Chapter 7

Representing Colour

7.1 Theory

We now turn to look at colour and its implementation. So far we have talked about red, green and blue components (as in colour TVs), but the real story is more complicated.

It is not commonly realised that there are many ways to represent colour. The RGB (red green blue) model is widely understood due to its use in (a) the eye, and (b) televisions, but there are at least two dozen other models in common use, largely depending on the application area. Many of these are industry standards.

This variety of models and standards is not because people can't agree on the best one to use. It's because there is no best: each device has very different needs. Thus, a colour printer (based on inks absorbing light) needs a different model from a screen (based on phosphor emitting light).

However, there is a more fundamental problem that we can't standardise, namely human perception. The difficulty is that different humans respond to the same colour in different ways. This can be physiologically based (someone can be red-green colour blind; colours you have seen recently can affect the way you view another colour), but also culturally based. For example, in the West, white is associated with marriage, black with funerals. In other cultures the opposite is true. In fact purple and black were the colours of mourning in the UK until around 50 years ago; now it is just black. We associate red with danger, but this is only a convention (though probably derived from the colour of blood or fire). Similarly, green is for "go" but there is no obvious reason for this.

Language is also important. Before the 16th century there was no word for "pink" or "orange," both were considered variants on red. In *Macbeth* the word "gilded" is used for the application of blood: gold, too, was thought to be a shade of red. Today we wouldn't think of relating gold and red, but this is merely an effect of language. Black is the absence of colour but is still called a colour.

The seven colour rainbow spectrum was invented by Newton, who thought seven was a nice number, to match with the seven intervals in a musical octave. Looking at a spectrum it is difficult to count seven distinct colours; it is of course a continuum.

Colour is thus a mix of psychology (the mind), physiology (the eye), language and culture.

We shall start by addressing the physiology of the eye.

7.1.1 The Eye

It is well established that the eye uses a combination of *rods* and *cones*. Near the centre of the eye there are many cones, which detect colour. Generally spread about are the rods, which detect brightness (shades of grey).

Cones only work well in bright light. In low light only the rods are contributing to what you see: this is why night time scenes are grey, though the brain does try to fill in with colours it thinks you ought to be seeing.

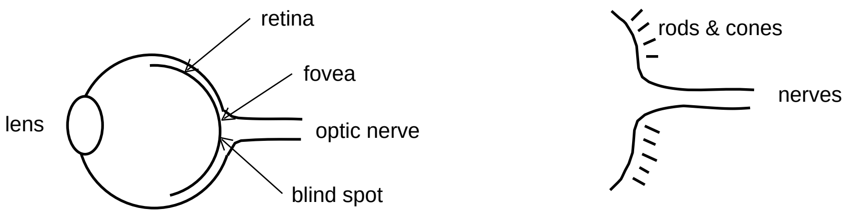

The optic nerve enters the eye from behind creating a *blind spot* where there is no visual sensitivity. The nerves then continue into the eye over the *top* of the rods and cones. Light has to pass through the nerve layer, and its associated blood vessels, to reach the rods and cones. This is a strange arrangement, but the brain edits what we think we are seeing, eliding away both the effect of the nerves and the blind spot. Even at this level the brain has an effect on the optics of the eye.

