Chapter 6

Colour

6.1 Visible Light and Colour

In a classic experiment, Isaac Newton passed sunlight through a triangular glass prism and decomposed it into a rainbow of colours. He then showed that the colours of the rainbow could not be decomposed further and that when they were recombined, using a second prism, they produced the original white sunlight.

This experiment was the beginning of the study which lead to our present understanding that visible light is produced by a small section of a larger electromagnetic spectrum. We now know that electromagnetic radiation is well described by a dual wave/particle theory; with wavelength determining how energetic the radiation is (long wavelengths are less energetic than short ones).

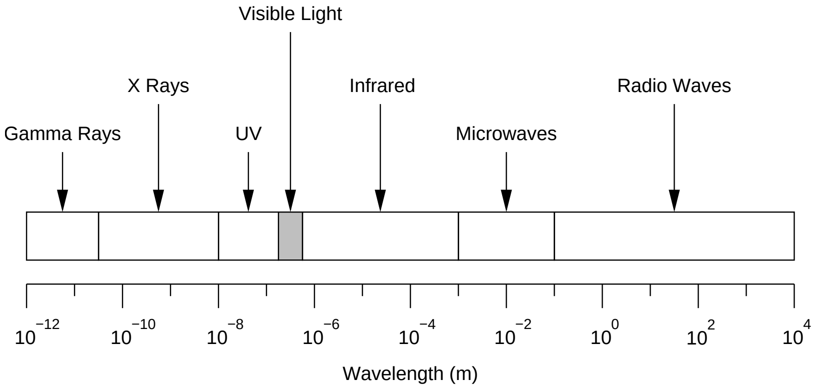

Figure 6.1: The electromagnetic spectrum.

Figure 6.1 shows a broad band of the electromagnetic spectrum. Long wavelengths correspont to radio waves of various kinds and slightly shorter wavelengths to microwaves and infrared radiation. Short wavelengths correspond to ultra-violet light, X-rays, and high-energy gamma rays. Between these extremes lies a comparatively narrow band containing visible light.

The visible part of the spectrum falls in a small interval centred at about 6×10^{-7} meters. To avoid dealing with such very small numbers, light wavelengths are traditionally measured in nanometres (nm), or 10^{-9} meters. Figure 6.2 shows an enlarged view of the visible spectrum, showing the wavelength bands corresponding to various colours.

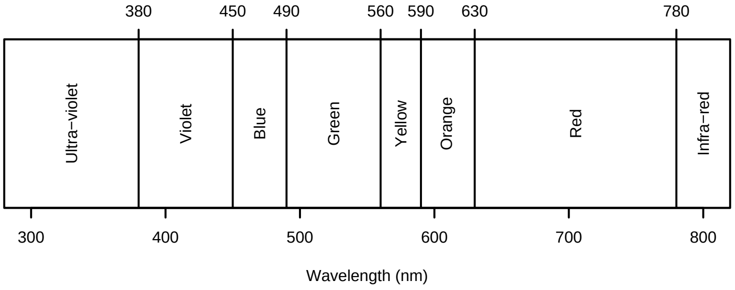

Figure 6.2: The visible portion of the electromagnetic spectrum.

When Newton decomposed white light into its spectrum, he claimed to see the seven colours: red, orange, yellow, green, blue, *indigo*, and violet. These days, the inclusion of indigo in this list seems contrived. It is believed that Newton chose to have seven colours to parallel the seven notes of the standard musical scale. We will see later that four or six are more natural numbers of colours to use to describe the rainbow.

While it is scientifically natural relate colour to wavelength, it does not entirely correspond to the way in we perceive colour. Our perception seems richer and more complex than a purely linear arrangement would suggest. For example, in terms of wavelength, red and violet are the two most different visible colours, but we perceive them as being quite similar or somehow related. During his experiments with light, Newton noticed this and was the first person to arrange the visible colours into a colour wheel, by wrapping the red end of the spectrum around to the violet.

6.2 Low-light and Daylight Vision

Humans possess two parallel vision systems; one for use in very low light levels, and one for normal daytime light levels. These vision systems are driven by different receptor cells in the retina.

The Low-light or *scotopic* vision system is driven by the rod-cells of the eye, (these were described in section 4.2.2). There is only one kind of rod cell in the retina and they are sensitive only to the presence or absence of light. This means that our scotopic vision systems are not able to differentiate between differing light wavelengths, and hence, scotopic vision provides no sensation of colour.

The normal light or *photopic* vision system is driven by the cone cells in the retina. In individuals with normal vision, there are three different kinds of cone-cells. Each type responds (best) to different light wavelengths. This means that our daytime vision can differentiate between different light wavelengths, and it is this which creates the sensation of colour.

The scotopic and photopic vision systems are at their peak performance over different ranges of wavelengths. The scotopic system has its peak efficiency at about 500 nm, which corresponds to a blue-green colour, and the system has very little sensitivity

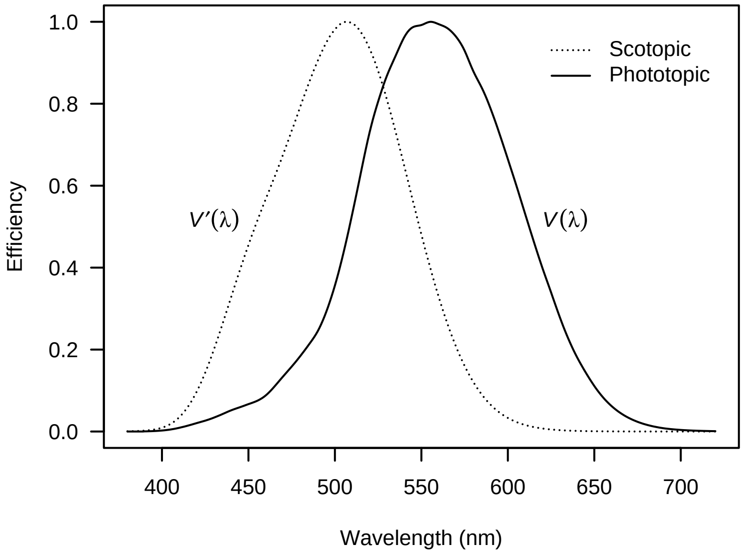

Figure 6.3: The spectral luminous efficiency curves.

to red light. The photopic system achieves its peak efficiency at 560 nm, which is at the boundary between green and yellow.

