Colour for Presentation Graphics

Ross Ihaka

Department of Statistics
University of Auckland

Abstract

Choosing a "good" set of colours for a graphical display is an important problem, but one which many data analysis practitioners are ill-equipped to solve. Most current graphics systems provide very little assistance in making good colour choices. Indeed, most systems require that a user specify their colours in ways which are closely related to the hardware representation of the colours rather than to the way we most naturally think about colour. This tends to encourage bad colour choices.

In this paper we'll examine some principles and software which enables us to make better colour choices. The ideas are close to those presented by Albert Munsell over 100 years ago, but they will be be developed in a more contemporary framework.

Note: The figures in this paper are designed to be displayed on a device with a "gamma" value of 2.2. Displays with a different gamma values will not faithfully reproduce the figures contained in this paper.

1   Colour Vision

It has been hypothesised (see [12] for example) that human colour vision evolved in three distinct stages. The original vision system was based on a single class of yellow/green sensitive "cone" cell. This initial system was monochrome only, sensing the world in terms of light and dark. This ancient system still provides the majority of the visual information we receive about the world, including information about geometric structure and motion. Colour is a later adaptation which provides us with additional information.

The oldest form of colour vision arose at a time probably predating the evolution of mammals. At that time a second class of cone cells became differentiated from the original yellow/green class. The cells of this new class were sensitive in the

Yellow

Green ← → Red

Blue

Figure 1: A graphical representation of the human colour axes.

blue/violet part of the spectrum. This adaptation provides us with the ability to discriminate colour on a yellow/blue axis. (It is interesting to note that this colour axis seems to be associated with our notion of warm and cool colours [12].)

A second adaptation occurred in primates about 30 million years ago. This divided the yellow/green sensitive cone cells into two kinds of cells, one more sensitive to green and the other more sensitive to red. The adaptation provides very fine colour discrimination in the red/green region of the spectrum. Such an adaptation is clearly advantageous when assessing the ripeness of many kinds of fruit and it hardly a surprise that it should have arisen in primates. This new adaptation provides us with colour discrimination on a red/green colour axis.

Because we seem to possess three natural colour axes (light/dark, yellow/blue and red/green), it is natural to describe colours as locations in a three dimensional space. A representation of the plane spanned by the yellow/blue and red/green axes is shown in figure 1. Although these axes have a natural physiological explanation, they do not correspond to our natural perception of colour. Our perception seems to correspond to the use of polar coordinates in this plane, with the angle, counterclockwise from red to a given colour correlating with our notion of colour *hue* and the radial distance from the origin to a given colour correlating with our notion of *chroma* or colourfulness.

2 Colour Synthesis and Colour Matching

The three dimensional nature of colour means that a wide range of colour sensations can be generated by mixing different amounts of three "primary" colours. Any three colours can serve as primaries, but the widest set of colour sensations can be created by single wavelength light sources placed close to the peak sensitivities of the three kinds of cone cells present in the eye. The most common choice for primary colours in lighting and display technology consists of monochromatic red, green and blue.

Much of what we know about our colour vision systems is based on colour matching experiments which seek to quantify the amounts of three given primaries needed to match a given colour or set of colours. Suppose that to match a given colour we require amounts R, G and B of a particular set of red, green and blue

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

One possible matching experiment consists of seeing how single wavelength colours can be matched by a given set of primaries. For a given wavelength λ, suppose that the tristimulus values for a match are $\bar{r}_\lambda$, $\bar{b}_\lambda$ and $\bar{g}_\lambda$. Regarded as functions of λ, $\bar{r}_\lambda$, $\bar{b}_\lambda$ and $\bar{g}_\lambda$ are known as the *colour matching functions* for the given set of primaries.