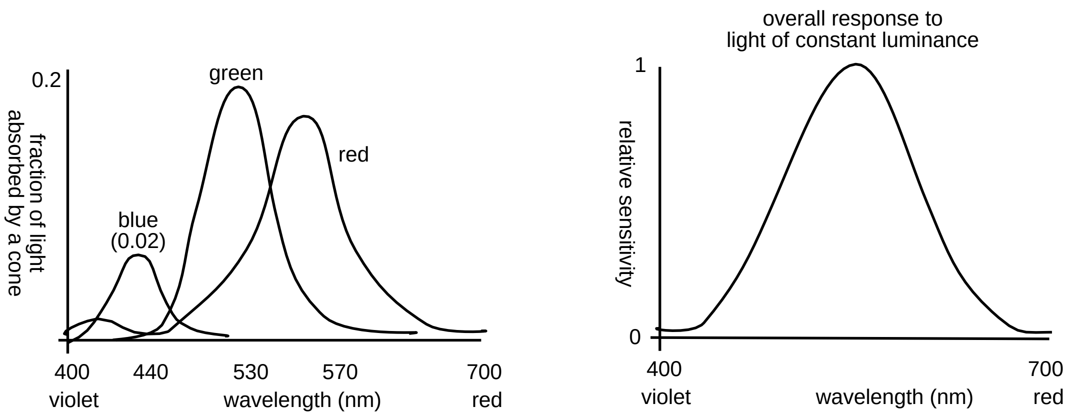

Cones come in three kinds, with peak sensitivities in different parts of the spectrum. These are conveniently called the red, green and blue cones. However, the responses are quite broad, with considerable overlap. Green cones have the strongest response, red is close but slightly weaker, while blue is 20 times lower, i.e., the eye's response to blue is much lower than to red and green. The blue peaks at about 440nm, green about 545nm, and red about 580nm. Interestingly, the red and green peaks are within the yellow part of the spectrum. Furthermore, red cones have a secondary response in the blue part of the spectrum. The eye's limits are about 400nm for violet and 700nm for red.

This use of three receptors to explain low-level colour is called the *trichromatic theory* or the *tri-stimulus model*. Basically this states that all perceivable colours can be represented as RGB stimuli to the eye. (This is almost true but not quite, as we will see shortly.) Thus, when we see, say, yellow light, what is really happening is that the red and green cones are being excited to certain relative levels. If we were to take red light and green light of the right intensities and mix them, the relative excitement of the cones would be the same, and we would perceive yellow. This, of course, is the principle of TVs. (The technical term for two colours which are visually the same but generated by different mixtures of wavelengths is 'metamers').

Overall, the eye's response to light of constant luminance has a peak at about 550nm, roughly the colour of sunlight.

Even though we have a completely non-linear response in the eye, we will start with a simple *tri-stimulus* system. This means we shall represent a colour by a mixture of standard red, green and blue sources, mixed in some proportion. These are our *primary colours*.

We get *secondary colours* by mixing the primary ones in equal measure:

$$\text{yellow} = \text{red} + \text{green}; \text{cyan} = \text{blue} + \text{green}; \text{magenta} = \text{red} + \text{blue}.$$

We get white by adding all three primaries in equal proportion:

$$\text{white} = \text{red} + \text{green} + \text{blue}.$$

It is important to understand that the basis of the tri-stimulus model is the human eye, not physics. The spectrum is a continuous, linear concept. If you were to invent something to approximate it, you might choose one light at each end of the visible part - blue and red say - and hope you could linearly interpolate the colours in-between. The reason this does not work fully is exactly because of the need to stimulate eye's cones. Equally, the existence of metamers is a consequence of the overlap of the cones' responses: there is no unique physical stimulus for any given perceived colour. One immediate consequence is that our computer screens depend on the viewing eyes having three types of cone: an alien - or even a bee - will not see what we see.

As an interesting digression, you might like to know that a bee has three receptors too, but one is in the ultraviolet and bees are less sensitive to red than we are. Many birds can also see in the ultraviolet but in their case they have four receptors and so have a wider viewable range than we do. So do some fish and turtles. Sometimes male-female pairs of birds look similar to us but are very different in the UV, so they can find the opposite sex more easily: the blue in a blue tit's head is different in the UV, according to the sex, but is very similar to us. Some white feathers reflect strongly in the UV, others do not, so they will look different to other birds. Blue tits, zebra finches and even starlings use UV in mate selection. Maybe that UV "black light", once popular at discos, had its uses after all.

7.1.2 Tri-Stimulus Experiments

We can test the usefulness of this model. We can show a large selection of colours to a sample of (normal vision) people under standardised conditions. We then get them to match each colour by mixing pure red, green and blue lights. If they can do this, then we *define* the proportions to be the tri-stimulus values for the colours.

But there is a snag. It's not possible to get every perceived colour by mixing red, green and blue lights. The strange thing is that we *can* get every colour if we allow *negative* proportions of light.

To explain this, we need to explain the experiments in more detail. The set-up presents the subject with two panels side-by-side. One has the colour to be matched, the other is a neutral panel illuminated by the three sources. The subject can control the intensities of the three sources. Given a match, the three intensities define the tri-stimulus values for the colour.