The relative efficiency of the two vision systems is shown in figure 6.3. The curves $V'(\lambda)$ and $V(\lambda)$ give the intensity of light of wavelength λ as perceived by the scotopic and photopic systems. The curves have been normalised so that their maximum value is one. In absolute terms, the scotopic system is far more sensitive than the photopic one. It is so sensitive that it becomes saturated at even very low light levels and is shut down. Activating the scotopic system takes a long time (up to an hour), and this is why it can be counter-productive to use a flashlight when carrying out night activities — you are left effectively blind for a long time after the light is switched off.

It is possible to see the effect of the different sensitivity of the scotopic and photopic systems by looking at adjacent areas of blue and red at sundown. In normal light red will appear brighter than blue but, as daylight fades, the blue will appear to brighten and glow and the red will fade. This shift in wavelength sensitivity is known as the *Purkinje shift*.

6.3 Colour Vision

There are three types of cone cells in the retina. These provide the basis of normal vision. The cones provide different sensitivities to light of different wavelengths. The three different cone types are traditionally known as red, green and blue cones to emphasise their sensitivity to differing wavelength bands. This is actually a misnomer as the red and green cones peak in the yellow and yellow-green regions of the spectrum. Because of this they are also known as L, M and S cones (for long, medium and short wavelength), or sometimes they are labelled with the Roman letters R, G, and B or the Greek letters ρ, γ, and β (rho, gamma, and beta). Figure 6.4 shows how the rela-

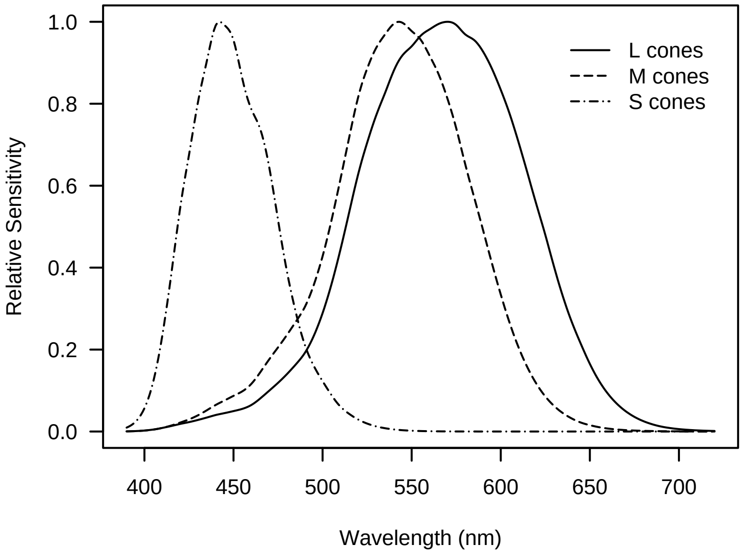

Figure 6.4: Spectral sensitivity curves for the three cone types.

tive sensitivity of each cone type varies with wavelength (again, the curves have been normalised to have a maximum of one).

The relative abundances of the L, M and S cones is in the ratio 40:20:1. The L and M cones are overwhelmingly concentrated in the fovea. There are clearly many fewer S cones, and they are more uniformly distributed through the retina.

The cones act as basic colour receptors for the vision system. There is evidence that a further processing step takes place before the received image is transmitted to the brain. This step consists of encoding the cone output in the form of *opponent colours*.

Suppose that the signals from the L, S and M cones are ρ, γ and β. The overall level of light received by the eye (ignoring colour) is encoded as the *achromatic signal*

$$A = 2\rho + \gamma + \beta/20$$

(because of the relative abundances of the three cone types). The colour content of the signal is encoded as *colour difference* signals. There are three possible colour differences

$$C_1 = \rho - \gamma, \qquad C_2 = \gamma - \beta, \qquad C_3 = \beta - \rho.$$

To transmit all three of these signals would be redundant because if two of them are know, the remaining one can be computed because $C_1 + C_2 + C_3 = 0$. There is evidence to suggest that the transmitted signals are, in fact,

$$C_1 = \rho - \gamma, \qquad \text{and} \qquad C_3 - C_2 = \beta - \rho - (\gamma - \beta) = 2\beta - (\rho + \gamma).$$

These two signals measure colour along red-green and blue-yellow axes.

It is interesting to note this corresponds to the theory of *opponent colours* described by Leonardo da Vinci, Goëthe, and others. The theory is based on the notion of opponent colours, which cannot exist together in a mixture. We don’t see colours which are “reddish-green” or “yellowish-blue”.

Although the RGB and opponent colour descriptions provide a full description of colour, they don’t correspond the the way in which we naturally think about colour. When equal parts of red and green are presented to the eye, we recognise the resulting colour as yellow, not as a mixture of red and green.

Our natural description of colour is in terms of *hue, purity* and *lightness*.

- Hue is the property of colour which allows us to distinguish red, yellow, green and blue. The hues in a a spectrum produced by the refraction of light through a glass prism correspond naturally to wavelength.
- Purity measures how close a given colour is to a pure spectral colour. Pink is “less pure” than red because it is obtained by mixing red with white, and so it contains more than just one light wavelength.
- Brightness measures the overall intensity of a colour. It is possible to see colours with the same hue and purity in brighter or dimmer forms (as though the position of a dimmer switch were being changed). These are different colours and they differ in brightness.

The naturalness of the hue, purity and lightness description of colour suggests that it is likely that the opponent colour signals transmitted from the eye to the brain are later converted to this form.

6.4 Computer Generated Colour

Computer displays produce colour by stimulating the cones of eyes with varying amounts of three primary colours — red, green and blue. For CRT displays these colours are emitted when phosphors coating the display screen are bombarded by electrons, while for more modern flat displays, the colours are produced by light emitting diodes.

The three primary colours used on computer displays produce a very wide range (or gamut) of colours. It is not possible to produce every colour which the eye can see, but selection is wide enough for producing most colour images. The primaries used are close to pure spectral colours with wavelengths 610nm, 540nm and 465nm; but are somewhat less saturated than spectral colours. Graphics display systems are calibrated so that when the maximum amounts of the three primary colours are combined, the result is white.

