

Copyright is owned by the Author of the thesis. Permission is given for a copy to be downloaded by an individual for the purpose of research and private study only. The thesis may not be reproduced elsewhere without the permission of the Author.

A CALCULATION OF COLOURS

TOWARDS THE AUTOMATIC CREATION
OF
GRAPHICAL USER INTERFACE COLOUR SCHEMES

A THESIS PRESENTED IN PARTIAL FULFILMENT OF THE REQUIREMENTS FOR THE
DEGREE OF
DOCTOR OF PHILOSOPHY
IN
COMPUTER SCIENCE
AT
MASSEY UNIVERSITY, PALMERSTON NORTH,
NEW ZEALAND.

GIOVANNI S. MORETTI

2010



Dedicated to Edris and Serafino, my parents.

Abstract

Interface colour scheme design is complex, but important. Most software allows users to choose the colours of single items individually and out of context, but does not acknowledge colour schemes or aid in their design. Creating colour schemes by picking individual colours can be time-consuming, error-prone, and frustrating, and the results are often mediocre, especially for those without colour design skills. Further, as colour harmony arises from the interactions between all of the coloured elements, anticipating the overall effect of changing the colour of any single element can be difficult.

This research explores the feasibility of extending artistic colour harmony models to include factors pertinent to user interface design. An extended colour harmony model is proposed and used as the basis for an objective function that can algorithmically assess the colour relationships in an interface colour scheme. Its assessments have been found to agree well with human evaluations and have been used as part of a process to automatically create harmonious and usable interface colour schemes.

A three stage process for the design of interface colour schemes is described. In the first stage, the designer specifies, in broad terms and without requiring colour design expertise, colouring constraints such as grouping and distinguishability that are needed to ensure that the colouring of interface elements reflects their semantics.

The second stage is an optimisation process that chooses colour relationships to satisfy the competing requirements of harmonious colour usage, any designer-specified constraints, and readability. It produces sets of coordinates that constitute abstract colour schemes: they define only relationships between coloured items, not real colours.

In the third and final stage, a user interactively maps an abstract scheme to one or more real colour schemes. The colours can be fine-tuned as a set (but not altered individually), to allow for such “soft” factors as personal, contextual and cultural considerations, while preserving the integrity of the design embodied in the abstract scheme. The colours in the displayed interface are updated continuously, so users can interactively explore a large number of colour schemes, all of which have readable text, distinguishable controls, and conform to the principles of colour harmony.

Experimental trials using a proof-of-concept implementation called the Colour Harmoniser have been used to evaluate a method of holistic colour adjustment and the resulting colour schemes. The results indicate that the holistic controls are easy to understand and effective, and that the automatically produced colour schemes, prior to fine-tuning, are comparable in quality to many manually created schemes, and after fine-tuning, are generally better.

By designing schemes that incorporate colouring constraints specified by the user prior to scheme creation, and enabling the user to interactively fine-tune the schemes after creation, there is no need to specify or incorporate the subtle and not well understood factors that determine whether any particular set of colours is “suitable”. Instead, the approach used produces broadly harmonious schemes, and defers to the developer in the choice of the final colours.

Acknowledgements

I would like to thank those who have contributed to this research project:

my family, for their support and understanding during what has turned out to be a longer than expected “absence”;

those who helped with the experimental work, especially: Dr Heather Meikle and the students from Palmerston North Girls High School; the staff at TeManawa Museum and Science Centre; staff and students at Massey University and at the UCOL School of Photography, Art & Design;

Dr Alasdair Noble, Dr Geoff Jones and Dr Jonathan Godfrey for their assistance with the statistics;

Dr Chris Phillips and to Dr Greg Allan their insightful questions and suggestions, and to my colleagues in the School of Engineering and Advanced Technology for their encouragement; and

my supervisors, Mr Paul Lyons and Dr Stephen Marsland, for their suggestions, and their patience and support with the so many aspects that, looking back, are now so obvious.

Contents

| | |
|--|------------|
| Abstract | v |
| Acknowledgements | vii |
| 1 Introduction | 1 |
| 1.1 Colouring user interfaces | 2 |
| 1.1.1 Colour selection for computer interfaces is important | 3 |
| 1.1.2 Colour selection is hard | 4 |
| 1.1.3 Colour selection tools are naïve | 4 |
| 1.2 Guidelines for the design of interface colour schemes | 5 |
| 1.2.1 User interfaces are not art | 6 |
| 1.2.2 Constraints on the use of colour in interface colour schemes | 6 |
| 1.2.3 The design process: client and designer | 7 |
| 1.3 The thesis | 8 |
| 2 Background and related work | 11 |
| 2.1 The impact of colour and colour schemes | 12 |
| 2.1.1 Colour: associations and emotions | 12 |
| 2.1.2 Emotive impact of multiple colours – colour schemes | 14 |
| 2.1.3 Colour and culture | 15 |
| 2.1.4 Colour scheme appeal is age-specific | 16 |
| 2.2 The use of colour in graphical user interfaces | 16 |
| 2.2.1 Aesthetics affect usability | 17 |
| 2.2.2 Colour affects usability | 17 |
| 2.2.3 Colour affects visual search and recognition | 19 |
| 2.2.4 Colour, culture, emotions and commerce | 19 |
| 2.2.5 Standards bodies and colour use | 21 |
| 2.2.6 Neurophysiological factors affecting the use of colour | 21 |
| 2.2.7 Colour use in graphical user interfaces – summary | 25 |
| 2.3 Human colour perception | 26 |
| 2.3.1 The derivation of perceptually important primaries | 27 |
| 2.3.2 Colour vision deficiencies (CVD) | 29 |
| 2.4 CIE colour spaces | 30 |
| 2.4.1 The CIEXYZ colour space | 31 |