If the subject cannot get a match, we add a colour *bias* to the original sample by adding, say, a little red to it. If the subject can now achieve a match, then we have a tri-stimulus value for Colour+red. So if we subtract the amount of red from the tri-stimulus value, we will have the tri-stimulus value for the original colour. This will produce a negative number.

To see what is happening, recall the responses of the cones are roughly

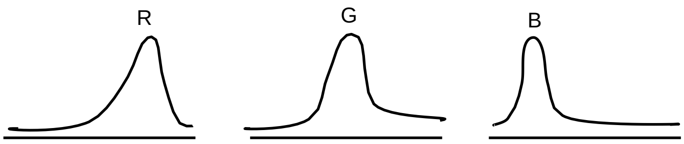

Suppose we wish to match a colour that should have this response in the eye

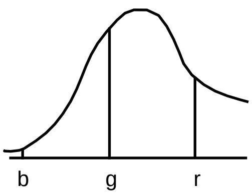

Mixing red and green we get too much appearance of blue due to the tails of the red and green responses.

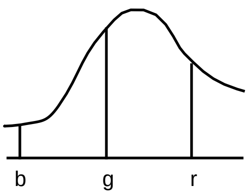

So adding in a little blue to the test colour we get

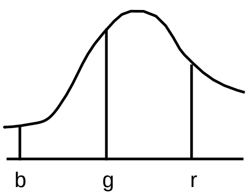

which matches what we get from the tri-stimulus. Thus we would deem this colour to have a negative blue value.

The reason this occurs is not mysterious. While we can get red, green or blue lights very close to a single frequency, the cones respond to a wide range of overlapping frequencies. Cones are not independent sensors. Our subject, trying to match a colour by increasing one primary, will also unintentionally stimulate the cone for another primary - or even all three. To offset this requires "negative light" to be added to that second primary.

Thus a colour TV or computer display *cannot* display every colour that the eye can see! However, the loss is not great for most purposes, as they *can* display a huge number of colours.

It seems that the eye can distinguish about 1/3 of a million colours if we judge them side by side; and the sensitivity of the eye is dependent on wavelength. For colours of equal brightness, the eye can distinguish colours when they are separated by

- about 2nm near blue (480nm) and yellow (580nm)
- generally 4-10nm elsewhere.

Thus the eye is fairly unresponsive to blue but can distinguish many shades if they are bright enough.

7.2 The CIE Chromaticity Diagram

This was the first attempt to standardise colour. It was based on the above theory, with a small number of practical tri-stimulus experiments to determine average human responses. It was devised by the Commission Internationale de L'Eclairage (International Lighting Committee) in 1931, though it has been updated several times since, notably in 1976. It is also known as the XYZ model.

It addresses the problem of having negative weights for colours: we would prefer only to have positive values. It solves it at the cost of having the three primary colours not physically realisable!

It starts with three *ideal* primary colours **X**, **Y** and **Z**, which are offset from red, green and blue in such a way we need only positive sums of **X**, **Y** and **Z** to produce all visible colours.

Suppose our colour **C** has weights X, Y, and Z:

$$\mathbf{C} = X\mathbf{X} + Y\mathbf{Y} + Z\mathbf{Z}.$$

We call (X, Y, Z) the *chromaticity coordinates* of the colour. These numbers are determined by experimentation with observers. We further define *chromaticity* values, or the *normalised chromaticity coordinates* of the colour as

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

This is dividing out the total energy of the colour, so the result is intensity independent. Now note that $x + y + z = 1$, so we can specify the colour by just x and y. We have lost intensity information, so we have dropped a dimension. To recover X, Y, and Z we need an extra piece of information, say Y, when

$$X = \frac{x}{y}Y,\ Y = Y,\ Z = \frac{1-x-y}{y}Y.$$

We can plot colours in the (x, y) space:

