
>> C * diag(J)
ans =
    0.4121    0.3204    0.2179
    0.1488    0.8395    0.0116
         0     0.0103    1.0785
```

The middle row of this matrix leads to the luminance relationship

$$Y = 0.1488R + 0.8395G + 0.0116B$$

which is similar to Eq. 10.10. The small variation is due to the different primaries used – CIE in this case versus Rec. 709 for Eq. 10.10.

The RGB tristimulus value of the redbrick was computed earlier and we can determine its XYZ tristimulus

```
>> XYZ_brick = C * diag(J) * RGB_brick';
ans =
    0.0351
    0.0224
    0.0039
```

which we convert to chromaticity coordinates by Eq. 10.9

```
>> tristim2cc(XYZ_brick')
ans =
    0.5729    0.3645
```

Referring to Fig. 10.12 we see that this *xy*-chromaticity lies in the red region and is named

```
>> colorname([0.5729    0.3645], 'xy')
ans =
    'englishred'
```

as might be expected for a “weathered red brick”.

### 10.2.6 What Is White?

In the previous section we touched on the subject of white. White is both the absence of color and also the sum of all colors. One definition of white is *standard daylight* which is taken as the mid-day Sun in Western/Northern Europe which has been tabulated by the CIE as illuminant  $D_{65}$ . It can be closely approximated by a blackbody radiator at 6500 K

```
>> d65 = blackbody(lambda, 6500);
>> lambda2xy(lambda, d65)
ans =
    0.3136    0.3241
```

which we see is close to the  $D_{65}$  chromaticity given in Table 10.2.

Another definition is based on white light being an equal mixture of all spectral colors. This is represented by a uniform spectrum

```
>> ee = ones(size(lambda));
```

which is also known as the equal-energy stimulus and has chromaticity

```
>> lambda2xy(lambda, ee)
ans =
    0.3333    0.3338
```

which is close to the defined value of ( $\frac{1}{3}$ ,  $\frac{1}{3}$ ).

## 10.3 Advanced Topics

In this section we will cover some advanced topics. The first is the effect of illumination on the apparent color of an object which is a very real problem for a robot using color cues in an environment with natural lighting. This leads to a discussion of the problem of white balancing. The next topic is an introduction to gamma encoding which is a very common non-linear relationship between tristimulus values and actual luminance. Finally we look at the distribution of colors in a real image and segment the image into regions of similar colors – this is a preview of techniques that we will cover in Chap. 12 and 13.

We adapt our perception of color so that the integral, or average, over the entire scene is grey. This works well over a color temperature range 5 000–6 500 K.

### 10.3.1 Color Constancy

Studies show that human perception of what is white is adaptive and has a remarkable ability to *tune out* the effect of scene illumination so that white objects always appear to be white. For example at night under a yellowish tungsten lamp the pages of a book still appear white to us, but a photograph of that scene viewed later under different lighting conditions will look yellow.

All of this poses real problems for a robot that is using color to understand the scene because the observed chromaticity varies with lighting. Outdoors the color of the morning or evening Sun (3 500 K) is different to that of the noon Sun (6 500 K), an overcast day is different to a clear day, and reflections from buildings or trees all conspire to change the illumination spectrum, and hence the luminance and color of the object. To illustrate this problem we revisit the red brick

```
>> lambda = [400:10:700]*1e-9;
>> R = loadspectrum(lambda, 'redbrick.dat');
```

under two different illumination conditions, the Sun at ground level

```
>> sun = loadspectrum(lambda, 'solar.dat');
```

and a tungsten lamp

```
>> lamp = blackbody(lambda, 2600);
```


and compute the *xy*-chromaticity for each case