Colour matching functions are important because it has been observed that colour matching is linear. This means that it is possible to determine the match for a mixture of wavelengths by using the matching functions to obtain a match at each frequency and then summing across frequencies. If the function s_λ gives the amount of each wavelength present in a colour, then the tristimulus values for the mixture are given by

$$R = \int_\lambda s_\lambda \bar{r}_\lambda \, d\lambda, \qquad G = \int_\lambda s_\lambda \bar{g}_\lambda \, d\lambda, \qquad B = \int_\lambda s_\lambda \bar{b}_\lambda \, d\lambda.$$

It is common to normalise the tristimulus values by dividing them by their sum to obtain the relative amounts of red, green and blue required to obtain a match. The normalised values

$$r = \frac{R}{R + G + B}, \qquad g = \frac{G}{R + G + B}, \qquad b = \frac{B}{R + G + B},$$

are referred to as the *chomaticities* for the match. Because chromaticities sum to one it is possible to compute any one of them from the other two. It is usual to reduce dimensionality by discarding the third chromaticity. This makes it possible to plot colours in two dimensions, which is often convenient.

If two colours have tristimulus values $\mathbf{C}_1 = (R_1, G_1, B_1)$ and $\mathbf{C}_2 = (R_2, G_2, B_2)$, then any mixture of these colours must have a tristimulus value which lies on the line joining $\mathbf{C}_1$ and $\mathbf{C}_2$. This property is also preserved for chromaticities, which makes chromaticity plots useful for describing colour mixtures.

3 Colour Standards

In the 1920s and 1930s a series of colour matching experiments by W. D. Wright and J. Guild showed that there was enough consistency in colour matching for it to be possible to define a "standard observer," whose (hypothetical) colour matching functions would represent a typical human colour response and provide an absolute reference standard for colour. This standardisation was carried out under the auspices of the *Commission Internationale de l'Éclairage* (the CIE) in 1931.

Although a standard observer had been agreed upon, there was still the issue of which set of primaries should be used for colour matching. With any real set of primaries, it is not possible to match all possible colours — there will always be some colours which can only be matched after adding a small amount of one of the primaries to that colour. In such cases, the tristimulus value of that primary is taken to be negative. This leads to matching functions which take on negative values for some ranges of λ.

The CIE felt that the use of negative matching functions was likely to be a source of confusion and errors. Their solution to this problem was to use a set of "imaginary" primaries to obtain colour matching functions. This made it possible to define all colours with positive tristimulus values and hence positive chromaticities.

Roughly speaking the CIE primaries can be thought of as supersaturated shades of violet-blue, yellow-green and orange. The tristimulus values corresponding to these primaries are denoted by X, Y and Z and chromaticities by x, y and z. A full colour description can be obtained by plotting the tristimulus values in a three-dimensional space and a partial one by plotting the x and y chromaticity values on a two dimensional chromaticity diagram. As part of their choice of primaries, the CIE arranged that the Y tristimulus value correspond to the apparent brightness of colours.

While the CIE tristimulus values provide a complete colour description they do not correspond to the way we perceive colours. In particular, distances between tristimulus triples need not reflect the degree to which we perceive the colours to differ. The CIE recognised this and introduced new spaces which more closely reflect the way we perceive colour.

The two perceptually based spaces introduced by the CIE in 1976 are the *CIELUV* and *CIELAB* spaces. The CIELUV space is generally preferred by those who work with emissive colour technologies (such as computer displays) and the CIELAB space is preferred by those working with dyes and pigments (such as in the printing and textile industries).

I will confine my remarks to the CIELUV space. This is because the presentation graphics I will describe are typically created on computer displays and this is where colour experimentation takes place.

The CIELUV space is based on the uniform chromaticity coordinates:

$$u' = 4X/(X + 15Y + 3Z)$$

$$v' = 9Y/(X + 15Y + 3Z).$$

These coordinates correspond to positions on a red/green and yellow blue axes. To produce a full colour description these coordinates must be scaled and combined with brightness information. The full CIELUV spaces is defined by the coordinates.

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

$$u^* = 13L^*(u' - u'_n)$$

$$v^* = 13L^*(v' - v'_n)$$

where Y_n is the Y tristimulus value and u'_n and v'_n are the uniform chromaticities for the white-point of the display.

This parameterisation means that the u^* axis corresponds to the horizontal axis of figure 1 and the v^* axis to the vertical one. The intersection of the two axes takes place at the colour-neutral white point.