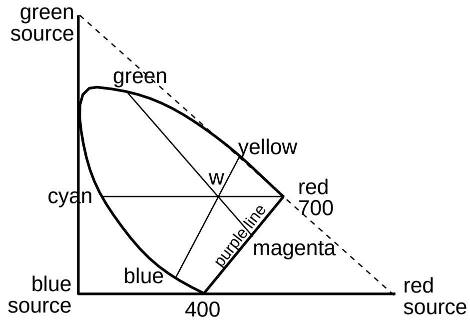

We have

- colours of the spectrum fall on the curved part of the boundary
- white is in the middle, somewhere close to $x = y = z = 1/3$
- the straight edge consists of *non-spectral purples*; these have no single free colour equivalent
- as we move towards w, colours are more desaturated (contain more white): these are pastel colours
- all colours are of unit intensity due to the normalisation, so there are no intensity related colours, such as brown

We can read off the spectral components and saturation of a colour from the diagram.

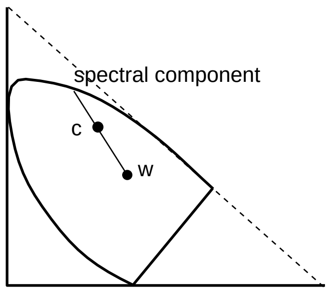

The saturation is the proportional distance to the colour from the white point, and the spectral component (fully saturated) is the intercept with the curved edge. If we completely desaturated any colour, it becomes white.

The position of white in this diagram is somewhat complex. Nominally it is at the (1/3,1/3) point, corresponding to equal parts of all three primaries. People's opinion of what is true white has changed over the years, and has drifted somewhat towards a more bluish white. So the "reference white" has been defined with reference to the black body radiation of 6774 Kelvin. This is close to natural daylight. Other whites have been defined for particular purposes.

7.2.1 RGB on the CIE Diagram

A computer display has three colours of phosphor red, green, and blue. Typical CIE values for them are

	x	y
R	0.628	0.330
G	0.285	0.590
B	0.1507	0.060

(In practice, no phosphor produces a single pure colour as this implies, so this is already a compromise.) This defines a triangle in the CIE diagram, and all the colours that the screen can display lie within this triangle.

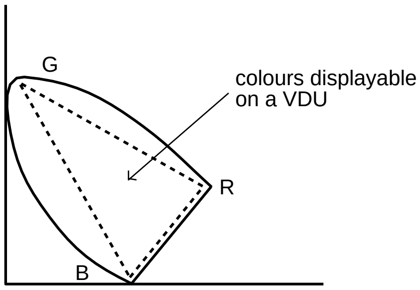

We can never get all the CIE diagram within a triangle created from real phosphors because real phosphors generate real visible light (within the CIE 'horseshoe'), not the hypothetical light of the CIE primaries. So (again) we cannot display all visible colours on a PC screen. In fact the typical triangle of a monitor is much smaller than we have pictured. We call the range of colours a device can display the *colour gamut* of the device. We can use the CIE diagram to discover the colour gamut of, say, a printer.

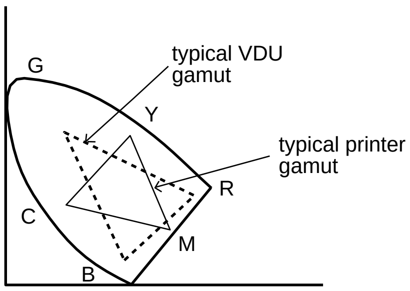

Only those colours in the overlap of the screen gamut and the printer gamut can be displayed on both. If we want to take any screen picture and print it, we have to decide what to do with those screen colours which are outside the printer gamut. This problem of mapping a picture on the screen onto the printer is a difficult one. There are several approaches.

- Perceptual Mapping. Compress the colour space so that colours outside the target gamut are mapped to inside. This alters colours inside the gamut but aims to retain the relationships between colours.
- Relative Colourimetric. Map a colour outside the target gamut to the closest within, leaving colours within the gamut unchanged. Two colours which appear different in the source colour space might be mapped to the same colour in the target gamut. Also known as *clipping*.
- Saturation. A mapping that keeps the original image's colour saturations. Used where bright colours are more important than accurate colours (e.g., business graphics).

A *colour management system* (CMS) tries to manage the mappings between devices while maintaining colour accuracy. Printer drivers usually do this for you and offer you the choice via a pop-up control panel. CIE is often used as an intermediate colour model.