| | | |
|----------|--|-----------|
| 2.4.2 | The CIE chromaticity diagram | 32 |
| 2.4.3 | The white point | 33 |
| 2.4.4 | Standard Illuminants | 33 |
| 2.5 | Adaptation of human colour vision to the viewing conditions | 34 |
| 2.5.1 | Colour constancy and adaptation | 35 |
| 2.6 | Colour display systems | 36 |
| 2.6.1 | Colour management | 38 |
| 2.7 | Colour spaces useful in colour scheme design | 40 |
| 2.8 | Perceptually-based colour arrangements | 43 |
| 2.8.1 | Munsell's colour order system | 44 |
| 2.8.2 | Measuring colour difference and perceptual uniformity | 45 |
| 2.8.3 | Human colour sensitivity and discrimination | 46 |
| 2.8.4 | The CIELAB colour space | 48 |
| 2.8.5 | RGB colour spaces | 50 |
| 2.9 | Abstract representations of colour harmony | 52 |
| 2.9.1 | The concept of colour harmony | 52 |
| 2.9.2 | Necessary conditions for colour harmony | 53 |
| 2.9.3 | Harmonious colour selection using colour wheels | 54 |
| 2.9.4 | Interpolation in colour space | 56 |
| 2.9.5 | The difference between a colour scheme and a palette | 58 |
| 2.9.6 | Munsell's theory of colour harmony | 59 |
| 2.9.7 | Other models of colour harmony | 62 |
| 2.9.8 | The readability of text | 62 |
| 2.10 | Computer-based tools for colour scheme design | 66 |
| 2.10.1 | General purpose colour selectors | 66 |
| 2.10.2 | Template-based colour scheme design | 68 |
| 2.11 | Automating colour scheme design | 70 |
| 2.11.1 | Automating colour palette design | 71 |
| 2.11.2 | Interactive design using colours in context | 75 |
| 2.11.3 | Automated generation of interface colour schemes | 76 |
| 2.12 | Colour schemes for viewers with colour vision deficiencies | 86 |
| 2.13 | Summary | 88 |
| 3 | The conceptual foundations of an automatic interface colouring system | 91 |
| 3.1 | Augmenting colour harmony heuristics for use in GUI design | 93 |
| 3.2 | The developer's input | 95 |
| 3.3 | The information resulting from the interface characterisation | 96 |
| 3.4 | The concept of an abstract colour molecule | 96 |
| 3.5 | The conflicting demands of colour harmony and GUI design | 100 |
| 3.6 | Colour schemes are independent of colour | 101 |
| 3.6.1 | Initial colour atom positions within a molecule | 103 |
| 3.7 | Improving a colour scheme – optimisation | 103 |
| 3.8 | Deriving real colours from abstract colour space coordinates | 104 |

| | | |
|----------|--|------------|
| 3.9 | Personalising the schemes created by the optimiser | 105 |
| 3.10 | The modular structure of the proposed architecture | 106 |
| 3.11 | Summary | 107 |
| 4 | The realisation of a colour harmoniser | 111 |
| 4.1 | Characterisation of an interface | 111 |
| 4.1.1 | Determining the area of the colourable interface elements | 112 |
| 4.2 | Including developer-specified colouring constraints | 114 |
| 4.2.1 | Colour and user interface semantics | 114 |
| 4.2.2 | Meta-objects and identically coloured items | 116 |
| 4.2.3 | Distinguishability constraints | 119 |
| 4.3 | Ensuring legibility of the coloured interface | 120 |
| 4.4 | The creation of an abstract colour molecule | 123 |
| 4.4.1 | Creating prototypical colour molecules in abstract colour space | 123 |
| 4.4.2 | Augmenting the wireframe with achromatic colours | 124 |
| 4.5 | The algorithmic evaluation of colour schemes | 126 |
| 4.5.1 | Normalising the fitness term scores | 127 |
| 4.5.2 | Colour strength balance contribution | 128 |
| 4.5.3 | The black-white wireframe option and disabling lightness balance | 130 |
| 4.5.4 | The wireframe alignment contribution | 130 |
| 4.5.5 | The distinguishability contribution | 131 |
| 4.5.6 | The readability contribution | 132 |
| 4.5.7 | Ensuring that critical colouring flaws are not overlooked | 133 |
| 4.5.8 | The fitness function | 136 |
| 4.6 | Colour molecule optimisation | 138 |
| 4.6.1 | Optimisation requirements | 138 |
| 4.6.2 | Genetic optimisation applied to colour molecules | 139 |
| 4.6.3 | Multiple solutions from the optimisation | 142 |
| 4.6.4 | Avoiding premature convergence | 142 |
| 4.7 | Real colour schemes from an abstract molecule | 147 |
| 4.7.1 | Mapping abstract colour space onto a perceptually uniform space | 149 |
| 4.7.2 | The LAB scaling factor | 151 |
| 4.8 | Colour scheme personalisation | 156 |
| 4.8.1 | Controlling the dominant hues in the scheme | 156 |
| 4.8.2 | Controlling the colour saturation | 156 |
| 4.8.3 | Allowing for the “natural” placement of colours | 158 |
| 4.8.4 | Implementing the personalisation controls | 161 |
| 4.8.5 | Direct manipulation colour scheme exploration | 162 |
| 4.8.6 | Colour Harmoniser implementation – summary | 162 |
| 5 | The experimental validation of the fitness function | 167 |
| 5.1 | An overview of the experiments and methods | 167 |
| 5.1.1 | The evaluation interface used in the colouring trials | 167 |
| 5.1.2 | The inclusion of CVD participants in the experimental trials | 172 |