4 Device Dependent Colour

Cheap colour computer displays became widely available in the 1980s. The earliest displays were very limited in the colours which could be displayed. Typically such displays were limited to a fixed set of just 8 or 16 rather lurid colours. A second generation of cheap colour displays could display 256 colours. While this

is considerably better than first-generation displays, it is still quite limiting. The restriction of web pages to a fixed palette of 216 colours is a legacy which dates from the heyday of this kind of display.

The current generation of colour displays is much more satisfactory. They can display as many colours it is possible to discriminate with the human eye. Unfortunately, the software which controls these displays often retains the effects of the restrictions earlier generations of graphics hardware.

Current colour displays are based on red, green and blue primaries. Colours are described by directly specifying the levels of red green and blue to be combined. It is convenient to think of the RGB levels as lying in the interval $[0, 1]$, although in practise they are more likely to be specified as an 8-bit integer in the range $0 - 255$. The colour produced by a given RGB description is also affected by the response characteristics of the display. For most displays, the principal characteristic of interest is the *display gamma*, γ. The intensity I of the on-screen primaries is related to the specified level L for that primary by

$$I = L^{\gamma}, \qquad L \in [0, 1].$$

It is possible for displays to have a distinct gamma value for each of their primaries, but in they are generally quite close in value.

If colours are to be made to appear identically on different displays, the display gammas must be taken into account. For most current displays the gamma value is close to 2.2. and there has been some attempt to standardise colour description so that they will display faithfully on this kind of display [9].

RGB colour descriptions do not correspond naturally to the way in which we think about colour. Because of this, a number of other equivalent colour descriptions have been formulated. These include HSV, HSL, HSB and others. These new spaces seek to transform the RGB coordinates to a more intuitive set. In the case of HSV, these coordinates are *hue*, *saturation* and *value*, which loosely correspond to dominant wavelength, colour purity and brightness. Although the coordinates do have a more intuitive interpretation than RGB they are device-dependent and do not offer the ease of interpretation of the perceptually based CIE spaces.

5   Colour Harmony

The colour descriptions of the previous sections provide a way of precisely specifying colours which are to appear in a figure or graph. They do not, however, address the issue of which colours should be used.

Choosing a set of colours which work well together is a challenging task for anyone who does not have an intuitive gift for colour. Some general guidance on colour choice is available in books on art and graphic design. These books suggest the use of complementary colours, split complementaries, triads and tetrads. Most of the advice is based on the use of a vaguely described "colour wheel," and does not recognise the fact that there are many different colour wheels to choose from.

The notable exception to this rule is to be found in the work of the noted 19th century colourist Albert Munsell. Munsell developed a colour notation system ([14], [5], [1]) which he used in teaching. The system is deeply rooted in how we perceive colour but has a strong quantitative basis. In addition Munsell gives well defined, quantitative rules which can be used to choose harmonious sets of colours. Munsell's work has always been appreciated in publishing and related graphic arts, and it now

R 5/5, BG 5/5, N5

R 6/5, BG 4/5, N5

Figure 2: Illustrations of colour balance, after Munsell. These figures are redrawn from [1].

appears to be undergoing a rediscovery by those working in user-interface design and visualisation (see [10], [13] and [8] for example).

Munsell describes colour in terms of *hue*, *value* and *chroma*; hue corresponding to dominant wavelength, value to brightness and chroma to colourfulness. Unlike saturation, which is a statement of colourfulness relative to the maximum possible for a given hue and value, chroma is an absolute measure of colourfulness (the maximum chroma possible for red is much greater than that for green). Munsell divided the circle of hues into 5 main hues — R, Y, G, B, P (for red, yellow, green, blue and purple). He placed 5 intermediate hues (YR, GY, BG, PB and RP) between them and provided a way of specifying finer hue divisions on a decimal scale. He also divided the value range into 10 equal equal steps and provided a similar quantification of chroma. Under Munsell's system the colour specification R 4/5 means a hue of red with a value of 4 and a chroma of 5. A specification of N 5 refers to a mid-range (neutral) grey.