7.3 Colour Models

7.3.1 The RGB Cube

Given the three additive primaries, we can devise a simple representation in which we have three axes, R, G and B. If we normalise the values needed by the output device generating these colours, then we can visualise a unit cube in this space. One corner of the cube corresponds to black (0,0,0) and sits at the origin of the coordinate system. The diagonally opposite corner corresponds to the brightest white (1,1,1) the device can produce. The three corners closest to the black corner are the primaries red (1,0,0), green (0,1,0) and blue (0,0,1). The remaining three corners are the secondaries yellow (1,1,0), cyan (0,1,1) and magenta (1,0,1).

This is a very easy representation to use for tri-stimulus output devices, such as screens.

7.3.2 CIE LAB Colour Space

RGB does not match well to human expectations and is definitely not perceptually linear. One improvement is to separate out the overall “lightness” of the colour L. We then need two more dimensions. Experiments show that the way the eye-brain analyses colours is on an *opponent colour* basis. Red and green are opponents, as are blue and yellow. CIE (1976) defined a colour space using these coordinates, known as CIE LAB. The “A” corresponds to the red-green axis (green is negative, red is positive) and the “B” to the yellow-blue (blue is negative axis, yellow is positive). Both have neutral grey at the zero point. The other axis, L, is shades of grey from black (negative) to white (positive).

There is non-linear transform between CIE XYZ and CIE LAB, which ensures that LAB is perceptually (nearly) uniform.

```
if (Y/Yn)>0.008856)
L= 116.0 x 1/3 x (Y/Yn) - 16.0
else
L= 903.3 x (Y/Yn)

a= 500.0 x 1/3 x ((X/Xn) - (Y/Yn))

b= 500.0 x 1/3 x ((Y/Yn) - (Z/Zn))
```

where Xn, Yn and Zn are the XYZ values of the reference white point.

Which all goes to show why you need a good reference book if you want to manipulate colour spaces!

Opponent colour theory explains visual effects such as after images, which trichromatic theory does not. Both theories are necessary to understand colour vision. You can think of trichromatic theory as explaining the low-level reception of colour, whereas opponent theory says something about how this low-level information is subsequently combined.

7.3.3 The HSV model

This is a user-oriented model. The CIE and RGB models are quite technical, and do not correspond to how normal people think of colours. Most people think of colours in terms of hue, rather than primary components. The *Hue*, *Saturation* and *Value* (HSV) model tries to accommodate this. It is also known as HSB, with B for brightness.

We define HSV in terms of RGB. Suppose we have an RGB triple (all three values from 0.0 to 1.0). We know that (in general) the hue will be somewhat diluted by white (unless the colour is fully saturated). White W is equal parts of R,G and B:

$$W = \min(R, G, B)$$

Then saturation is

$$S = 1 - W = 1 - \min\{R, G, B\}.$$

This is the additive equivalent of subtracting black from CMY inks.

Subtracting the white from the colour, we get the basic hue. In RGB, the hue H is

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

But one of these values is zero (again, just as in CMYK), so we have only two numbers to represent. Conventionally this is done as an angle around a circle.

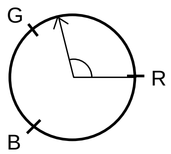

Put R at 0°, G at 120°, and B at 240° and work from there. Any point on the circumference is between two of the primaries, so we only need one value – the angle – to identify the non-zero primaries and the proportion of each.

Example The RGB triple $(0.3, 0.6, 0.2)$.

First we subtract the white component. The minimum of the three components is 0.2, so there is (0.2, 0.2, 0.2) of white in this colour. Subtracting leaves:

$$(0.1, 0.4, 0.0)$$

The two non-zero components are R and G. The proportions tell us that the colour is

$$0.4/(0.4 + 0.1) = 4/5$$

of the way around the arc between R and G. We can either express it as an angle:

$$120^\circ \times 4/5 = 96^\circ$$

or as a proportion of the whole circle:

$$96/360 = 0.27$$

The third component is *value*. This is

$$V = \max\{R, G, B\},$$

and is a measure of the brightness of the colour. We use maximum rather than the sum $R + G + B$ (which is a better value for the brightness energy), so that we control brightness uniformly across the spectrum. Without this, pure yellow (= red + green) would have a V value twice that of pure red.