A inherent difficulty with colour displays is that the colours they produce depend on the colours of the particular primaries used. For many purposes this is unimportant, but if exact colour matching is crucial then a more sophisticated colour description model is needed. If exact colour matching is not important, then there are a number of alternative ways in which colour can be described.

6.4.1 The RGB Colour Model

The RGB model is the most direct colour specification system. It works by directly specifying the levels of the red, green, and blue primaries which are to be used to produce a colour.

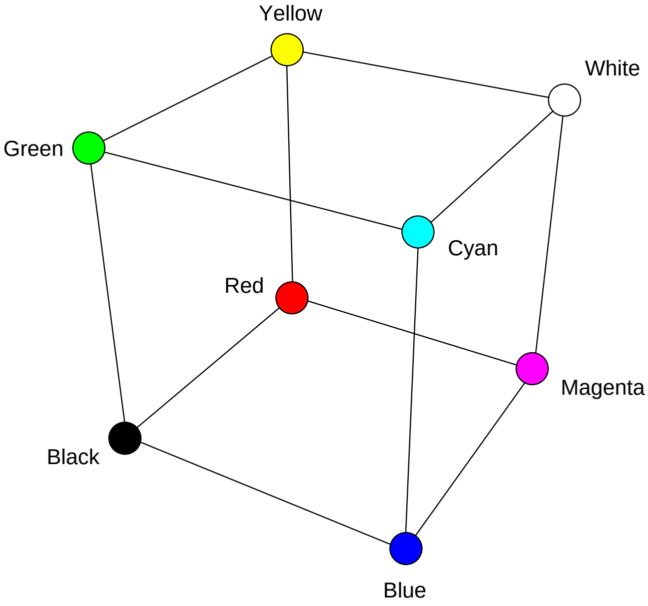

Figure 6.5: The RGB colour cube.

An RGB colour specification typically uses a value between 0 and 1 to indicate the intensity for each of the three primary colours. The set of possible colours which can be created in this way can be thought of as a “colour cube” as shown in figure 6.5.

A common alternative specification for primary intensity is as an integer value in the range 0 to 255. In the latter case it is common to write the integer colour specification as a two digit hexadecimal value. For example, the value `FFFF00` (corresponding to $(255, 255, 0)$) implies full intensity of the red and green primaries and no blue primary. This combination of primaries will produce the colour yellow.

The most important colours in the RGB specifications are shown in the table below.

Colour Name	**Hexadecimal**	**RGB Intensities**
Black	`000000`	$(0,\ 0,\ 0)$
Red	`FF0000`	$(1,\ 0,\ 0)$
Green	`00FF00`	$(0,\ 1,\ 0)$
Blue	`0000FF`	$(0,\ 0,\ 1)$
Yellow	`FFFF00`	$(1,\ 1,\ 0)$
Cyan	`00FFFF`	$(0,\ 1,\ 1)$
Magenta	`FF00FF`	$(1,\ 0,\ 1)$
White	`FFFFFF`	$(1,\ 1,\ 1)$
Gray	`808080`	$(.5,\ .5,\ .5)$

Although the RGB model provides a reasonable description of all the colours which can be produced with a standard colour display, it is not intuitive to use. For example, it is not obvious how to describe a rainbow in RGB coordinates.

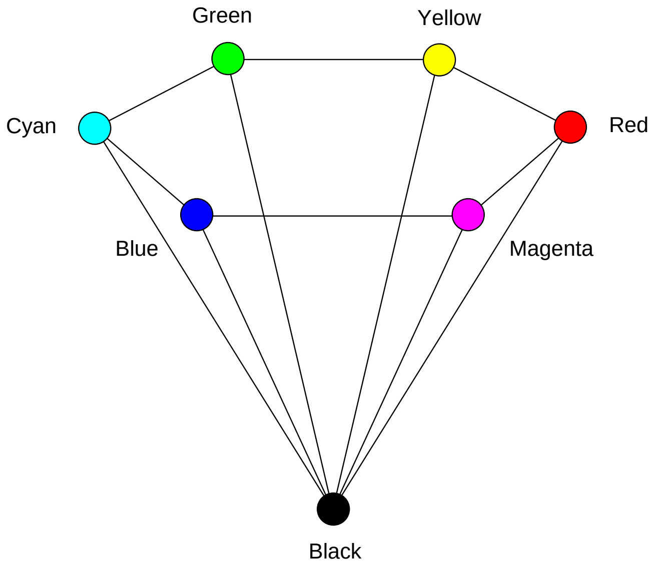

Figure 6.6: The HSV colour hexcone.

6.4.2 The HSV Colour Model

An important alternative colour specification system is based on the intuitive ideas of hue, saturation and lightness. The most common of these is the HSV system, with HSV an abbreviation for *Hue*, *Saturation*, and *Value*.

Under the HSV specification, the hues of the rainbow are arranged in a circle with the value of hue ranging through the values 0 to 1 corresponding to angles 0 to 2π around the circle. The values for saturation range from 0 to 1, as colours range from white to a fully saturated hue. The final parameter, value, ranges from 0 to 1 as colours range from black to full intensity.

The HSV system is very intuitive and simple to use. Most image processing and drafting packages provide a colour control in this form.

6.5 Colour Theory

Different light wavelengths stimulate the three kinds of cones in the eye to different degrees, and it is this that provides our perception of the different colours associated with each wavelength. When we are presented with the rich light spectra of objects in the real world we see them as having distinct colours.

Although exposing the eye to a rich mixture of light wavelengths does produce a clear colour sensation, there is a simpler way of producing different colour sensations. By choosing three “primary” colours (traditionally red, green and blue) and stimulating the cones with varying amounts of these primaries it is possible to produce the same sensation of colour as is produced by differing wavelengths.