Munsell's ideas on colour harmony are rooted in the notion of "colour balance." Munsell uses the term balance in a variety of ways, but usually it means that set of colours is chosen to be centred on a mid-range or neutral value. Figure 2 shows an illustration of the Munsell concept of balance. The picture in top part of the figure uses the colours R 5/5 and BG 5/5 on an N5 background. R 5/5 and BG 5/5 are at opposite sides of the Munsell colour wheel and hence they balance at N5. Using these two colours together with an N5 background produces a colour scheme which balances at N5. The lower figure shows the picture as above, but with the value of the red raised by a step and the value of the blue-green lowered by a step. Again, N5 lies at the centre of these two colours and this again creates a balanced colour choice. (It is interesting to note that these small changes in colour value make the elements of the picture stand out. This is because we rely on light/dark contrast to demarcate basic geometric structure.)

Munsell was particular fond of the 5/5 colours in his system. The five principal Munsell 5/5 colours are shown in figure 3. An appreciation of the balance in this colours can be obtained by comparing these colours with the same hues drawn at full saturation from HSV space. The colours in this second figure differ wildly in value and chroma and cannot be considered to be balanced or in harmony.

In addition to the principal of "balance on grey" Munsell gives an number of other ways in which balanced colours can be chosen. Most of these suggestions amount to choosing colours at equally spaced points along smooth paths through his (perceptually uniform) colour space. This is qualitatively the same type of recommendation made by Cynthia Brewer [2], [3], [4].

Munsell makes another point which is useful to keep in mind. The intensity of colour which should be used is dependent on the area that that colour is to occupy. Small areas need to be much more colourful than larger ones. Again, Munsell describes this as a matter of balance.

Although Munsell's colour balance recommendations are specific to his colour space, they apply equally to any perceptually uniform space based on correlates of hue, value and chroma. In particular, reasonable results are obtained by applying Munsell's principles to the CIELUV and CIELAB spaces. This is convenient because while there are approximate methods for mapping Munsell's colour notation to modern devices (see [11]) a precise mapping is not easy.

Figure 3: The principal Munsell 5/5 colours. From the top these are R 5/5, Y 5/5, G 5/5, B 5/5 and P 5/5. This figure is redrawn from [1].

Figure 4: The same images as figure 3, but drawn with full saturation HSV colours.

6   Colour Use In Presentation Graphics

Colour can be a very useful addition to presentation graphics such as bar charts and mosaic plots. The typical use of colour in such graphs is to indicate which of several groups various areas of a graph correspond to.

Choosing a good set of colours means addressing both perceptual and aesthetic issues. The perceptual issues usually impose clear constraints on colour choices. Some examples of these constraints follow:

- It is best to avoid large areas of high-chroma colours in graphs which must hold a users attention for extended periods. This is because such colours tend to produce after-image effects which are can be distracting.

- If the size of the areas presented in a graph is important, then the areas should be rendered with colours of similar luminance. This is because lighter colours tend to make areas look larger than darker colours. See Cleveland and McGill [6] for example.

- When colours are used to indicate group membership, the colours should be easy to distinguish.

Aesthetic considerations are much harder to state, but we would do well to heed Munsell's advice and choose colours in a methodical, balanced way.

There is a simple way of choosing colours which meets all the constraints above. The method can be stated simply as follows. In a perceptually uniform space, choose colours which have equal luminance and chroma and correspond to set of evenly spaced hues.

To be more precise let us consider how we might achieve such a colour choice in CIELUV space. First we must consider transforming the u^*, v^* plane to polar coordinates, the radial distance C corresponds to chroma and the angle h to hue. First we must settle on a choice of luminance and chroma. The values $C = 55$ and $L^* = 75$ provides colours which are about as colourful as possible on an RGB monitor under these constraints. The values $C = 35$ and $L^* = 85$ provide a slightly more subdued set of colours. The corresponding colours are shown in figure 5. Both sets of colours are rather lighter and brighter than the Munsell 5/5 colours but they do provide the same type of "balance on grey" that the Munsell colours do, it is just a lighter grey.