| | | |
|----------|---|------------|
| 5.1.3 | The participants in the experimental trials | 173 |
| 5.1.4 | The experimental conditions | 173 |
| 5.2 | The experiments used to evaluate the fitness function | 174 |
| 5.2.1 | Assessing the relationship between the fitness function score and human colour scheme assessment | 175 |
| 5.2.2 | A pilot fitness function validation experiment | 175 |
| 5.2.3 | The main fitness function validation experiment | 182 |
| 5.2.4 | An experiment to establish the independence of fitness function terms | 187 |
| 5.2.5 | Estimating the fitness function terms weights from experimental data | 189 |
| 5.3 | Conclusion | 195 |
| 6 | The experimental evaluation of the Colour Harmoniser | 197 |
| 6.1 | Compare the Methods: traditional vs. Harmoniser-based scheme creation | 198 |
| 6.1.1 | Aims of the experiment | 198 |
| 6.1.2 | Experimental design | 198 |
| 6.1.3 | Experimental method | 198 |
| 6.1.4 | The interfaces used to control the experiment | 201 |
| 6.1.5 | The results of the comparative colour scheme creation experiment | 212 |
| 6.2 | Evaluating colour schemes created using the Colour Harmoniser | 220 |
| 6.2.1 | Aim of the experiment | 220 |
| 6.2.2 | Experimental design | 220 |
| 6.2.3 | Experimental method | 221 |
| 6.2.4 | Results | 224 |
| 6.3 | Discussion | 234 |
| 7 | Conclusions | 237 |
| 7.1 | Summary | 237 |
| 7.2 | The research outcomes | 238 |
| 7.2.1 | It is possible to design a computationally tractable model of colour harmony that can be used to evaluate user interface colour schemes | 238 |
| 7.2.2 | It is possible to create harmonious schemes in an abstract colour space | 241 |
| 7.2.3 | It is possible to build a software tool that incorporates the aug- mented model of colour harmony | 243 |
| 7.2.4 | Users of an algorithmic colour harmony tool do not need colour design expertise | 245 |
| 7.2.5 | The Colour Harmoniser method produces good colour schemes . | 246 |
| 7.2.6 | The Colour Harmoniser method of colour adjustment is easy to understand and makes finding good colour schemes easier | 247 |
| 7.3 | Confirmation of the thesis | 249 |
| 7.4 | Contributions to knowledge | 250 |

| | | |
|----------|---|------------|
| 7.5 | Possible extensions | 250 |
| 7.5.1 | Extending the model of user interface colour harmony | 250 |
| 7.5.2 | Extensions to the holistic creation and adjustment of colour schemes | 253 |
| 7.6 | General conclusions | 254 |
| A | Expanded results table for the “Compare the Results” experiment | 257 |
| A.0.1 | Cross-tabulation of the t-test results of the difference in means for differing colour scheme creation methods | 257 |
| B | How professional and How artistic were the schemes from each method? | 259 |
| C | Colour-related publications and presentations by the author | 261 |
| D | A Calculation of Colours – Vincent van Gogh | 263 |
| | Bibliography | 269 |

List of Figures

| | | |
|------|---|----|
| 2.1 | Nine colours, all the same hue, but with differing lightness and saturations in each image. Left: high saturation; centre, lighter (tinted – mixed with white); right, desaturated and darker (shades). The emotional reaction to each of the three figures is likely to be quite different. All three redrawn from Hornung (2005). | 13 |
| 2.2 | The adjacency of high saturation complementary colours can cause unpleasant shimmering effects at the boundary, an example of a visual artifact related to high saturation chromatic contrast. | 22 |
| 2.3 | The gradations in colour apparent in the left-hand image are visual artifacts – each band is a single colour. As shown by the right-hand image, the gradation disappears if the regions are separated. | 22 |
| 2.4 | An example of the conflict between foreground, background and negative space in high contrast areas (based on an example by Tufte 1990). The approximately equal width of the text and the white spacers in the left image (black in the right) can cause the spacers to appear to be the foreground elements instead of the text. | 23 |
| 2.5 | Left: warm–cool contrasts can cause tensions in figure–ground perception. A cool (blue) foreground object can make a warm background object appear to be wrapped around the cool object rather than behind it. Right: lighter colours appear closer to the viewer, even when of the same hue. | 25 |
| 2.6 | An indication of the correspondence between colour and wavelength in the range of human vision (380–750nm). | 26 |
| 2.7 | LMS cone responses to the varying wavelengths of light. The shorter (lower) wavelengths correspond to blues, the longer to reds. Plotted from Stockman and Sharpe (2000) 2° cone fundamentals. | 27 |
| 2.8 | The outputs from the long, medium and short sensitive cones in the eye are combined by the ganglia to derive luminance and opponent colour signals, redrawn from Fairchild (1998). | 29 |
| 2.9 | The dimensions resulting from the opponent processing: one for lightness, two for colour. | 29 |
| 2.10 | The CIEXYZ colour matching functions that indicate the contribution of three imaginary primaries to create the appearance of colours for the range of human colour vision. | 31 |