In the early part of the twentieth century, this observation was used as the basis for an extensive study of human colour perception. Much of the information gained during

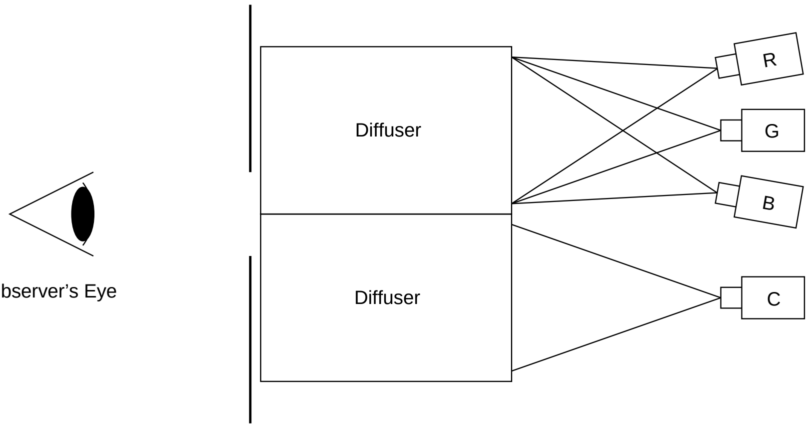

Figure 6.7: The type of apparatus used in colour matching experiments.

this study forms the basis for current theories about colour, and also underlies much of the technology used for colour manipulation and reproduction.

One of the problems attacked by the early colour scientists was that of *colour matching* — determining the amounts of the primaries needed to match the colour produced by each of the wavelengths of visible light. These matching experiments were carried out with the type of apparatus shown in figure 6.7. In these experiments, a pure colour is shown on one half of a screen and the subject is asked match this with a colour created by mixing varying amounts of red, green, and blue primaries and shining this on the half of the screen. By varying the wavelength of the colour being matched, it is possible to obtain amount of red, green, and blue primaries required for a match as a function of wavelength λ. These functions of λ are denoted $\bar{r}(\lambda)$, $\bar{g}(\lambda)$, and $\bar{b}(\lambda)$, and are known as the *colour matching functions* of the primaries.

The results of a colour matching experiment are specific to the particular set of primaries used, and a unique set of colour matching functions will thus be produced for each set of primaries. Figure 6.8 shows the results of a colour matching experiment [9] with primaries chosen at the wavelengths 645.16 nm (R), 526.32 nm (G), and 444.44 nm (B). Note that more red primary is required than green or blue. This is because the eye is less sensitive at its wavelength than it is at the wavelengths of the other two primaries.

It is clear that the matching functions are negative in some regions. This is because it is not possible to match all pure spectral colours with these three primaries. When the value for a particular primary is negative, that primary must be added to the colour being matched, rather than to the other primaries.

The colours we usually encounter are not the pure spectral colours of the rainbow, but are composed of a mixture of such colours (just as Newton's white light was composed of a mix of rainbow hues). Such a colour is described by its spectral distribution function $S(\lambda)$, which for each wavelength λ is proportional to the amount of power which the light radiates in a small interval about λ. The total amount of light being radiated is

$$E = \int S(\lambda) d\lambda$$

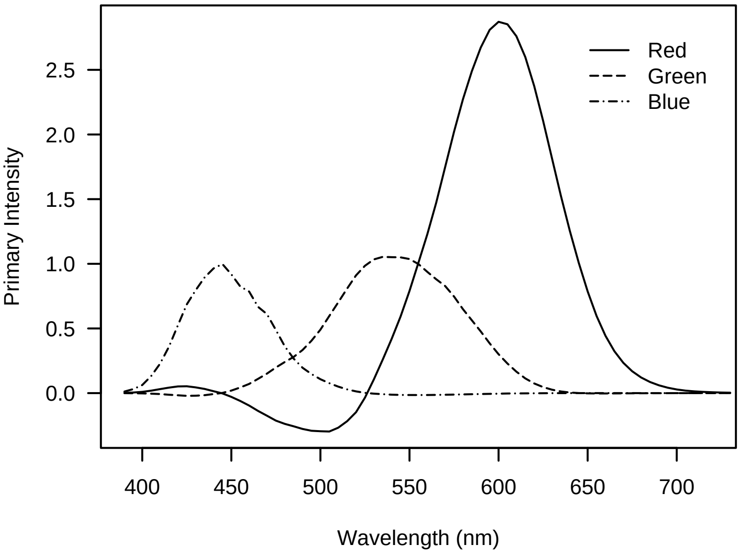

Figure 6.8: RGB colour matching functions.

Given a particular light spectrum, the colour matching functions can be used to compute how much of each of the RGB the primaries are required to match that colour. For a spectrum $S(\lambda)$, the amounts are

$$R = \int S(\lambda)\,\overline{r}(\lambda)d\lambda, \qquad G = \int S(\lambda)\,\overline{g}(\lambda)d\lambda, \qquad B = \int S(\lambda)\overline{b}(\lambda)d\lambda.$$

The values R, G, and B are called the *tristimulus values* for the colour.

6.6 The Chromaticity Diagram

In the 1930's a standard was developed for the description of colour by the *Commission Internationale de l'Éclairage* (the CIE). The use of colour matching functions as a basis for colour description seemed useful, and a *standard observer* model was adopted. The standard obsever was defined by an idealized set of colour matching functions for a given set of RGB primary colours. The matching functions were broadly representative of those actually observed for a variety of individuals.

It was decided that fact that the RGB matching functions took on negative values was likely to cause confusion and errors, and it was decided to replace the RGB primaries by a set of "imaginary" primary colours which corresponded to matching functions which were always positive.

This is a purely mathematical description, but it has the advantage that all visible colours can be obtained as a simple mixture of the primaries. The CIE primaries are know as X, Y, and Z. Their defining colour matching functions $\overline{x}(\lambda)$, $\overline{y}(\lambda)$, and $\overline{z}(\lambda)$ are shown in figure 6.9. The $\overline{y}(\lambda)$ matching function is identical in shape to the scotopic luminous efficiency $V'(\lambda)$. It can therefore be taken as providing brightness information about colours.

