Lecture 8: Colour Generation in Computer Graphics

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.

The physical description of colour

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. Diagram 10.1 shows a distribution of light which we would perceive as red in colour along with 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.

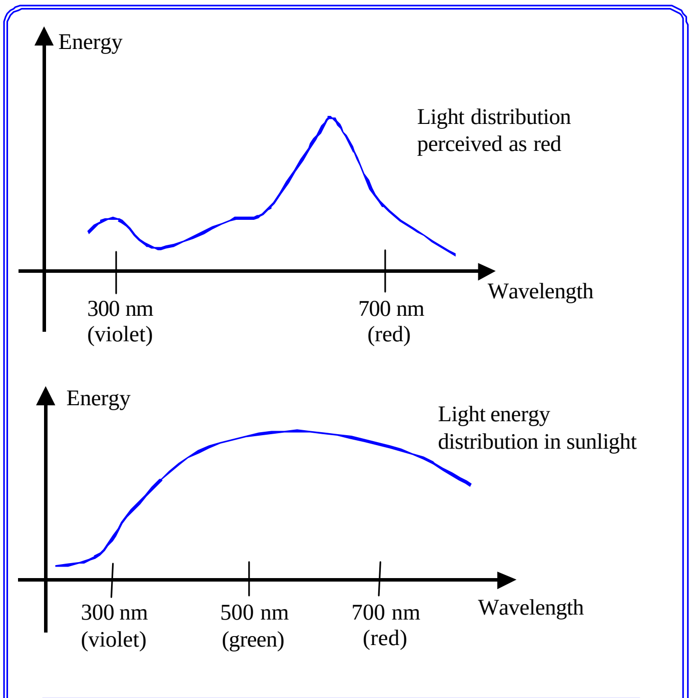

Diagram 10.1: Colour as an energy distribution

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. 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:

$$X = r*R + g*G + b*B$$

where r, g, b are intensities of the red, green and blue light sources. 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.

$$X + r*R = g*G + b*B$$

or $$X + g*G = r*R + b*B$$

or $$X + b*B = r*R + g*G$$

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

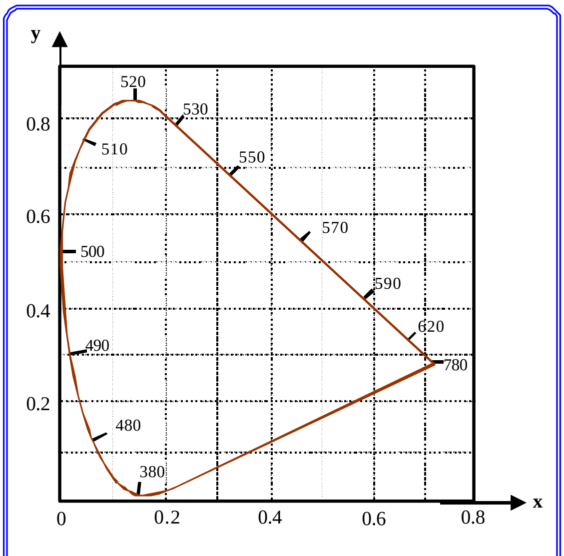

Diagram 10.2: The CIE diagram

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 use 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 by:

COLOUR = r + g + b

where r, g, and b. If we divide both sides of this equation with (r+g+b) and get:

$x = r/(r+g+b)$ the normalised red component,
$y = g/(r+g+b)$ the normalised green component
$z = b/(r+g+b)$ the normalised blue component

and $x + y + z = 1$

So, knowing two of the three normalised colour components (x,y) the third one can be determined, i.e. z=1-x-y. All normalised colour values can now be shown on a two-dimensional figure of diagram 10.2

The horseshoe shaped curve represents the visible pure colours with their approximate physical wavelength running from 380 to 780 nm. 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 part of the diagonal line on the right represents the condition that x+y=1 which means that the blue component is zero. The straight line at the bottom is a combination of red and blue, which has no equivalent single wavelength colour. It is called the magenta line. The usefulness of this diagram is that it is based on a linear additive system and the addition of colour components can be done graphically.

Graphical Manipulation of Colour Parameters

Diagram 10.3 shows a CIE diagram. 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 595 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 (485 nm in this case), which is called the complement because mixing the two makes white.