```
>> xy_sun = lambda2xy(lambda, sun .* R)
xy_sun =
    0.4760    0.3784
>> xy_lamp = lambda2xy(lambda, lamp .* R)
xy_lamp =
    0.5724    0.3877
```


and we can see that the chromaticity, or apparent color, has changed significantly. These values are plotted on the chromaticity diagram in Fig. 10.15.

### 10.3.2 White Balancing

Photographers are well aware of the importance of illumination and refer to the color temperature of a light source – the equivalent black body temperature from Eq. 10.1. Compared to daylight an incandescent lamp appears more yellow, and a photographer



**Fig. 10.15.**  
Chromaticity of a red-brick  
under different illumination  
conditions



**Fig. 10.16.** The calibration target used for the Mars Rover's PanCam. Regions of known reflectance and chromaticity (red, yellow, green, blue and shades of grey) are used to set the white balance of the camera. The central stalk has a very low reflectance and also serves as a sundial. In the best traditions of sundials it bears a motto (photo courtesy NASA/JPL/Cornell/Jim Bell)

would use a blue filter (on the camera) to attenuate the red part of the spectrum to compensate. We can achieve a similar function by choosing the matrix  $J$

$$\begin{pmatrix} R' \\ G' \\ B' \end{pmatrix} = \begin{pmatrix} J_R & 0 & 0 \\ 0 & J_G & 0 \\ 0 & 0 & J_B \end{pmatrix} \begin{pmatrix} R \\ G \\ B \end{pmatrix}$$

to adjust the gains of the color channels. ▶ For example, boosting  $J_B$  would compensate for the lack of blue under tungsten illumination. This is the process of white balancing – ensuring the appropriate chromaticity of objects that we know are white (or grey).

Some cameras allow the user to set the color temperature of the illumination through a menu, typically with options for tungsten, fluorescent, daylight and flash which select different preset values of  $J$ . In manual white balancing the camera is pointed at a grey or white object and a button is pressed. The camera adjusts its channel gains  $J$  so that equal tristimulus values are produced  $R' = G' = B'$  which as we recall results in the desired white chromaticity. For colors other than white these corrections introduces some color error but this nevertheless has a satisfactory appearance to the eye. Automatic white balancing is commonly used and involves heuristics to estimate the color temperature of the light source but it can be fooled by scenes with a predominance of a particular color.


The most practical solution is to use the tristimulus values of three objects with known chromaticity in the scene. This allows the matrix  $C$  in Eq. 10.13 to be estimated directly, mapping the tristimulus values from the sensor to XYZ coordinates which are an absolute lighting-independent representation of surface reflectance. From this the chromaticity of the illumination can also be estimated. This approach is used for the panoramic camera on the Mars Rover where the calibration target shown in Fig. 10.16 can be imaged periodically to update the white balance under changing Martian illumination.

Typically  $J_G = 1$  and  $J_R$  and  $J_B$  are adjusted.

### 10.3.3 Color Change Due to Absorption

A final and extreme example of problems with color occurs underwater. For example consider a robot trying to find a docking station identified by colored targets. As discussed earlier in Sect. 10.1.1 water acts as a filter that absorbs more red light than blue light. For an object underwater this filtering affects both the illumination falling on

**Fig. 10.17.**  
Spectrum of the red brick when viewed underwater. The spectrum without the water absorption is shown in red



the object and the reflected light, the luminance, on its way to the camera. Consider again the red brick

```
>> [R,lambda] = loadspectrum([400:5:700]*1e-9, 'redbrick.dat');
```

which is now 1 m underwater and with a camera a further 1 m from the brick. The illumination on the water's surface is that of sunlight at ground level

```
>> sun = loadspectrum(lambda, 'solar.dat');
```

The absorption spectrum of water is

```
>> A = loadspectrum(lambda, 'water.dat');
```

and the total path length through the water is

```
>> d = 2
```

The transmission  $T$  is given by Beer's law Eq. 10.2.

```
>> T = 10 .^ (-d*A);
```

and the resulting luminance of the brick is

```
>> L = sun .* R .* T;
```

which is shown in Fig. 10.17. We see that the longer wavelengths, the reds, have been strongly attenuated. The apparent color of the brick is


```
>> xy_water = lambda2xy(lambda, L)
xy_water =
    0.3722    0.3813    0.2465
```

which is also plotted in the chromaticity diagram of Fig. 10.15. The brick appears much more blue than it did before. The reality underwater is more complex than this due to the scattering of light by tiny suspended particles. These result in additional color filtering of light on its way to the camera. They also reflect ambient light into the camera that has not been reflected from the target.

### 10.3.4 Gamma

In an old fashioned CRT monitor the luminance produced at the face of the display is non-linearly related to the control voltage  $V$  according to

$$L = V^\gamma \quad (10.14)$$



**Fig. 10.18.**  
The linear intensity wedge

where  $\gamma \approx 2.2$ . To correct for this non-linearity cameras generally apply the inverse non-linearity  $V = L^{1/\gamma}$  to their output signal which results in a system that is linear from end to end. ▶

Both operations are commonly referred to as gamma correction though more properly the camera-end operation is gamma encoding and the display-end operation is gamma decoding. ▶ LCD displays have a stronger non-linearity than CRTs but correction tables are applied within the display to make it follow the *standard*  $\gamma = 2.2$  behavior. ◀

To show the effect of display gamma we create a simple test pattern

