

Lecture 7: Colour

The physical description of colour

Colour vision is a very complicated biological and psychological phenomenon. It can be described in many different ways, including by physics, by subjective observation, or by the tri-stimulus representation commonly used in computers. As we noted previously, light has a wavelength (λ), which can be measured accurately, and an energy or intensity. Lasers produce light of one particular wavelength, but generally visible light is a distribution of energy into a band of wavelengths. Figure 1(a) shows a distribution of light energy which we would perceive as red in colour, and figure 1(b) shows the the distribution of energy found in sunlight, which we perceive as light yellow. The wavelength range of the visible spectrum is approximately from 400 to 700 nano-meters.




Figure 1: Energy Distributions for light

In the eye we have just three distinct 'cone' cells for detecting light energy. These respond to a band of wavelength centred around red (600), green (560) and blue (440); respectively. The bands overlap, so, for example green light excites all three types. Typical responses of the cone cells are shown in figure 2 The consequence of this is that each type of cell may be excited similarly from very different energy distributions. Any distribution of wavelengths will be perceived by us as a single colour, but two entirely different distributions of intensities could be perceived as the same colour.

Subjective observation of colour

Two colours are said to be equivalent when an observer says that they are equivalent. Selecting three pure light sources (R, G, B) and mixing them together while varying their respective intensities, one may be able to create a large number of colours. Each colour satisfies the following linear combination: $X = r + g + b$ where r, g, b are the intensities of the red, green and blue light sources in the mix. One way to approach colour definition is through a matching experiment. For an unknown colour we can compare it to a mixture X , and adjust r, g and b until the mixture matches the unknown colour. As we shall see later, not all colours can be matched in this way. However, by adding one of the pure colours to an unknown, unmatchable colour, we can make a match. This is in effect subtracting a colour from the mix.

$$X + r = g + b$$

$$\text{or } X + g = r + b$$

$$\text{or } X + b = r + g$$

The representation of colours as a mixture of three components is called the tri-stimulus representation, and is




Figure 2: Typical responses of individual cone cells

very commonly used in monitors and other active colour devices. The pure colours used are red green and blue, and are referred to as the additive primary colours.

The Standard Additive Linear System (CIE)

Matching allows us to determine an rgb representation of any colour, even though the scales include negative numbers. The scales used depend on the wavelengths of the sources for matching and are not standard. In order to create a manageable standard system, a committee of scientists devised the CIE Chromaticity Diagram, in which colours are specified by a linear combination of three light sources normalised to the positive range [0..1] to avoid any negative values. The three principle colours would normally require a three dimensional diagram, but by using further normalisation a two dimensional diagram is produced. Suppose a colour is defined, as above, by: $C = r + g + b$. If we divide each component of this equation with $(r + g + b)$ and get a normalised set of coordinates:

$$\begin{aligned}
 x &= r / (r + g + b) && \text{the normalised red component,} \\
 y &= g / (r + g + b) && \text{the normalised green component,} \\
 z &= b / (r + g + b) && \text{the normalised blue component,}
 \end{aligned}$$

and $x + y + z = 1$

So, knowing two of the three normalised colour components (x, y) the third one can be determined, for example $z = 1 - x - y$.