| | | |
|------|--|----|
| 2.11 | The CIE Chromaticity diagram | 33 |
| 2.12 | The sources of variability in colour scheme creation include both the designer's and the viewer's preferences, equipment and viewing conditions. Nevertheless, colour schemes can be created and disseminated, with a large degree of agreement as to their visual appeal, using only triplets of RGB values. | 35 |
| 2.13 | The ordering and relationship of colours of Franciscus Aguilonius giving prominence to red, blue and yellow, and ordered by the lightness of the colours (from 1613), based on a diagram in Norman (1990). | 40 |
| 2.14 | Newton's sketch of his colour wheel. The size of each segment is related to its width in the spectrum. | 41 |
| 2.15 | Goethe's colour wheel with each colour positioned opposite its perceptual complement. | 41 |
| 2.16 | The arrangement of colours as a sphere by Runge (1810). The spherical arrangement is an idealisation, and is not derived from measurements of perception. | 43 |
| 2.17 | Munsell's illustrations from "A Color Notation" (Munsell, 1907) showing the placement of colours within his Color Tree. Left: an illustration showing the colours arranged taking into account lightness (vertically), and chroma (radially), and with the hues positioned to reflect the complementary relationships. Right: an illustration showing that, compared to red and green, pure yellow is lighter, while blue is darker. At the shown rotation, red has the maximum chroma (colour purity or intensity). | 45 |
| 2.18 | A colour wheel derived from Munsell's placement of hues, illustrating the five main colours, their interpolation, and their naming. Each of five primaries (red, yellow, green, blue and purple, can be subdivided to give colours between the primaries. The naming is from 1–10 of each primary or mixed pair of primaries, e.g. 10Y, 1GY, 2GY – 5GY – 10GY, 1G. | 46 |
| 2.19 | An indication of the variability in human colour perception. Overlaid on the Munsell hue circle are coloured bands indicating the range of colours perceived as the pure primary colour. Derived from Kuehni (2004). | 48 |
| 2.20 | The NCS colour system uses forty experimentally derived hues, with the opponent colours having prominent places. | 51 |
| 2.21 | A vertical slice of the NCS colour system gives the colours of a single hue (Y90R) mixed with black and white. The full saturation colour is always at the same vertical position, irrespective of its lightness. | 51 |
| 2.22 | Examples of different types of colour schemes that can be derived using geometric shapes and a colour wheel. The monochromatic scheme is not viable unless, as shown, the colour wheel is modified to vary lightness or saturation (or both) along the radial axis. | 55 |
| 2.23 | An circular arrangement of colours modelled on a colour wheel that shows variations of saturation, lightness and hue, from http://www.visibone.com/color/poster2x.html , accessed May 21, 2010. | 56 |

| | | |
|------|--|----|
| 2.24 | An alternate expanded 2D visualisation featuring perceptually uniform spacing, from http://www.visibone.com/color/kilochart%5F900.jpg , accessed May 21, 2010. | 56 |
| 2.25 | The HSL (hue-saturation-lightness) colour space is biconical with the full saturation colours aligned with mid-grey. | 57 |
| 2.26 | The HSB (hue-saturation-brightness) space aligns the full saturation hue with white. | 57 |
| 2.27 | A 3D view of the hue-saturation-lightness colour space, with the cutaway section allowing a view into the interior of the space (from software by the author, inspired by a diagram in Varley (1980)). | 58 |
| 2.28 | The use of complementary colours for text and its background can impair readability. | 63 |
| 2.29 | Two colour selectors typical of those in current use, both from Microsoft <i>PowerPoint 2003</i> TM . On the left, a flattened HSL representation, and the right, a form resembling a colour wheel. Neither supports selecting sets of colours, and both are modal, forcing the choice of colours out of context. It is very common for repeated invocations of either dialog to be necessary, as the new colour (visible in the swatch at bottom right of the dialog) can appear different when viewed in context, the visual impression being affected by the area of the colour and by nearby coloured objects. | 66 |
| 2.30 | One of the non-modal colour selectors available in the InkScape drawing tool. The combined colour wheel and hue slice colour selector on the right enables all tints and shades of the selected hue to be seen simultaneously. Dragging the little circle (the selected colour) within the triangular area immediately updates the colour of the selected object. The colours are therefore seen in context, which helps to avoid colour selection errors. | 67 |
| 2.31 | Two colour selectors from the Adobe <i>Illustrator CS4</i> TM package illustrating the features offered by professional graphical design tools. On the left, the contents of the thin vertical rainbow-coloured area can be set to show all variations of a particular colour property. The selector on the left displays all the hues in the vertical strip, with the larger area showing the other two colour space dimensions (for the example shown, lightness vertically and saturation horizontally). On the right, an illustration of the palette management facilities. The multi-coloured grid on the web page can be instantly recoloured by clicking on any of the palette sets. | 68 |