```
>> wedge = [0:0.1:1];
>> idisp(wedge)
```

that is shown in Fig. 10.18 and is like a photographer’s *greyscale step wedge*. If we display this on our computer screen it will appear differently to the one printed in the book. We will most likely observe a large change in brightness between the second and third block – the effect of the gamma decoding non-linearity Eq. 10.14 in the display of your computer.

If we apply gamma encoding

```
>> idisp( wedge .^( 1/2.2 ) )
```

we observe that the intensity changes appear to be more linear ▶ and closer to the one printed in the book.

The chromaticity coordinates of Eq. 10.8 and Eq. 10.9 are computed as ratios of tristimulus values which are linearly related to luminance in the scene. The non-linearity applied to the camera output must be corrected, gamma decoded, *before* any colometric operations. The Toolbox function `igamma` performs this operation. Gamma decoding can also be performed when an image is loaded using the `'gamma'` option to the function `iread`.

Today most digital cameras ▶ encode images in sRGB format which uses the ITU Rec. 709 primaries and a gamma encoding function of

$$E' = \begin{cases} 12.92L, & L \leq 0.0031308 \\ 1.055L^{1/2.4} - 0.055, & L > 0.0031308 \end{cases}$$

which comprise a linear function for small values and a power law for larger values. The overall gamma is approximately 2.2.

The important property of colorspace such as *HSV* or *xyY* is that the chromaticity coordinates are invariant to changes in intensity. Many digital video cameras provide output in *YUV* or  $Y C_B C_R$  format which has a luminance component *Y* and two other components which are often mistaken for chromaticity coordinates – they are not. They are in fact color difference signals such that  $U, C_B \propto B' - Y'$  and  $V, C_R \propto R' - Y'$  where  $R', B'$  are gamma *encoded* tristimulus values, and  $Y'$  is gamma *encoded* intensity. The gamma nonlinearity means that *UV* or  $C_B C_R$  will not be a constant as overall lighting level changes.

The tristimulus values from the camera must be first converted to linear tristimulus values, by applying the appropriate gamma decoding, and then computing chromaticity. There is no shortcut.

Many cameras have an option to choose gamma as either 1 or 0.45 (= 1 / 2.2).

Gamma encoding and decoding are often referred to as gamma compression and gamma decompression respectively, since the encoding operation compresses the range of the signal, while decoding decompresses it.

Macintosh computers are an exception and prior to MacOS 10.6 used  $\gamma = 1.8$  which tends to make colors appear brighter and more vivid.

For a Macintosh prior to MacOS 10.6 use 1.8 instead of 2.2.

The JPEG file header (JFIF file format) has a tag `Color Space` which is set to either `sRGB` or `Uncalibrated` if the gamma or color model is not known. See page 289.

### 10.3.5 Application: Color Image

In this section we bring together many of the concepts and tools introduced in this chapter. We will also preview a number of functions that will be properly introduced in the next chapter. We consider a garden scene

```
>> flowers = imread('flowers4.png', 'double', 'gamma', 'sRGB');
```

shown in Fig. 10.19a comprising three different colored flowers and background greenery. Importantly we have applied gamma decoding so that the tristimulus values are proportional to the luminance of the original scene. The image `flowers` has 3 dimensions as shown in Fig. 12.2. The first two dimensions are the vertical and horizontal pixel coordinate, and the third is the color plane that selects the red, green or blue pixels.

We can convert the image to hue, saturation and value

```
>> hsv = colorspace('RGB->HSV', flowers);
```

and the result is another 3-dimensional matrix but this time the color planes represent hue, saturation and value. We can display hue

```
>> idisp( hsv(:,:,1) )
```

and saturation

```
>> idisp( hsv(:,:,2) )
```

as images which are shown in Fig. 10.14b and c respectively.

If we plot the chromaticity of each pixel as points in the chromaticity plane we would observe clusters of points corresponding to different parts of the scenes such as red flowers, yellow flowers, green leaves and shadows. We convert the color RGB image to an XYZ image