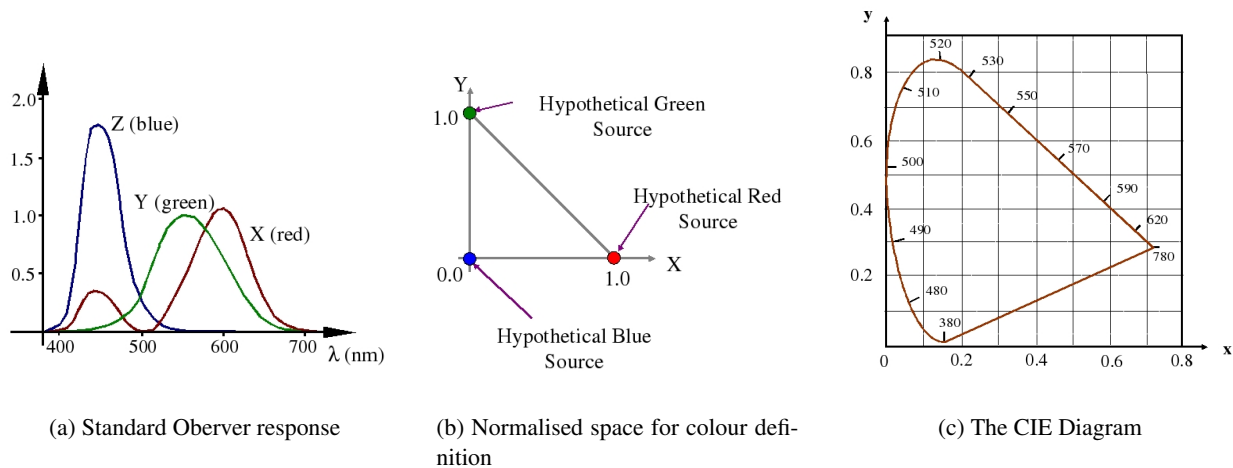


Figure 3: Defining the CIE Diagram

The process by which the CIE diagram was defined is illustrated in in figure 3. First a standard observer response was defined. This takes into account not only the typical responses of the three types of cone cell, shown in figure 2, but also the density of the cells in the foveal region of the eye. This standard response is shown in figure 3(a) and represents the typical overall cone cell responses to the light arriving in a small solid angel at the eye. Any wavelength in figure 3(a) can now be associated with a tri-stimulus value defined by the relative magnitudes of the responses of the three types of cell. For example at the wavelength 580 we can read off the graph that blue is zero and red and green are approximately equal at 0.8. The three values are normlised so that they sum to 1 and plotted in a normalised cartesian space which is illustrated by figure 3(b). The resulting wavelengths plotted in this space are shown on the CIE diagram of figure 3(c). All visible colours can be defined in this way, though only their normalised values are shown on colour prints of the diagram.

The pure visible colours are distributed on the edge of the horseshoe shaped curve with their physical wavelength running from 380 to 780 nm. At top and bottom ends of the visible spectrum, where the cell responses fall off, the wavelengths converge very fast towards a vanishing point. For example the wavelength 700 nm is virtually coincident with the 780 nm point in figure 3(c). We can see that just below 500 nm the red component is zero and the blue and green are approximately equal which would correspond to the colour cyan. The straight line from 780 to 540 on the right hand side represents the condition that $x + y = 1$ which means that the blue component is zero. The straight line at the bottom, joining 380 to 780 is a combination of red and blue, which has no equivalent single wavelength colour. It is called the magenta line. All visible colours are

contained within the horseshoe, which must be convex, since all colours are a mixture of pure wavelengths. The usefulness of this diagram is that it expresses a linear additive system and the addition of colour components can be done graphically.

Graphical Manipulation of Colour Parameters

Some definitions are shown on the CIE diagram of figure 4. The white point, $x = y = z = 1/3$, has equal strength of all three basic components. Pure colours, around the edge of the horseshoe, are called pure or fully saturated colours. If we take a pure colour, say 488 nm on the diagram, the line connecting this point and the white point contains all the possible colours which can be achieved by adding white to this pure colour. If we extend this line and on the other side it hits the curve at another pure colour point (598 nm in this case), which is called the complement because mixing the two makes white.

Saturation, or the degree of purity, of a given colour point can also be determined graphically. The white point is completely unsaturated (has no colour). The ratio of distances between the colour and the white point and the pure colour and the white point gives us saturation:

$$saturation = \frac{|W - U|}{|W - P|}$$

The saturation is always taken as 1 on the horseshoe (fully saturated) and zero at the white point. If we take any colour on the diagram, for example [0.2, 0.3, 0.5] shown in figure 4, we can measure this ratio, which in this case is approximately equal to 0.4. This tells us the white content of the point. We can also determine a pure colour by considering the line through the point and the white point. From the slope of the line we can calculate that it cuts the y axis at [0, 0.25, 0.75]. This point is not in the visible region of the CIE diagram. The nearest pure colour will be seen to be around [0.04, .26, .736]. The complement of this colour can be found by determining the colour which when added to it will produce white. This will be on the z=0 line, so the blue component is zero, and so we have that [0.04, .26, 0.736] + [r,g,0] = [k,k,k] from the blue component we can find that k = .736, whence 0.04 + r = 0.736, so r = 0.696 and g = 0.476. If we normalise these points we get [0.6, 0.4, 0.0] as shown on the diagram. Notice that the normalisation on the diagram gives a weighting to the colour coordinates. Since on the diagram the line joining a colour and its complement always goes through the white point, it is always possible to find a linear combination of the two that produces white. It does not mean that adding the co-ordinates and normalising produces white.