| | | |
|------|---|-----|
| 2.32 | Left and centre: the template-oriented approach to colour scheme management from Microsoft's web design tool <i>Frontpage 2003</i> TM . The facilities in <i>PowerPoint 2003</i> TM are similar. Complete schemes can be changed or recoloured in toto, but only with the indicated coloured elements. Right: the facilities offered to GUI application developers in Microsoft <i>Visual Studio Express 2008</i> TM . The options, unless altered programmatically, are limited to choosing from the limited palette shown. Microsoft products are illustrated as they are very widely used, and their colour selection facilities are typical of those found in products from other manufacturers. | 69 |
| 2.33 | A full colour image. | 87 |
| 2.34 | The same image with all colours forced to the same lightness. | 87 |
| 3.1 | A wireframe tilted at 45° to the vertical lightness axis (white at the top, black at the bottom) will allow the colours along the wireframe path to include colours with varying lightness (from the different vertical position) and, as can be seen from the figure on the right, varying saturations from the differing radial distances. | 97 |
| 3.2 | The Chromotome: a cutaway view of the HSL colour space showing the internal colours, illustrating how the placement of a wireframe with the dots (beads) on the wireframe indicating positions in the colour space, which correspond to sets of colours. The wireframe lies on a plane that may be tilted to vary the lightness of some beads, and the size of the wireframe itself (indicated by the circle on the plane) may be altered to vary saturation. The six circles outlining the beads and the two arrows indicating control handles have been added to the image for clarity in this thesis. Software by the author. | 99 |
| 3.3 | Data from the user interface is extracted, and used to allow the developer to specify groupings and constraints between the colourings of the interface elements. | 100 |
| 3.4 | The data from the interface characterisation is sufficient to allow the construction of colour molecules (each representing a colour scheme) in an abstract colour space. | 103 |
| 3.5 | The initial colour schemes are differing orderings of fixed wireframe spacings that can be considered as candidates or prototype colour molecules. These are optimised to create abstract colour molecules representing colour schemes that are harmonious, readable, and incorporate the design characteristics and constraints specified by the developer. | 105 |

-
- 3.6 The proposed colour scheme design cycle, from the extraction of data from the user interface and specification of constraints by the developer, through to the creation of an optimised colour scheme, the colours of which are used to recolour the interface elements. The developer can holistically alter the abstract-to-real colour mapping to change the overall scheme colours to get the most pleasing result. Also illustrated is the flow of data within the Colour Harmoniser architecture. The dotted lines (6 & 7) are external: 6 is the user looking at and evaluating the displayed colour scheme, and 7 is the user adjusting the personalisation controls to holistically alter the scheme, which will cause the displayed colours to be immediately updated. 107
- 4.1 A prototypical web page. 115
- 4.2 The left hand panel shows the interface element names extracted automatically from the web interface (fig. 4.1). Multiple elements of the same type are placed into an automatically created group (named using the underlying class names, which unfortunately are not usually self-explanatory). In the interface, the textual elements are in several different containers, but the automatic classifier groups them together, e.g. “HeaderTextColor” and “HomeButtonTextColor”, which are items of text have been placed together into the automatically created group “TLabelTextColor”, and items of class “TLabelColour” (the rectangular backgrounds to the text) have also been grouped. Unfortunately, the structure of the automatically created groups does not match the visual structure of the web page, but this can be remedied by renaming the groups and regrouping elements (fig. 4.3). 116
- 4.3 The user indicates those interface elements whose colours are part of the colour molecule by dragging them from the left-most panel to the centre, where they can be grouped into a hierarchical structure and given more meaningful names. The Y-shaped icon with yellow background indicates the colour scheme wireframe (“Y” denotes split-complementary), and the octagonal black-edged icons indicate “same coloured” groups. The right-most panel mirrors the centre panel. Its use is explained in figure 4.5. 118
- 4.4 A popup menu can be used to create new groups, to rename elements, and to specify the wireframe to be used. If desired, the black-white axis can be included as part of the wireframe. 118

-
- 4.5 The groups can be collapsed to hide unnecessary detail. Distinguishability constraints can be defined by first selecting an item in the centre panel (e.g. *Button Colour*), then right-clicking on an item on the right-hand panel (e.g. *LeftNavBar Background*) and selecting the appropriate constraint from the popup-menu. As shown above, the “is-distinct-from” is selected, meaning the *Button Colour* should be coloured differently from the *LeftNavBar Background*. Any constraint between the two selected items is also shown in the right-hand panel: the “≠” icon indicates “is different from”, “=” is used to indicate the elements have the same colour (are part of a group), and no icon means there is no colouring constraint between the selected elements. Instead of using the popup menu, the radio buttons below the centre and right panels can also be used. 119
- 4.6 The constraints shown will force the colour of a button to be different from the header/footer background, the left navigation panel, and the main body background. The “fR” (in the right-hand panel) indicates an automatically defined “forced” readability constraint between the button and the button text. 120
- 4.7 Once the interface has been characterised, the interface colouring is determined from the properties of the “top level items” – the children of the “Colour Harmonised” parent. To ensure constraints are handled correctly, distinctions between any children (e.g. footer-text and footer-background) in different groups require the constraint be also applied between the user-defined groups (Header/Footer Text and Header/Footer Background). 121
- 4.8 Two initial arrangements of colour atoms on different wireframes. On the left, a complementary wireframe with the items equally spaced. On the right, a split-complementary wireframe extended to include the black-white axis as part of the wireframe. 123