```
>> XYZ = colorspace('RGB->XYZ', flowers);
```

and then to *xy*-chromaticity coordinates

```
>> [x,y] = tristim2cc(XYZ);
```

where `x` and `y` are each images the same size as `flowers`. Next we compute a 2-dimensional histogram of (*x*, *y*) values with 100 bins in each dimension

```
>> xbins = [0 0.01 100]; ybins = [0 0.01 100];
>> [h,vx,vy] = hist2d(x, y, xbins, ybins);
```

where `h` is the number of points in each bin and `vx` and `vy` are the *x*- and *y*-coordinates of the corresponding bins. We display the histogram as a contour map overlaid on the *xy*-chromaticity diagram

```
>> xycolorspace
>> hold on
>> contour(vx, vy, h)
```


as shown in Fig. 10.19b. We see that there are 5 or 6 peaks: a broad peak in the red area, a narrow white peak and a number peaks in the green area.

Next we will perform unsupervised clustering using the k-means algorithm

```
>> [cls, cxy] = colorkmeans(flowers, 7);
```

where the second argument specifies the number of pixel chromaticity classes or clusters which we have set to seven. The k-means algorithm is iterative, it adjusts its estimate of the centre of each cluster and the assignment of pixels to clusters until equilibrium is reached. The initial cluster centres are chosen randomly which means that the function will give a different result each time it is run.

A limitation of k-means is that the number of clusters must be known in advance, typically guided by domain knowledge.



**Fig. 10.19.** Color image segmentation. **a** Original color image; **b**  $xy$ -chromaticity plane with overlaid frequency contours; **c** label image in false color; **d** cluster centroids on chromaticity diagram; **e** class 5 pixels (yellow) only

The result is another image `cls` where each pixel value indicates the cluster to which the corresponding pixel in `flowers` has been assigned – its color classification, or class, which is an integer in the interval 1 to 7. We can display this

```
>> idisp(cls, 'colormap', 'jet', 'nogui')
```

as shown in Fig. 10.19c. In this case class 5 corresponds to yellow flowers and we can display just those pixels

```
>> idisplabel(flowers, cls, 5)
```

as shown in Fig. 10.19e. In addition to the yellow flowers there are many very small groups of dark pixels that correspond to parts of the foliage – these have the same chromaticity as the flowers but a low luminance.

The `colorkmeans` function also returns the centre of the clusters and these are

```
cxy =
  0.6082  0.5378  0.3328  0.3466  0.4017  0.3176  0.3238
  0.3327  0.3436  0.4599  0.5195  0.4874  0.3788  0.5744
```

in *xy*-space. We plot these on the *xy*-chromaticity diagram

```
>> xycolospace
>> plot_point(cxy, '*', 'sequence', 'textsize', 10, 'textcolor', 'b')
```

as shown in Fig. 10.19d and we can see that class 5 does indeed lie in the yellow area, and also that classes 1 and 2 are red, class 6 is white, while the rest are different greens.

Finally we convert these cluster centres from *xy*-coordinates to human meaningful names