This method of choosing colours can be carried out with a small R function called `hcl` available from `http://www.stat.auckland.ac.nz/~ihaka/colour`. The function has three required (vector) arguments, `h` giving *hue* as an angle in $[0, 360]$, `c` which gives the *chroma* and `l` which gives the *lightness* as a value in [0,100]. The function also has an optional argument `gamma` which defaults to 2.2 and gives the gamma for the display being used to display the colours.The function returns the specified colours as a hexadecimal character string of the form `"#RRGGBB"` which is suitable as a colour description in many applications. The function can be easily modified to produce an alternative colour description.

Although the method outlined above is quite constrained, it is can be used to produce quite a wide range of colour selections. As noted above, lightness and chroma can be varied and it is also possible to restrict the range of hues to obtain various kinds of effects. We will illustrate this by showing some variations on a bar plot presented in [7].

Figure 5: Colour selections made from LUV space.

There are four classes in this data set and this means choosing four colours to represent them. In the examples that follow, we will fix $C = 35$ and $L^* = 85$. As noted above, this produces relatively muted colours. If stronger colours are desired, the value of C can be increased and the lightness decreased.

Figure 6: These colours correspond to four equally spaced hues at the angles 30°, 120°, 210° and 300°. This is what is known as a colour "tetrad." Such a colour choice generally described as dynamic or exciting.

Figure 7: These colours correspond to four equally spaced hues at the angles 60°, 120°, 180° and 240°. These hues only run half-way around the colour circle. The colours are still easily distinguished and the choice is more harmonious that the previous one.

Figure 8: These colours correspond to four equally spaced hues at the angles 210°, 160°, 110° and 60°. These hues similar to the previous ones, but have been modified slightly and reordered so that they provide a visual metaphor for the season they represent.

Figure 9: These colours correspond to four equally spaced hues at the angles 270°, 230°, 190° and 150°. These are "cool" colours, chosen from the blue side of the colour wheel.
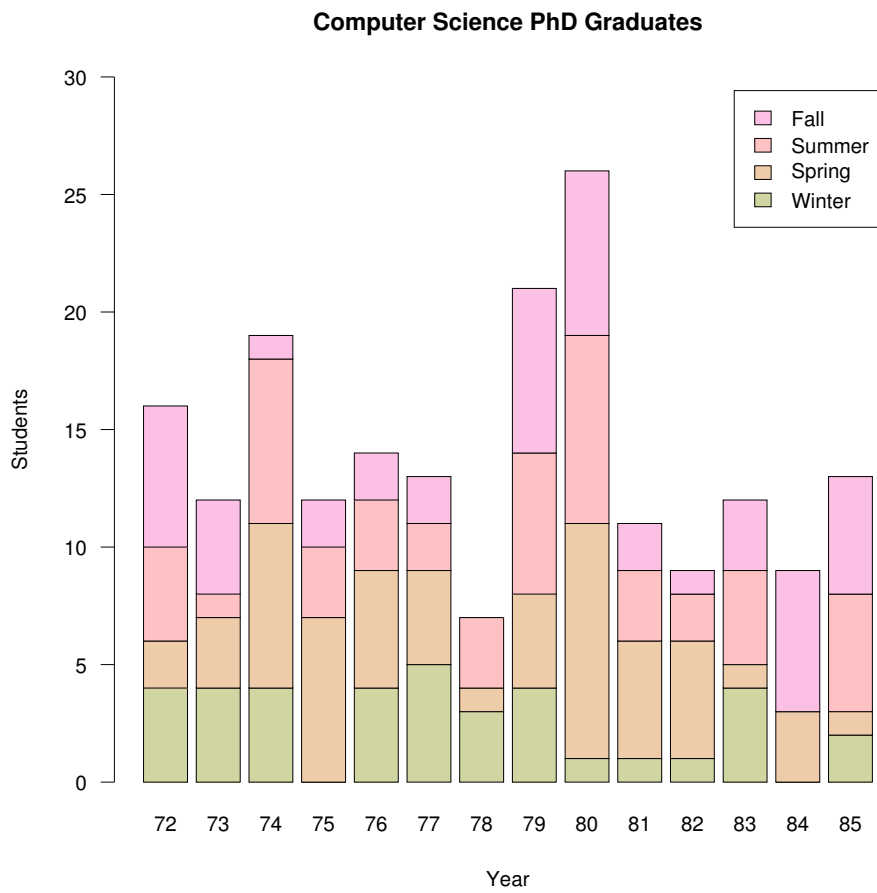
Figure 10: These colours correspond to four equally spaced hues at the angles 90°, 50°, 10° and −30°. These are "warm" colours, chosen from the yellow side of the colour wheel.