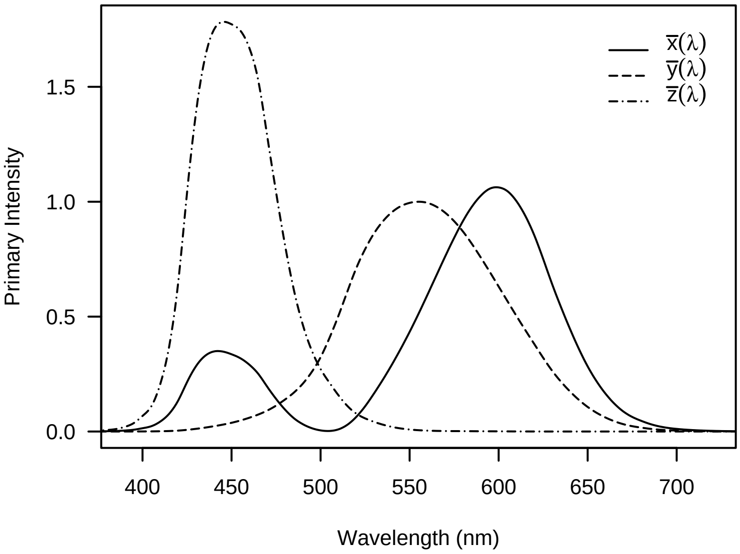

Figure 6.9: XYZ colour matching functions.

As in the RGB case, it is possible to compute tristimulus values X, Y, and Z.

$$X = \int S(\lambda)\,\bar{x}(\lambda)d\lambda, \qquad Y = \int S(\lambda)\,\bar{y}(\lambda)d\lambda, \qquad Z = \int S(\lambda)\,\bar{z}(\lambda)d\lambda.$$

These values characterise the colour being viewed.

The tristimulus values provide a complete description of colour, but they mix the effects of brightness and hue. It is conventional to transform them into *chromaticity* values which separate out the hue information from the brightness. The chromaticity values are defined by

$$x = \frac{X}{X+Y+Z}, \qquad y = \frac{Y}{X+Y+Z}, \qquad z = \frac{Z}{X+Y+Z}.$$

Because the chromaticities have been normalised in this way, they represent information about colour which is independent of brightness. Because $x+y+z=1$, they have the added advantage that one of them can be dropped without loss of information. By convention, it is the z value which is dropped, and colour information is usually presented as a plot of y against x. This is called a *chromaticity diagram*.

Figure 6.10 gives a chromaticity diagram showing the locations of the pure spectral colours (the *spectral locus*). The rightmost point of the spectral locus corresponds to red, and the lowest point to violet. The regions corresponding to other colours can be determined by comparing figure 6.10 with figure 6.2.

The line joining the two ends of the spectral locus is called the *purple boundary*. The region enclosed by the spectral locus and the purple boundary contains the complete set of visible colours. Colours in the interior of the visible region are mixtures of pure spectral colours. The chromaticity diagram provides an excellent tool for describing such mixtures as we will see in the next section.

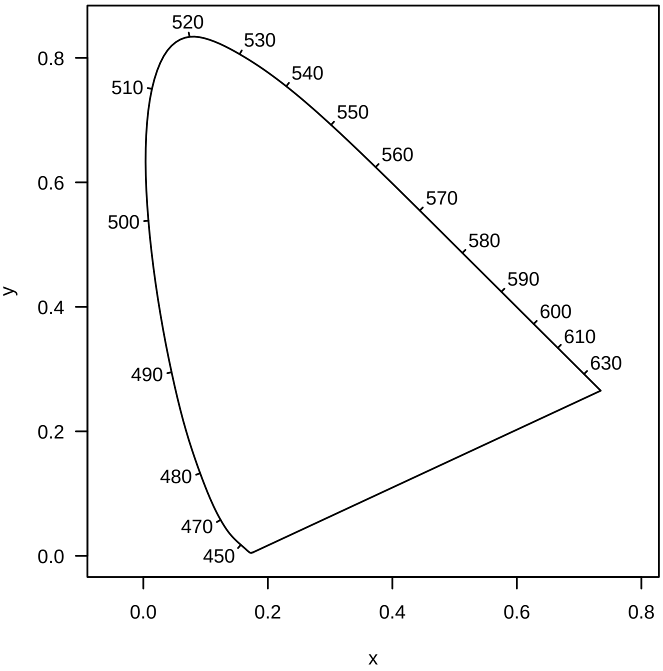

Figure 6.10: The chromaticity diagram showing the spectral locus.

The point at the centre of the diagram, with $x = 1/3, y = 1/3$ and $z = 1/3$, is called the equal stimulus point and denoted S_E. It is the point corresponding to the detection of equal amounts of the three primaries. It is tempting to think of this as the point where “white” is located, but the concept of white is more complex, because it is influenced by later processing of the visual signal.

In general, it is very important to remember that chromaticity diagrams show relationships between colour stimuli, not colour perceptions. Exact colour perceptions depend strongly on the viewing conditions, and on the adaptation and other characteristics of the observer.

6.7 Manipulation of Chromaticities

The chromaticity diagram is particularly useful because it provides a way of describing and manipulating colour mixtures in an intuitive way, but which has a firm analytical basis.

Suppose that we have two coloured lights C_1 and C_2 with tristimulus values (X_1,Y_1,Z_1) and (X_2,Y_2,Z_2). If we set $S_1 = X_1+Y_1+Z_1$ and $S_2 = X_2+Y_2+Z_2$, then we can write the chromaticities as:

$$(x_1,y_1,z_1) = \frac{1}{S_1}(X_1,Y_1,Z_1),$$

and

$$(x_2,y_2,z_2) = \frac{1}{S_2}(X_2,Y_2,Z_2).$$

If we mix the two colours by shining the two lights simultaneously, we obtain a colour C_M with tristimulus values $(X_M, Y_M, Z_M) = (X_1 + X_2, Y_1 + Y_2, Z_1 + Z_2)$, and chromaticities

$$(x_M, y_M, z_M) = \frac{1}{S_1 + S_2}(X_1 + X_2, Y_1 + Y_2, Z_1 + Z_2).$$

This means that the chromaticites must satisfy

$$\begin{aligned}(x_M, y_M, z_M) &= \frac{S_1}{S_1 + S_2}(x_1, y_1, z_1) + \frac{S_2}{S_1 + S_2}(x_2, y_2, z_2) \\ &= \alpha(x_1, y_1, z_1) + (1 - \alpha)(x_2, y_2, z_2),\end{aligned}$$

showing that the chromaticity for the mixture must lie on the line segment joining the individual chromaticities, with the position on the line segment depending on the relative brightness of the two colours being mixed.