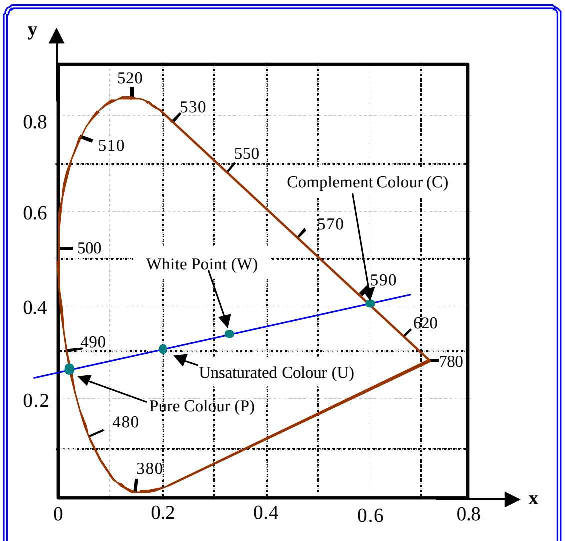

Diagram 10.3: Colour computations

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 = |**W-U**| / |**W-P**|

The saturation is always taken as 1 on the horseshoe (fully saturated) and zero at the white point. If we consider a colour on the diagram at point [0.2, 0.3, 0.5] we can see that this ratio 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, and so we have that [0.04, .26, 0.736] + [r,g,0] = [k,k,k] Solving for r and g we have first 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.

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 principle 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.

Diagram 10.4 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 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.

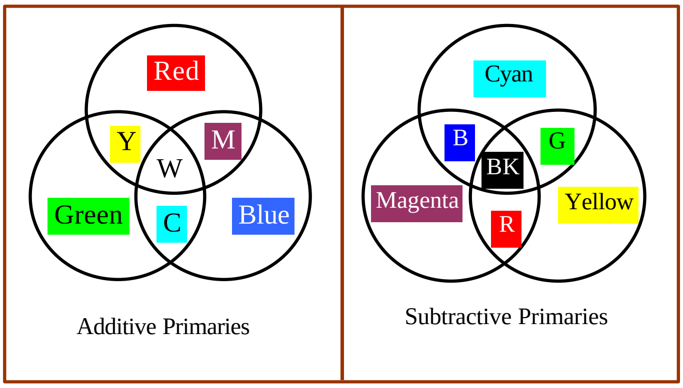

Diagram 10.4 Colour Mixing

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.

Practical Colour Displays

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 on diagram 10.5. Only the points inside this triangle may be reproduced by the display device.

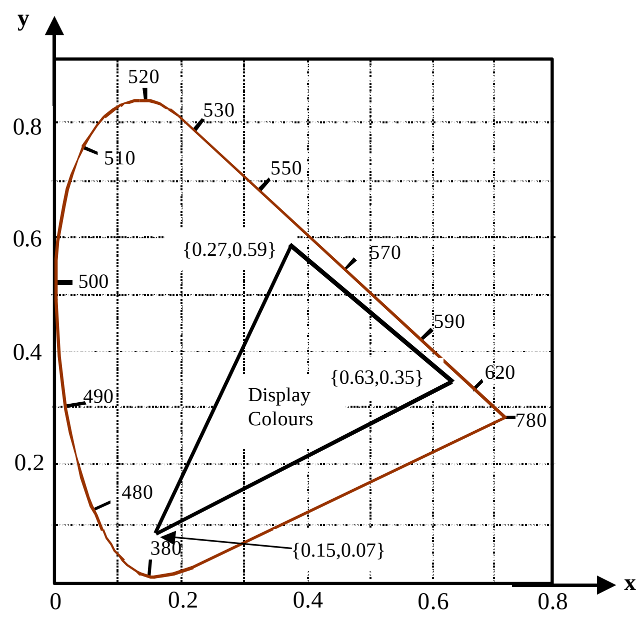

Diagram 10.5: Display Colours

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:

$$(x, y, z) = \begin{pmatrix} 0.628 & 0.268 & 0.15 \\ 0.346 & 0.588 & 0.07 \\ 0.026 & 0.144 & 0.78 \end{pmatrix} \begin{pmatrix} R \\ G \\ B \end{pmatrix}$$

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 HSI (Hue, Saturation and Intensity) 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:

$I = (r + g + b)/3$

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

Hue (which is an angle between 0 and 360^{o}) 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))

Note that, every colour in the RGB system is made 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. (Diagram 10.6) 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.

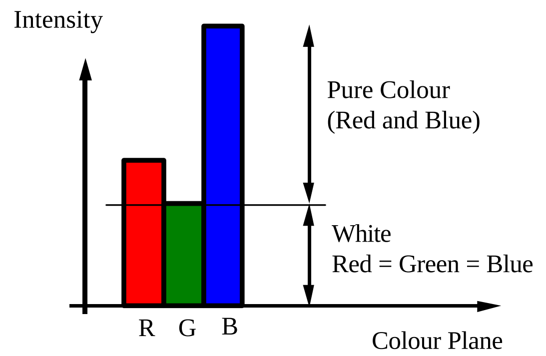

Diagram 10.6
A tri-stimulus colour represented as a pure colour and white

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.