Figure 4: Calculations on the CIE Diagram


Subtractive Colour representation

Paints and printing ink absorb colour reflecting only part of the incident light. For this reason they are called subtractive primaries. The three principal colours of paints are magenta (purple), cyan (pale blue) and yellow. These three can be created by combining the additive colour primaries. Magenta is red and blue, cyan is blue and green, and yellow is red and green. Black is created by a blend of the three subtractive primaries, but is usually also used in colour printing for convenience. It reflects none of the incident light. White paint (or paper) reflects all the incident light and so does not change its colour. Black and white paints can be blended to produce different shades of grey. With grey and red paint you may make brown, and a milliard other colours which have no equivalents in the preceding systems. Paints are much richer than colours produced by red/green/blue lights.

Figure 5 indicates how the different, pure basic colours are generated by the additive (monitor) and subtractive (printer) colour systems. Each circle represents one basic additive (R,G,B), or subtractive (M,C,Y) colour. The other regions indicate pure



Additive Primaries



Subtractive Primaries

Figure 5: The Primary Colours

colours generated by two basic colour addition or subtraction, respectively. The centre colour (three addition or subtraction) in the additive system is white, while it is black in the subtractive system (all colours are absorbed). The two systems show additive or subtractive complementary colours in the same region; i.e., red + cyan light makes white while red + cyan paint makes black.

Practical Colours

Approximately 128 different hues may be distinguished by the human eye. For each hue around 20 to 30 different saturation values may be seen as different colours. The eye is capable of distinguishing between 60 and 100 different brightness levels. If we multiply these three numbers, we get approximately 350,000 different colours. 24 bits provides us over 1.6 million colour shades, however, the eye is very sensitive to differences and under some circumstances small differences can be detected by it.

When a transformation is necessary from the CIE standard colour chart to the colour produced on a specific display device, we must express the primary colours of the display device by x and y quantities. Good quality monitors will be calibrated for the CIE colour chart. For example, a colour CRT monitor may have the following primary colour sources:

	x	y	z
Red	0.628	0.346	0.026
Green	0.268	0.588	0.144
Blue	0.150	0.07	0.780

These three points define a triangle on the CIE chromaticity diagram, as shown in figure 6. Only the points inside this triangle may be reproduced by the display device.

Colour Spaces and Transformation of Colours

There are many different ways of representing colour, and often they can be expressed as linear combinations of each other. For example, the transformation from computer colour intensities R,G,B (expressed in the range from 0.0 to 1.0) may be done by the matrix:

$$[x, y, z] = \begin{bmatrix} 0.628 & 0.268 & 0.150 \\ 0.346 & 0.588 & 0.07 \\ 0.026 & 0.144 & 0.780 \end{bmatrix} \begin{bmatrix} R \\ G \\ B \end{bmatrix} \quad (1)$$

The pure red on the monitor has representation of [1.0, 0, 0] which then produces the correct CIE values of [0.628, 0.346, 0.026] for the monitor. In order to find out how to represent a known CIE colour on the display, we have to invert the matrix and then right multiply it with the vector [x,y,z].

One useful colour space is called the HSV (Hue, Saturation and Value) space. It is of interest since it is thought to correspond more closely to our perception of colour than the RGB space we have been discussing. Like the RGB system, a colour is represented by three values which can be calculated from the r g b values as follows:

$$V = \max(r, g, b)$$

$$S = (\max(r, g, b) - \min(r, g, b)) / \max(r, g, b)$$

Hue (which is an angle between 0 and 360°) is best described procedurally:

if (r = g = b) Hue is undefined, the colour is black, white or grey.

if (r > b) and (g > b) Hue = 120 * (g - b) / ((r - b) + (g - b))

if (g > r) and (b > r) Hue = 120 + 120 * (b - r) / ((g - r) + (b - r))

if (r > g) and (b > g) Hue = 240 + 120 * (r - g) / ((r - g) + (b - g))

There are other, very similar, perceptual colour space called HSL (lightness) and HSI (intensity) which have slightly different definitions for saturation and intensity. Note that, every colour in the RGB system is made




Figure 6: Display Colours

up of a white component whose magnitude is the minimum out of r , g and b , and a pure colour defined by the ratio of the two non-zero colours when the white component has been subtracted (Figure 7). Note that the pure colour is not a coherent wavelength as in the CIE diagram. The saturation is the proportion of pure colour. As in the RGB system colours may be adjusted by changing any of the three values. However, with the HSI system the three variables each correspond to properties of colour we can easily understand, and thus they are easier to use to adjust colour palettes and balance colour images.

Alpha Channels

Colour representations in computer systems sometimes use four components to represent each colour. The first three are R, G and B and the fourth or alpha value is simply an attenuation of the intensity. This does not add anything further to the representation, but allows greater flexibility in representing colours and can avoid truncation errors at low intensities. It can also be used for providing special features such as masking certain parts of an image.




Figure 7: A tri-Stimulus Colour