- 4.9 The colour scheme wireframes supported by the Colour Harmoniser prototype. The top row shows the wireframes for most frequently used colour schemes: complementary, split-complementary, and monochromatic (where only one hue is used). The monochromatic wireframe includes a pragmatic optimisation: it stops short of the black-white axis, so that the lightest and darkest elements contain a touch of colour, resulting in a more sophisticated looking scheme. On the bottom-left is the elliptical path suggested by Munsell. The bottom-centre and bottom-right figures show the analogous wireframe used in the prototype. The analogous scheme is often shown as an arc on a full saturation colour wheel, but this is too restrictive for use in the prototype: changes in lightness are obtained by tilting the wireframe from the horizontal; changes in saturation by using a different radius for the top and the bottom of the wire, as shown in the bottom-centre figure; and changes in hue by tilting the wire from the vertical, which changes the rotational angle for the different atoms, as shown in the bottom-right figure. 125
- 4.10 The penalty function applied to the overall scores for both distinguishability and readability to ensure all items are easily readable or distinguishable 134
- 4.11 The top four images show colour schemes from a single population before optimisation: some molecules have the atoms spaced evenly on the wireframe (top row), and some have the colour atoms positioned randomly in the colour space (middle row). As can be seen the schemes have different mixes of readability, distinguishable elements, and pleasing colour schemes, and some are more usable than others. This is reflected in their fitness function scores (out of 4). After optimisation, as shown by the schemes in the bottom row, readability has improved, all the items are distinct, and the saturation of colours is appropriate to their area. . . . 145
- 4.12 The ten colour schemes resulting from optimising ten different populations, all using hues defined by the current wireframe rotation angle. The images illustrate the diversity of schemes possible using the same underlying colours. The basis for all of these colour schemes is a split-complementary wireframe augmented to allow the positioning of items on the black-white axis. All would be suitable for online use. 146
- 4.13 Two views of the sRGB gamut in the perceptually uniform CIELAB space. As can be seen, there is a significant mismatch between the shapes of the spherical abstract colour space and the shape of sRGB projection. 150

-
- 4.14 Only colours within the dotted sphere are within gamut at all possible orientations of the wireframe. If it was strictly required that all the colour transformed from the abstract colour space be “correct” at all orientations of the wireframe, only those colours within the sphere (most of which are quite desaturated) could occur in the resulting colour schemes. Sacrificing the wide range of colours outside the sphere is too great a penalty for the use of a small fixed scaling factor as an abstract-to-real transformation method to be seriously considered. 151
- 4.15 Colours inside the sRGB gamut for an LAB Scaling factor of 56. Only colours inside the dotted sphere can be present in the real colour scheme. This still excludes some high saturation colours, many of these can be included if the developer uses the personalisation controls (to be discussed later) to increase the colour intensity (saturation) of the recoloured interface. This will increase the abstract-to-LAB scaling factor which will increase the overall saturation of the scheme, but may force more colours out of gamut, which will require approximations to be made. 154
- 4.16 From top left, the different colour schemes resulting from rotating a molecule in 30° steps about the lightness axis as it is projected in the CIELAB space. 157
- 4.17 Varying the LAB scaling factor as the abstract colour scheme is projected in CIELAB space can act as a saturation control, enabling smooth control of the overall intensity of the resultant scheme. The six upper images show a monochromatic colour scheme of varying saturation, with LAB Scaling factors of 15, 30, 45, 60, 90, 150. The lowest scheme has a LAB Scaling factor of 0. This removes all colour, without impairing usability. 159
- 4.18 Each row shows the same colour scheme, but one of the pair has the lightness inverted, which changes the colours without affecting the colour relationships. The schemes are based on a split-complementary wireframe including the black-white axis. The top six schemes have a LAB scaling factor of 56, the lower four ≈ 145 160
- 4.19 Desktop application colour schemes created by the Harmoniser prototype. Extending the wireframe model to include the black-white axis (e.g. the bottom scheme) allows the use of less colour, and gives more familiar schemes. 164