```
>> colorname(cxy, 'xy')
ans =
  Columns 1 through 5
  'cadmiumreddeep'  'brown3'  'olive'  'terreverte'  'yellow4'
  Columns 6 through 7
  'darkseagreen4'  'yellowgreen'
```

The color names for class 1 “cadmiumreddeep” and class 4 “yellow4” corresponding to the red and yellow flowers respectively are quite apt. The color name for the white flowers “darkseagreen4” is surprising, implying dark-green rather than white, but the *xy*-chromaticity of this color is quite close to the white point. Dark-green has low luminance but the color name matching is based on chromaticity not luminance.

---

## 10.4 Wrapping Up

We have learnt that the light we observe is a mixture of frequencies, a continuous spectrum, which is modified by reflectance and absorption. The spectrum elicits a response from the eye which we interpret as color – for humans the response is a tristimulus, a 3-vector that represents the outputs of the three different types of cones in our eye. A digital color camera is functionally equivalent. The tristimulus can be considered as a 1-dimensional brightness coordinate and a 2-dimensional chromaticity coordinate which allows colors to be plotted on a plane. The spectral colors form a locus on this plane and all real colors lie within this locus. The three primary colors form a triangle on this plane which is the gamut of those primaries. Any color within the triangle can be matched by an appropriate mixture of the primaries. No set of primaries can define a gamut that contains all colors. An alternative set of imaginary primaries, the CIE XYZ system, does contain all real colors and is the standard way to describe colors. Tristimulus values can be transformed using linear transformations to account for different sets of primaries. Non-linear transformations can be used to describe tristimulus values in terms of human-centric qualities such as hue and saturation.

We also discussed definition of white, the problem of white balancing, the non-linear response of display devices and how this effects the common representation of images and video. Finally we used chromaticity information to classify pixels in a colorful real-world image.

### Further Reading

At face value color is a simple concept that we learn in kindergarten but it is actually a complex topic. In this chapter we have only begun to scrape the surface of photometry and colorimetry. Photometry is the part of the science of radiometry concerned with measurement of visible light. It is challenging for engineers and computer scientists since it makes use of uncommon units such as lumen, steradian, nit, candela and lux. One source of complexity is that words like intensity and brightness are synonyms in everyday speech but have very specific meanings in photometry. Colorimetry is the science of color perception and is also a large and complex area since human perception of color depends on the individual observer, ambient illumination and even the field of view. Colorimetry is however critically important in the design of cameras, computer displays, video equipment and printers.

The computer vision textbooks by Gonzalez and Woods (2008) and Forsyth and Ponce (2002) each have a discussion on color and color spaces. The latter also has a discussion on the effects of shading and inter-reflections. Comprehensive online information about computer vision is available through CVonline at <http://homepages.inf.ed.ac.uk/rbf/CVonline>, and the material in this chapter is covered under the sections *Image Physics* and *Sensors and their Properties*.

Readable and comprehensive books on color science include Hunt (1987) and from a television or engineering perspective Benson (1986). A more conversational approach is given by Hunter and Harold (1987), which also covers other aspects of appearance such as gloss and lustre. The CIE standard (Commission Internationale de l'Éclairage 1987) is definitive but hard reading. The work of the CIE is ongoing and its standards are periodically updated at [www.cie.co.at](http://www.cie.co.at). The color matching functions were first tabulated in 1931 and revised in 1964.

Charles Poynton has for a long time maintained excellent online tutorials about color spaces and gamma at <http://www.poynton.com>. His book (Poynton 2003) is an excellent and readable introduction to these topics while also discussing digital video systems in great depth. Gamma is also described online at <http://www.w3.org/Graphics/Color/sRGB>.

Other MATLAB® tools include the ColorLab toolbox at [http://cs.joensuu.fi/colorlab\\_toolbox/](http://cs.joensuu.fi/colorlab_toolbox/) and the `colorspace` function at <http://www.math.ucla.edu/~getreuer/colorspace.html>.

**Infra-red cameras.** Consumer cameras are functionally equivalent to the human eye and are sensitive to the visible spectrum. Cameras are also available that are sensitive to infra-red and a number of infra-red bands are defined by CIE: IR-A (700–1 400 nm), IR-B (1 400–3 000 nm), and IR-C (3 000 nm–1 000 μm). In common usage IR-A and IR-B are known as near infra-red (NIR) and short-wavelength infra-red (SWIR) respectively, and the IR-C subbands are medium-wavelength (MWIR, 3 000–8 000 nm) and long-wavelength (LWIR, 8 000–15 000 nm). LWIR cameras are also called thermal or thermographic cameras. **Ultraviolet cameras** typically work in the near ultra-violet region (NUV, 200–380 nm) and are used in industrial applications such as detecting corona discharge from high-voltage electrical systems.

**Hyperspectral cameras** have more more than three classes of photoreceptor, they sample the incoming spectrum at many points typically from infra-red to ultra-violet and with tens or even hundreds of spectral bands. Hyperspectral cameras are used for applications including aerial survey classification of land-use and identification of the mineral composition of rocks.

**Table 10.3.**  
Various spectra provided with the Toolbox. Relative luminosity values lie in the interval [0,1], and relative spectral power distribution (SPD) are normalized to a value of 1.0 at 550 nm