So if we have an RGB colour $(0.3, 0.6, 0.2)$ this corresponds to $W = 0.2$, so saturation $S = 0.8$. Next, value $V = 0.6$. Finally hue: subtracting W from the RGB value we get $(0.1, 0.4, 0.0)$, which is 96°. Thus RGB $(0.3, 0.6, 0.2)$ is HSV $(96°, 0.8, 0.6)$, a half-bright fairly solid slightly reddish green (green with a yellow tinge).

Mixing red light and green light to get yellow is not very intuitive. However, tools for choosing colours using the HSV model can be easier to use than tools that use RGB directly. We directly pick a colour's hue, its saturation to determine how pure we want it, and its brightness. No need to try to remember that yellow is red plus green, we choose chose yellow from the palette offered by the tool, then choose the other dimensions around it.

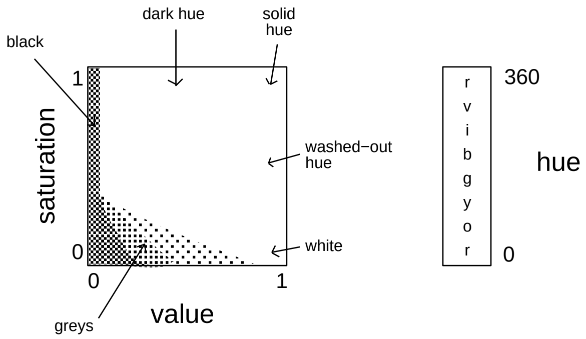

Colour Picker in Gimp

In RGB, the colour space is cube. In HSV, we can get an analogue of the RGB colour cube. This adds a third dimension to the hue circle:

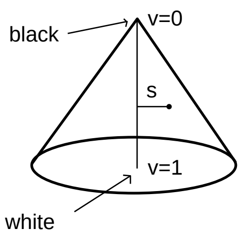

The vertical axis measures value, while radial distance measures saturation. Colours on the edge of the cone are fully saturated. All colours become black when they have zero value, which is why it is a cone rather than a cylinder.

Problems with HSV include:

- Non-uniformity as we vary the HSV parameters. For example as we vary H the perceived rate of change of hues is not uniform (slow rate of change near the primaries, faster elsewhere).

- Non-independent parameters. As above, the value of V, say, can have an effect on the values S can take.

There are several variants on the basic HSV model. For example, saturation can be measured as

$$S = (\max\{R, G, B\} - \min\{R, G, B\})/\max\{R, G, B\} = 1 - W/V$$

assuming $V \neq 0$. This adjusts saturation in proportion to the maximum amount of white available for a given colour (a colour with a high V can have a larger amount of white in it). In the basic model we must have $S \leq 1 - V$, while the adjusted model allows S the full range of values.

7.3.4 Other Models

There are many other models, each for (generally) good reasons.

- CIELUV, CIELAB: these are variants on CIE. CIELAB is used for subtractive, and CIELUV for additive displays. We find that *just noticeable differences* vary across the CIE diagram: a small movement in one part of the diagram may lead to a larger perceptual change of colour than an equal-sized step in another part of the diagram. Near the top of the 1931 diagram you have to move further to see a just noticeable difference (JND) than near the bottom. In fact the noticeable difference depends not only on the position but also on the direction. If you plot JNDs at various points over the 1931 diagram, you see ellipses of different size and different alignments. The CIELUV and CIELAB models distort CIE so that equal changes on the diagram lead to more-or-less equal perceptual changes: the ellipses become nearly circular and nearly the same size. These models are better for colour interpolation between close colours, not so good for relating distant colours.

- RGBA: the basic RGB plus an extra parameter called an *alpha* channel. This governs the opacity of a colour: 0 is fully transparent. For example, covering a blue with a green we get just the green in basic RGB. Using alpha channels we can declare the green to be 50% transparent. Then covering the blue again we get a mixture of 50% red with 50% green. This produces much more natural mixing of colours, e.g., green leaves of a tree let through a little of the blue sky behind. It is commonly available in PC graphics cards, being easy to implement.