| | | |
|------|---|-----|
| 4.20 | Colour Harmoniser-adjusted colour schemes for presentation slides. Almost all the area ($\sim 90\%$) of this interface is background), so schemes that balance around mid-grey (e.g. complementary/split-complementary/elliptical) can only balance by placing the background very close to the origin (and it is therefore greyish) to minimise its effect. If the “use black-white” option is enabled (e.g. the bottom-left scheme), the background can be placed on the achromatic axis, and the smaller objects balanced against one another. It would seem that using area balance to determine the lightness and saturation of the elements is better suited to interfaces whose elements have a more even distribution of areas (or can be made to have this by grouping). | 165 |
| 4.21 | The data flow within the prototype Colour Harmoniser, starting with extracting data from user interface to be coloured and the user’s constraints, through to the updating of the same interface with an optimised colour scheme. The solid lines indicate data flows within the software, the dotted lines are user inputs. | 166 |
| 5.1 | The colour schemes to be ranked by the participants were chosen from six different bands of fitness. The lowest scoring colour schemes in a population are used as the $rank_1$ sample. The $rank_6$ scheme was the best at the end of the optimisation, with a score of at least 3.5, with the remaining four being the first schemes during the optimisation to exceed the fixed fitness thresholds of 1.5, 2, 2.5 & 3. The darkness of the shading indicates the region more likely to have samples in each rank band. | 177 |
| 5.2 | The distribution of Spearman rank correlation coefficient scores between rankings of human evaluators and ranking derived from the fitness function score, for all sets of colour schemes. | 179 |
| 5.3 | The possible range of fitness function values is partitioned in equally size bands and gaps. The colour schemes for $rank_1, rank_2$ etc. were those whose fitness fell in the rank bands. Schemes whose fitness fell in the inter-band gaps were not used. | 184 |
| 5.4 | The distribution of Spearman correlation coefficient scores (ρ) between rankings of colour schemes having stratified levels of fitness by human evaluators and the ranking derived from the fitness function score, for the all randomly generated colour schemes in data set Main Fitness Function Validation trial. | 186 |
| 5.5 | The histogram of fitness function scores of one million colour scheme generated by a Monte Carlo method, using a split-complementary wireframe, with a wireframe affinity of 0.9. Even with such a high wireframe affinity, very few ($\sim 0.3\%$) acceptable interface colour schemes result from random colourings. | 189 |

| | | |
|-----|---|-----|
| 5.6 | A scatterplot between the scores of each fitness function terms for randomly generated schemes having a high fitness function score. To get high scoring schemes, a forced wireframe affinity was incorporated into the random scheme generation. | 190 |
| 5.7 | The summary statistics for the fitness function scores plotted against user-assigned image rank for the colour schemes in data set <i>Main Fitness Function Validation Experiment</i> . The darker line inside each box indicates the median fitness score and outlier values are shown by the small circles. | 191 |
| 6.1 | Manual colouring trial: this brightly coloured instruction page to the “one-colour-at-a-time” experimental interface shows the colourable regions of the web page, and the simple four-step instructions. | 202 |
| 6.2 | Manual colouring trial: the test interface as initially seen is achromatic. Clicking on any web page element will cause a colour selector to appear. The ”Show Help” will cause the previous instruction page to reappear. | 203 |
| 6.3 | Manual colouring trial: clicking on a page element (the header is shown here) displays the standard <i>Windows</i> TM colour selection dialog, expanded to show the RGB and HSL options, not just the basic colours. Whatever colour is selected will become the new colour of the clicked-upon element. | 203 |
| 6.4 | Manual colouring trial: the body text and button text are difficult to click on precisely. Therefore, for these elements, the interaction has an additional step: clicking anywhere on body or on a button will cause a popup-menu to appear. When either the text or background is chosen from this menu, the colour dialog is displayed. The additional step ensures the item to be recoloured is the intended one. | 204 |
| 6.5 | Manual colouring trial: the participant can change element colours as often as desired until a satisfactory scheme has been found. Clicking “Next” will move to the post-colouring questionnaire. | 204 |
| 6.6 | Harmoniser-based colouring: the introductory page showing the experimental interface the participants used to choose Harmoniser-generated schemes and invoke the adjustment interface. | 206 |
| 6.7 | Harmoniser-based colouring: the initial view displays a set of randomly selected Colour Harmoniser-created schemes as “starter” schemes for the user to adjust. The user can drag any thumbnailed scheme to the lower A or B areas to have the scheme enlarged, so subtle colour differences can be seen more clearly. | 207 |
| 6.8 | Harmoniser-based colouring: Dragging a thumbnail to the lower display areas will enable the “modify colour scheme” button, allowing the colour scheme to be modified. The A & B show the currently enlarged thumbnails. | 207 |