6.8 Using the Chromaticity Diagram

Many operations on chromaticities involve the use of the *white point* on chromaticity diagram. This is the point corresponding to the colour perceived as white — the colour with no perceptual bias toward any spectral colour. The white point is determined by computing the chromaticity of the ambient light in which viewing takes place. Under normal viewing conditions, this will generally be close to the equal stimulus point $S_E = (1/3, 1/3)$.

Any colour in the chromaticity diagram lies on a line from the white-point to the boundary of the diagram. This means that any colour can be repesented as a mixture of white and a pure spectral colour (or a purple). This characterization of colours by their position relative to the white point and the spectral locus is a very natural one. Figure 6.11 shows a colour C and its relation to the white-point W and a pure spectral colour H. The wavelength of H is called the *dominant wavelength* of the colour C and the relative distance of C from W to H is called its *excitation purity*.

Dominant wavelength and excitation purity correspond directly to the intuitive notions of the *hue* and *saturation* of a colour. Colours which lie on the spectral locus or purple boundary are as pure or intense as they can possibly be, and are referred to as saturated. Less saturated colours are otained by mixing pure spectral colours with white to obtain “tints”.

The CIE diagram also provides a way to formalize the notion of complementary colours. Artists have known for a long time that there are pairs of colours which provide a vibrant contrast in graphic designs, and referred to them as complementary. Examples of such colour pairs are yellow/blue, red/green and orange/purple. Looking at the CIE diagram reveals that each of these colour pairs has one colour on one side of the white-point and one on the other. This allows us to formalise and generalise the concept. We define colours to be complementary if they have the same excitation purity and lie at opposite ends of a line segment which passes throught the white point. The mixing property of colours on the CIE diagram means that complementary colours can be mixed to produce white. This provides an alternative way of defining complementary colours.

The set of colours which is is possible to produce for a graphics device is called the *colour gamut* of that device. One important use for the chromaticity diagram is

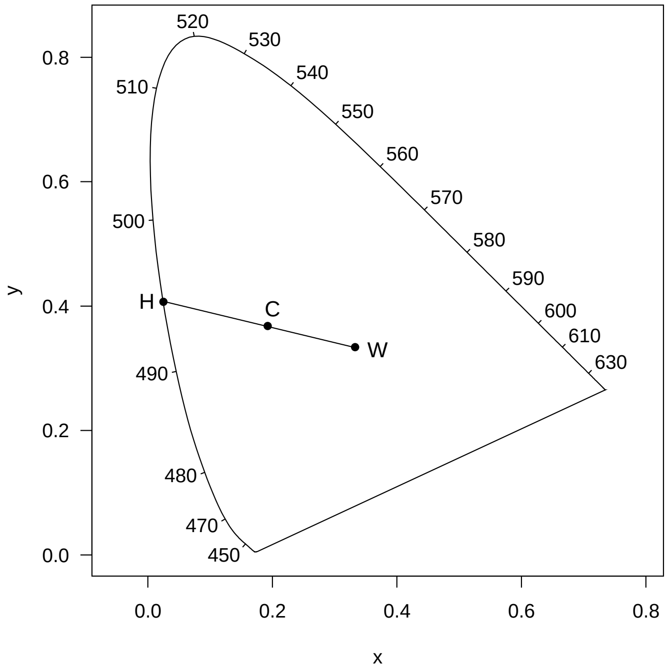

Figure 6.11: Saturation and hue on the CIE diagram.

providing a description of the gamut for devices such as colour computer displays and printers. Figure 6.12 shows the colur gamut for a typical computer display. The vertices of the triangle in the display are the chromaticities of the primary colours which drive the display.

Note that there are large areas of colour which cannot be produced on this display. In particular, it seems that there are many greens and blues which cannot be produced on-screen. Although there are colours which cannot be shown on this kind of display, the situation is not as bad as figure 6.12 would suggest. The chromaticity diagram is not *perceptually uniform* and the the blues and greens lying outside the triangle are difficult to distinguish from those on its boundary. In fact, the main colour deficit for this display occurs in the region close to the purple boundary.

A final use of the chromaticity diagram which we will mention is to describe the effect of “colourblindness” on the perception of colour. Individuals suffering from colourblindness are unable to distinguish between colours which appear quite different to someone with normal vision. The most common forms of colourblindness affect the ability to distinguish between red and green. About 5% of the population (almost all male) suffer from this kind of problem. The chromaticity diagram makes it possible for to give a precise description of the effects of colourblindness by showing the sets of colours which are confused by individuals with various forms of colourblindness.

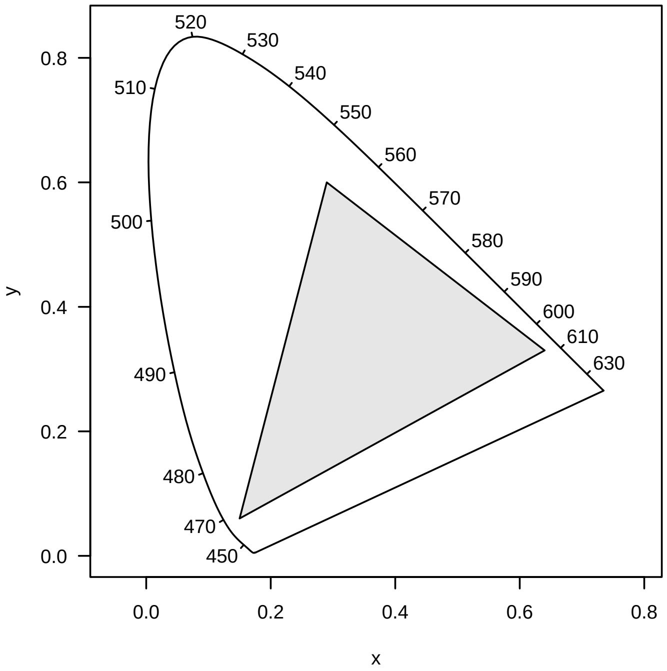

Figure 6.12: The colour gamut for a typical colour display.

6.9 Perceptual Uniformity

A major problem with the (1931) CIE chromaticity diagram is that it is not *perceptually uniform*. This means that distances between colours in diagram do not correspond to how different we perceive the colours to be. Figure 6.13 shows line segments whose lengths are proportional to the *least noticable differences* between colours in the CIE diagram. The line lengths vary greatly in different regions of the diagram, indicating that there is a great deal of distortion.