7 Conclusions and Future Work

Using perceptually uniform colour spaces and some simple ideas on colour balance, it is possible to make the process of choosing colours for area fills in presentation graphics a relatively simple one. Perceptual and aesthetic considerations place natural constraints on how colours should be chosen. Although the sets of colours which can be chosen is highly constrained, there is still sufficient flexibility for a variety of colour effects to be achieved.

If the choice of colours outined in this paper is too constrained, it is possible to use the methods outlined by Brewer [2], [3], [4], to produce a less constrained set. This is generally a much harder exercise than the method outlined above and it does not protect against colour related size illusions.

This paper addresses just one particular colour choice problem. There are other important problems, such as the selection of colours for glyphs and lines. It is likely that perceptual and aesthetic considerations will provide natural constraints in these cases too, and provide a way of simplifying the colour choices which must be made.

References

[1] F. Birren. *A Grammar of Color.* Van Nostrand Reinholt, 1969. This is a repackaging of [5].

[2] C. A. Brewer. Color use guidelines for mapping and visualization. In A.M. MacEachren and D.R.F. Taylor, editors, *Visualization in Modern Cartography*, chapter 7, pages 123–147. Elsevier Science, Tarrytown, NY, 1994.

[3] C. A. Brewer. Guidelines for use of the perceptual dimensions of color for mapping and visualization. In J. Bares, editor, *Color Hard Copy and Graphic Arts III, Proceedings of the International Society for Optical Engineering (SPIE)*, volume 2171, pages 54–63. San Jose, 1994.

[4] C. A. Brewer. Color use guidelines for data representation. In *Proceedings of the Section on Statistical Graphics, American Statistical Association*, pages 55–60, Alexandria VA, 1999.

[5] T. M. Cleland. *A Grammar of Color.* Strathmore Paper Company, 1921.

[6] W. S. Cleveland. A color-caused optical illusion on a statistical graph. *The American Statistician*, 37(2):101–105, 1983.

[7] J. D. Foley, A. van Dam, S. K. Feiner, and J. F. Hughes. *Computer Graphics: Principles and Practice.* Addison-Wesley, Reading MA, 1990.

[8] C. G. Healy. Choosing effective colours for data visualization. Technical Report 10, Department of Computer Science, University of British Columbia, 1996.

[9] Hewlett-Packard Inc. `http://www.srgb.com`.

[10] P. Lyons, G. Moretti, and M. Wilson. Colour group selection for computer interfaces. In B.E. Rogowitz and T.N. Pappas, editors, *Human Vision and*

Electronic Imaging V, volume 3959, pages 302–313, San Jose, USA, 2000. SPIE – The International Society for Optical Engineering.

[11] M. Miyahara and Y. Yoshida. Mathematical transform of (r, g, b) color data to munsell (h, v, c) color data. In *SPIE Conference on Visual Communications and Image Processing*, volume 1001, pages 650–657. 1988.

[12] J. Mollon. Seeing colour. In T. Lamb and J. Bourriau, editors, *Colour: Art and Science*. Cambridge Unversity Press, 1995.

[13] G. Moretti and P. Lyons. Tools for the selection of colour palettes. In *New Zealand Symposium On Computer-Human Interaction*, University of Waikato, New Zealand, July 2002.

[14] A. H. Munsell. *A Color Notation*. Munsell Color Company, Baltimore, twelfth edition, 1947.