Filename	Units	Description
cones.dat	Rel. luminosity	Spectral response of human cones
photopic.dat	Rel. luminosity	CIE 1924 photopic response
scotopic.dat	Rel. luminosity	CIE 1951 scotopic response
redbrick.dat	Reflectivity	Reflectivity spectrum of a weathered red brick
solar.dat	$\text{W m}^{-2} \text{nm}^{-1}$	Solar spectrum at ground level
water.dat	$1 \text{ m}^{-1}$	Light absorption spectrum of water
D65.dat	Rel. SPD	CIE standard $D_{65}$ illuminant

### Data Sources

The Toolbox contains a number of data files describing various spectra which are summarized in Table 10.3. Each file has as its first column the wavelength in metres. The files have different wavelength ranges and intervals but the helper function `loadspectrum` interpolates the data to the user specified range and sample interval.

Several internet sites contain spectral data in tabular format and this is linked from the book's web site. This includes reflectivity data for many materials provided by NASA's online ASTER spectral library and the Spectral Database from the University of Eastern Finland Color Research Laboratory. Data on cone response and CIE color matching functions is available from the Colour & Vision Research Laboratory at University College London. CIE data is also available online.

### Exercises

1. You are a blackbody radiator! Plot your own blackbody emission spectrum. What is your peak emission frequency? What part of the EM spectrum is this? What sort of sensor would you use to detect this?
2. Consider a sensor that measures the amount of radiated power  $P_1$  and  $P_2$  at wavelengths  $\lambda_1$  and  $\lambda_2$  respectively. Write an equation to give the temperature  $T$  of the blackbody in terms of these quantities.
3. Using the Stefan-Boltzman law compute the power emitted per square metre of the Sun's surface. Compute the total power output of the Sun.
4. Use numerical integration to compute the power emitted in the visible band 400–700 nm per square metre of the Sun's surface.
5. Why is the peak luminosity defined as  $683 \text{ lm W}^{-1}$ ?
6. Given typical outdoor illuminance as per page 229 determine the luminous intensity of the Sun.
7. Sunlight at ground level. Of the incoming radiant power determine, in percentage terms, the fraction of infra-red, visible and ultra-violet light.
8. Use numerical integration to compute the power emitted in the visible band 400–700 nm per square metre for a tungsten lamp at 2 600 K. What fraction is this of the total power emitted?
9. Plot and compare the human photopic and scotopic spectral response.
10. Can you create a metamer for the red brick?
11. Prove the center of gravity law.
12. On the  $xy$ -chromaticity plane plot the locus of a blackbody radiator with temperatures in the range 1 000–10 000 K.
13. Plot the XYZ primaries on the  $rg$ -plane.



**Fig. 10.20.** The Gretag Macbeth Color Checker is an array of 24 printed color squares, which includes different greys and colors as well as spectral simulations of skin, sky, foliage etc. Spectral data for the squares is available online via <http://www.cis.rit.edu/research/mcsl/online/cie.php>

14. The Gretag Macbeth Color Checker shown in Fig. 10.20 is an array of 24 printed color squares. Spectral data for the Color Checker is available at <http://www.rmimaging.com/information/colorchecker.html>. Compute and plot the  $xy$ -chromaticity for each square.
15. For Fig. 10.11 determine the chromaticity of the feasible green. Determine the brightest possible tristimulus value assuming that the value of any primary lies in the range  $[0, 1]$ .
16. Modify the function `xycolorspaces` to generate an  $rg$ -chromaticity plane.
17. Determine the tristimulus values for the red brick using the Rec. 709 primaries.
18. Take a picture of a white object using incandescent illumination. Determine the average RGB tristimulus value and compute the  $xy$ -chromaticity. How far off white is it? Determine the color balance matrix  $J$  to correct the chromaticity. What is the chromaticity of the illumination?
19. What is the name of the color of the red brick when viewed underwater (page 242).
20. Image a target like Fig. 10.16 that has three colored patches of known chromaticity. From their observed chromaticity determine the transform from observed tristimulus values to Rec. 709 primaries. What is the chromaticity of the illumination?
21. Consider an underwater application where a target  $d$  metres below the surface is observed through  $m$  metres of water, and the water surface is illuminated by sunlight. From the observed chromaticity can you determine the true chromaticity of the target? How sensitive is this estimate to incorrect estimates of  $m$  and  $d$ ? If you knew the true chromaticity of the target could you determine its distance?
22. Is it possible that two different colors look the same under a particular lighting condition? Create an example of colors and lighting that would cause this?
23. Use one of your own pictures and repeat the exercise of Sect. 10.3.5. Can you distinguish different objects in the picture?
24. Show analytically or numerically that scaling a tristimulus value has no effect on the chromaticity. What happens if the chromaticity is computed on gamma encoded tristimulus values?