Perceptual uniformity is important when choosing colours for applications such as graph drawing. Because of this, a number of attempts were made to modify the chromaticity diagram to make it more uniform. In 1976, the CIE introduced the *Uniform Chromaticity Scale Diagram* (UCS Diagram), commonly known as the u', v' diagram. It is obtained by plotting u' against v' where:

$$u' = 4X/(X+15Y+3Z) = 4x/(-2x+12y+3)$$

$$v' = 9Y/(X+15Y+3Z) = 9y/(-2x+12y+3)$$

To obtain x and y from u' and v' the following equations can be used:

$$x = 9u'/(6u'-16v'+12)$$

$$y = 4v'/(6u'-16v'+12)$$

Despite the fact that it results from a non-linear transformation of the original chromaticity diagram, the UCS diagram retains the most useful property of the original

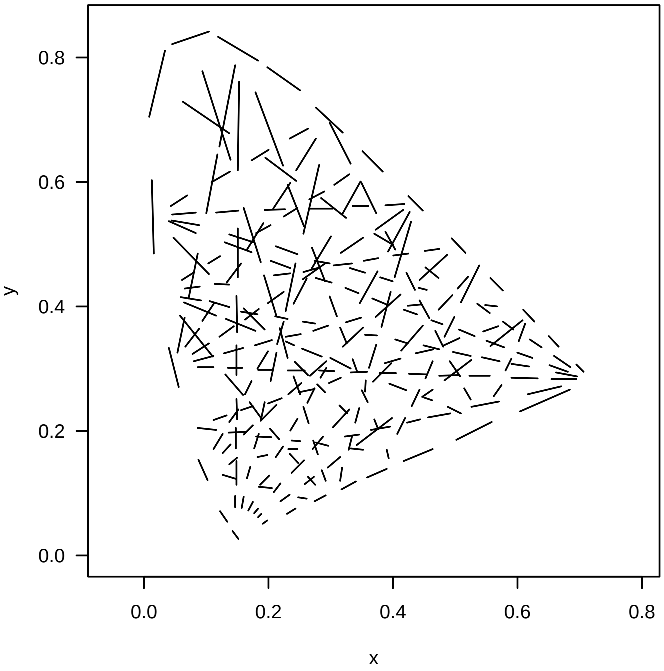

Figure 6.13: Perceptual non-uniformity of the CIE chromaticity diagram.

diagram – the chromaticity of a mixture of colours lies on the straight-line joining the chromaticities of the two colours. The UCS diagram with the spectral locus superimposed is shown in figure 6.14. The transformation compresses the green region of the diagram and raises its red region relative to others.

Figure 6.15 shows that the UCS diagram is much more perceptually uniform than the original chromaticity diagram, but it is still not completely uniform. The diagram represents a compromise between preserving the useful properties of the orginal x, y diagram and perceptual uniformity.

Both the x, y and u', v' diagrams provide a description of hue and saturation, but not brightness. For the original diagram, the brightness information is available in the tristimulus value Y, so that a complete colour description is given by the triple (Y,x,y). However Y does not give a perceptually uniform coding of brightness and the CIE has recommended a transformation which makes the brightness scale more uniform. The transformed value is the CIE 1976 lightness L^*, defined by

$$L^* = \begin{cases} 116(Y/Y_n)^{1/3} - 16 & \text{for } Y/Y_n > 0.008856, \\ 903.3(Y/Y_n) & \text{otherwise,} \end{cases}$$

where Y_n is the Y tristimulus value for a *reference white*. Roughly speaking, this is the appearance of white under ambient lighting conditions. The lightness scale is chosen so that the lightness of the reference white is 100.

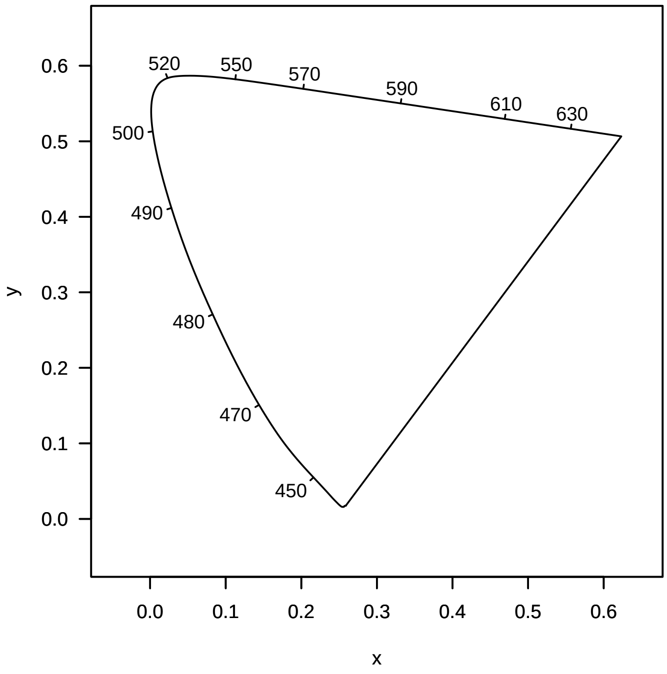

Figure 6.14: The spectral locus in UCS chromaticity diagram.

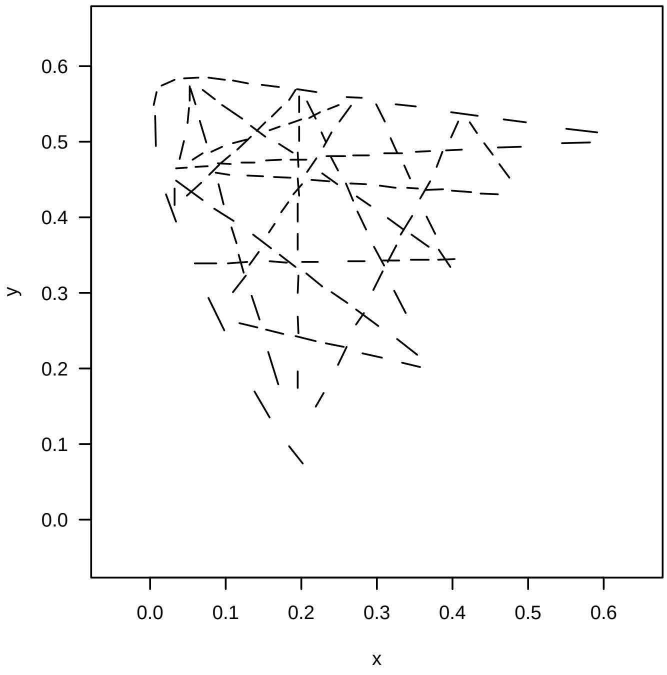

Figure 6.15: Perceptual non-uniformity of the CIE UCS diagram.