  Alpha values are only supported poorly in the popular graphics file formats. GIF allows a limited form (one value can be used to mean transparent, so a given pixel can have transparency of 0% or 100% only, i.e., fully opaque or fully transparent). The PNG file format does support alpha values properly.

- Bath Colour Model: this takes colour mixing further. It is a 7 parameter model. A colour is regarded as consisting of pigment (reflective, 3 RGB parameters), plus a coloured medium (transparent, 3 RGB parameters), plus another parameter (the proportion of pigment to medium). This is good for mixing colours, and painting over other colours. More recently (2006) this was extended to the Projective Alpha Colour model, which incorporates the RGBA model.

- NCS: the Natural Colour System is used by architects, interior decorators, and paint manufacturers. Another user-oriented system, fairly close to HSV. It is based on the relationship of a colour to six *elementary* colours: red, green, blue, yellow, black and white. Colours form a 3D cone, horizontal sections of which are *colour circles*:

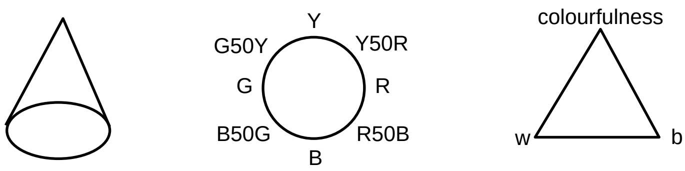

The colours are arranged in spectrum order about the circle, uniformly between red, yellow, green and blue.

A vertical section gives a *colour triangle*, providing a scale between *whiteness* and *blackness* (value); and also *colourfulness* (saturation). Colours are described as Y70Rs70c20, meaning a dark brown orange colour with 70% red; $s = 70$ means it has 70% blackness and $c = 20$ that it has 20% colourfulness.

- And so on.

Interestingly, tests on colour matching using various models seem to indicate that RGB is not so bad after all. Using RGB, users made matches the fastest, but the results were not so accurate. Using HSV they were slower to make matches, but the results were more accurate.

7.4 Practical Applications

7.4.1 Colour Television

How do we encode colour for transmission to a colour TV? This differs from the RGB that is used for the actual display tube. One reason is backwards compatibility with black and white television. Another reason is limited transmission bandwidth: sending R,G and B separately takes more bandwidth than is really necessary. What is interesting to us is that it takes advantage of the human's poor spatial perception of colour (via the eye's cones), whereas we are good at seeing monochrome contrast (via he eye's rods).

The PAL system (and other systems: SECAM in France, NTSC in North America) was designed to interoperate with the existing monochrome television services. Black and white deals with *luminance* (brightness), so PAL splits the colour information into a luminance and a *chroma* (colour) part. Old black and white TVs could use the luminance signal and ignore the chroma.

The black and white transmission still had to operate as before, so some way had to be found of including the colour information within the existing signal. This was done by putting the colour information in the part of the signal where the highest spatial frequencies would be. The bandwidth of the chroma signal is thus really rather small, so colour cannot change as rapidly as brightness. Thus we have a system with a fairly high definition luminance and a smear of low resolution colour over it (rather like a hand-tinted black and white photograph). This is not so bad as it corresponds to the relative sensitivity of the eye to these two things: the eye detects sharp changes in brightness at edges better than sharp changes in colour alone. By exploiting known characteristics of the eye, we can deliver a good engineering solution.

There is a problem with high spatial frequencies in the original scene, such as fine stripes on clothing. This produces high frequency components in the luminance signal, which are decoded as colour. Look at a test card to see the way fine lined areas shimmer in colour. You can sometimes see this effect if someone is wearing finely-striped clothing. Professional broadcasters try to avoid such materials.

Gamma correction (to correct for the non-linear response of the display output to input voltage) is applied at the studio before transmission as this was cheaper than requiring every TV set to have correction circuitry. (These days this would not be a problem, of course.)

The PAL luminance signal, named Y, is constructed as

$$Y = 0.30R + 0.59G + 0.11B,$$

for a colour with components R, G and B. These scale factors reflect the relative response of the eye: the blue component is much less important than red and green. This Y value is what is displayed on a black and white TV.

We now need to send the chroma values. We already sent Y, so we send values $R - Y$, and $B - Y$. These values are smaller than sending R and B directly, and so take less bandwidth. From this we can reconstruct immediately R, B, and then G from the equation for Y. We send differences for red and blue rather than green, as the difference for green would be larger, taking more bandwidth.

The YUV system, as used by PAL, has

$$Y = 0.30R + 0.59G + 0.11B,\ U = 0.49(B - Y),\ V = 0.88(R - Y),$$

being a collection of coefficients that fitted well with the old monochrome system.

The American NTSC is similar to PAL, but with slightly different coefficients. It is called the YIQ system. The colour signal has less bandwidth than PAL and the way of encoding them is simpler, so NTSC colours tend to be worse and we get more interference with the luminance signal.

7.5 Printing

Light absorbing devices, in particular printers, need to use a subtractive colour model.

Primaries for printers are the secondaries for displays, and vice-versa.

$$\begin{aligned} \text{cyan} &= 1 - \text{red} \\ \text{magenta} &= 1 - \text{green} \\ \text{yellow} &= 1 - \text{blue}. \end{aligned}$$

This is the CMY colour model.

Thus yellow ink absorbs only blue; cyan absorbs only red; magenta absorbs only green. In practice, a magenta ink will absorb a bit of blue, sometimes even 50% of the blue light. Thus the simple "one minus RGB" formula can produce picture that are yellower than they should be.

Adding these three inks in equal proportion should produce black, but experience in play school teaches us we only get a muddy brown. It is difficult to add the three colours in precisely the correct amounts. There is a further complicating factor: the quality (and colour) of the paper you print on. There tends to be a transmissive element *through* the paper which introduces an additive element to the equation (higher quality papers try to reduce the transmitted light). To combat this printers tend to use a *four* colour system, with cyan, magenta, yellow and black, the CMYK colour model. The K value (the *key*) is the minimum of C, M and Y: it is the amount of black in the colour. The C, M and Y values are the remaining amount of colour after we have subtracted the black:

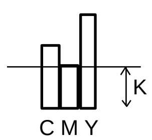

$$\begin{aligned} K &= \min\{C, M, Y\} \\ C &= C - K \\ M &= M - K \\ Y &= Y - K. \end{aligned}$$

one of C, M or Y is always zero. This produces a much better control over colours.

There are additional advantages to CMYK, too:

- coloured inks are typically more expensive than simple black ink
- if we have a black area composed of overlaid C, M, and Y mechanical tolerances will produce edges with coloured fringes (*registration errors*). The eye is most demanding at black and white edges so it is better to use black ink.
- printing a single black rather than three inks will dry faster, and the print can run faster.

Some printers use the *hexacolour* model, which takes CMYK and adds light cyan and light magenta, which produces better colour control.

7.6 Colour Interpolation

We use colour interpolation in many places. For example, Gouraud shading, anti-aliasing and blending images.

It is generally found that it is better to use RGB to interpolate between distant colours, e.g., in anti-aliasing, and to use HSV to interpolate between close colours, e.g., in Gouraud shading. Other colour models may be more appropriate for other situations. In practice, most computer graphics applications use RGB. Some newer applications, especially those which sit between computer graphics and image processing, do however use LAB or related "perceptual" spaces.

We can regard the interpolation as choosing colours on a straight line between two points in colour space: but which colour space? RGB, HSV, CIE, etc.? This is an important question as each gives different results for the interpolation, even when starting with the same pair of colours.

Example

Interpolate pure red to pure green. Using RGB, this is moving from point $(1, 0, 0)$ to $(0, 1, 0)$. For example, take the mid-point at $(0.5, 0.5, 0)$ (yellow).

Using HSV we move from $(0^\circ, 1, 1)$ to $(120^\circ, 1, 1)$, and midway is $(60^\circ, 1, 1)$.

Converting the RGB midway point to HSV we get $(60^\circ, 1, 0.5)$, which is not so bright as the HSV midway point.

In general, the result of interpolation depends on the underlying colour space.