6.10 CIELUV – A Uniform Colour Space

The u', v' diagram provides uniform chromaticities, but it is moe useful to have a uniform desciption of *colour*. This is possible if lightness is incorporated in the right way. The *CIE 1976* (L^*, u^*, v^*) *colour space* (or *CIELUV colour space*) is and example of doing this. It is produces by plotting at right angles the quantities

$$\begin{aligned}
L^* &= \begin{cases} 116(Y/Y_n)^{1/3} - 16 & \text{for } Y/Y_n > 0.008856, \\ 903.3(Y/Y_n) & \text{otherwise,} \end{cases} \\
u^* &= 13L^*(u' - u'_n) \\
v^* &= 13L^*(v' - v'_n)
\end{aligned}$$

where u'_n and v'_n are the uniform chromaticities for the reference white.

Although CIELUV provides a uniform space in which to choose colours, the u^*, v^* axes are not intuitive ones. It is more useful to change to polar coordinates in the u^*, v^* plane. The values

$$\begin{aligned}
h_{uv} &= \arctan v^*/u^* \\
C^*_{uv} &= (u^{*2} + v^{*2})^{1/2}
\end{aligned}$$

are know as the CIE 1976 *u*, *v hue angle* and *chroma*. Together with L^* they provide intuitive variables with which to choose colours.

As an example of using CIELUV, consider the problem of choosing colours to fill a number of areas in such a way that the different colours are easily distinguished and have equal visual impact. One way to proceed is to choose the colours to be at equally spaced hue angles and with equal lightness and chroma. Not every combination of L^*, u^* and v^* correspond to a colour, but it is possible when $C^*_{uv} = 59$ and $L^* = 75$. Less intense colours can be obtained by choosing a lower C^*_{uv} value and darker values by choosing a lower value of L^*.

6.11 General Considerations When Using of Colour

6.11.1 Light and Dark Contrast

In making use of colour in graphical (and text displays) it is important to consider the light-dark contrast between the colours being used. The most fundamental part of the vision system works purely in terms of the apparent brightness of the elements of the scene or image being examined. It is this monochrome part of the vision system which extracts basic information about shape and position. The absence of light/dark contrast makes it very difficult to see anything in an image.

A useful check of whether the structure present in an image is easy to see can be made by translating the image to grayscale. After the translation, the structure of the image should remain clear. In the printing industry the slogan "*get it right in black-and-white*" is often quoted.

Some Simple Rules

There are a number of very simple principles which, when used, help to produce clear graphic images.

1. Draw with dark colours on light backgrounds.
2. Draw with light colours on dark backgrounds.
3. Use colours of contrasting lightness to emphasise boundaries between colours of similar or equal lightness.

6.11.2 Effects Associated With Particular Colours

Blue

There are relatively few blue cone cells in the fovea. It is thus harder to see fine detail when it is presented in blue. Because of this, blue is best used as a background colour.

The eye's blue cone cells have the slowest response to light changes. This makes it hard to see blue objects in motion.

Red and Blue

Red and Blue lie at opposite ends of the visible spectrum. They are refracted different amounts by the lens of the eye. A consequence of these differing levels of refraction is that most people cannot bring Red and Blue into focus simultaneously. The majority of people see red as being in font of blue, a smaller group see blue in front of red and a few special individuals see them as being at the same distance.

Purple

Purple is a mixture of red and blue. The red and blue components of purple undergo different amounts of refraction when they pass through the lens of the eye. This means that the lens must be adjusted differently to focus the components. This makes purple a hard colour to focus on.

6.11.3 Intense Colour and Cone Bleaching

When the eye looks at areas of intense colour for extended periods, the sensitivity of the of cone cells responding to the colour is temporarily lost. To restore the sensitivity, the eye must be rested.

This means that it is best to avoid large areas of intense colour in presentation graphics. If intense colour is used, the eyes quickly become fatigued. To quote Albert Munsell, one of the founding fathers of modern colour theory:

> *The circus wheel and poster, although they yell successfully for momentary attention, soon become so painful to the vision that we turn from them.*

Note that recommendation is frequently violated in advertising. This is because the advertisers are desperately striving for your attention. Having gained your attention they need only hold it for a split second to get their message across.

6.12 Specific Reasons for Using Colour

There are a number of specific ways in which colour is used in graphical presentation and visualisation. Some of these are clearly useful and some are problematic.

6.12.1 Encoding Numerical Values

It is tempting to try to use colour to represent values on a numerical scale – e.g. red for high values and blue for low ones. In practice such an encoding of numerical values does not work well.

Research by visualisation researhers at IBM has shown that luminance is the only effective colour-based way of encoding numercial information – hue and saturation are not useful. Even luminance has to be considered of limited value. The Cleveland-McGill experiments place it towards the bottom of the perceptual scale.

The use of a rainbow encoding is common, particularly amoung practitioners of “scientific visualisation.” It is however quite misleading. After all, hues form a circle, not a straight line.

6.12.2 Encoding Ordered Categorical Values

Color can be used effectively to encode a set of ordered categories, provided there are not too many categories. To preserve the visual ordering it is important that only a limited range of colour be used. This is because of the wrap-around effect of colour noted above.

6.12.3 Encoding General Categorical Values

Colour is a very effective way of differentiating between different categories. Provided that no more than six colours are used, it it possible to decide which information corresponds to which category with just a single glance.

6.12.4 Differentiating or Grouping Graphical Elements

As in the categorical case, colour can provide a very effected way of grouping or differentiating the elements of a graphical display. It is much more effect that using shape to indicate relatedness.

6.12.5 Making Dull Plots Look Interesting

If the graphs are really that dull, no amount of window-dressing is going to help. Its better not to display them at all.