| | | |
|------|--|-----|
| 6.9 | Harmoniser-based colouring: Using the holistic adjustment controls, users can alter a scheme by varying all the colours (wireframe rotation), altering the colour intensity (saturation) or swapping the light and dark colours. The non-modal interface immediately updates the colours as soon as any control is altered. The mid-grey background minimises simultaneous contrast effects. | 208 |
| 6.10 | Harmoniser-based colouring: The holistic colour scheme adjustment interface showing the same colour scheme as the previous image, but with the lightness inverted. | 208 |
| 6.11 | Harmoniser-based colouring: Showing the previous colour scheme after varying the colours and increasing the saturation. Clicking the indicated button to finish the colouring closes this window, and, after updating the colours of the enlarged image and its related thumbnail, redisplay the main interface. | 209 |
| 6.12 | Harmoniser-based colouring: On return from the adjustment interface, the selected scheme and its thumbnail are updated to the colours chosen in the colour-change interface. Different thumbnails can be dragged to the lower A/B positions to explore other colouring and to see subtle colour differences more clearly. Clicking one of the “Choose A/B as best” buttons ends the colouring and displays the post-colouring questionnaire. | 210 |
| 6.13 | Manual colouring trial: having created their colour scheme, the user is shown the interface and asked to evaluate the process and the result. . . | 211 |
| 6.14 | Harmoniser-based colouring: having adjusted a Harmoniser-based scheme, the user asked for their opinion on the process and the result. | 211 |
| 6.15 | Having created two schemes, both are displayed and the user is asked to select the one with the more professional colour scheme. | 212 |
| 6.16 | In the concluding questionnaire, the user is asked to compare both methods of colour scheme creation. | 213 |
| 6.17 | Rating the final colour schemes, one prepared by choosing individual colours for interface elements, the other by holistically adjusting a Harmoniser-based scheme. | 213 |
| 6.18 | A histogram of the completion times for creating a colour scheme manually (left) and optimising a Harmoniser-based scheme (right). Each bar is one minute. Most participants took five minutes or less whichever method was used. | 214 |
| 6.19 | Responses to “would it be quicker to find a professional looking scheme using the manual or “all-colour-at-once” (Harmoniser-based) selector?” . | 215 |
| 6.20 | The responses to a question asking whether it would be easier to find good colours schemes using the Harmoniser method. It is clear that the majority of respondents think that finding good schemes would be easier using the Harmoniser method than by choosing colours individually. . . | 216 |

| | | |
|------|--|-----|
| 6.21 | Which colour scheme is more professional? – the responses after the participants had created a colour scheme using each of the manual and Harmoniser-based methods. The difference between the assessments of the resulting schemes is not statistically significant, in contrast to the results shown in figure 6.19 | 217 |
| 6.22 | This chart shows the degree to which the users found the effect of the Harmoniser-based holistic colouring controls understandable. | 218 |
| 6.23 | This chart shows the degree to which the users found that the Harmoniser-based schemes could be improved using the holistic controls. | 218 |
| 6.24 | The flip-light-dark option inverts the lightness of colours in Harmoniser-based colour schemes. The responses indicate the degree to which the control were perceived to be beneficial. | 219 |
| 6.25 | The colour scheme evaluation page used in the “Compare the Results” experiment. | 223 |
| 6.26 | The participants’ evaluation of how professional (upper) and how artistic (lower) the colour schemes were for each creation method. The categories within each plot indicate increasing quality from left-to-right: <i>unskilled</i> to <i>very professional</i> (top row), and <i>not-at-all</i> to <i>very artistic</i> (bottom row). From left to right, the columns of plots indicate schemes coloured: randomly, by a human non-artist, human artist, Harmoniser-raw, and Harmoniser-adjusted. There is a clear improvement in quality from left to right. An enlarged version of these images is included in Appendix B.1, p260. | 225 |
| 6.27 | The times required to rate a colour scheme for all four evaluation criteria, with the data trimmed to exclude values greater than sixty seconds. This indicates that most users took only a few seconds to score each criterion. | 235 |
| B.1 | The participants’ evaluation of how-professional (upper) and how-artistic (lower) the colour schemes were for each creation method (from left to right): randomly, by a human non-artist, human artist, Harmoniser-unadjusted, and Harmoniser-adjusted. | 260 |

List of Tables

| | | |
|-----|---|-----|
| 2.1 | The incidence of colour vision deficiencies in Western races, derived from Fairchild (1998). | 29 |
| 5.1 | Three sets of experiments were used to evaluate the viability of the Colour Harmoniser approach to colour scheme design. The first set (shown in this table) are oriented towards validating and adjusting the parameters of the fitness function. The next two (Table 5.2) assess the usability of the Harmoniser method and the quality of its schemes. . . . | 168 |
| 5.2 | An overview of the second and third experiments used to evaluate the viability of the Colour Harmoniser approach to colour scheme design. The second experiment assesses the usability of the Harmoniser method of colour scheme selection and adjustment in comparison to more conventional methods of colour scheme design, while the final experiment compares the quality of colour schemes produced by the Colour Harmoniser (both before and after adjustment), with schemes created by people. | 169 |
| 5.3 | A list of the colourable elements on the prototypical web page (fig. 4.1, p115). The middle column shows the total number of colours if each element can be coloured differently. The right-most column shows the effect of grouping related elements and using the same colour for the elements of a group. The number of colourable elements drops significantly, and the use of common colours for multiple interface elements also increases the visual coherency of the scheme. | 170 |
| 5.4 | A breakdown of the interface elements found in a survey of fifty personal web sites by Wu (2000). The elements in the prototypical web page used in the experiments are indicated with an asterisk: all the elements with essential semantic content are included. | 172 |
| 5.5 | Correlation (human vs fitness function) for the images within each image set. | 180 |
| 5.6 | Overall Spearman correlation (ρ) between the human rankings of colour schemes within an image set and those of the fitness function, grouped by subject. | 180 |
| 5.7 | The Spearman rank correlations between ranking a colour scheme by human evaluators and the rank derived from the fitness function score. . | 181 |

| | | |
|------|---|-----|
| 5.8 | The Kendall-W coefficient of concordance assessing the degree of agreement between participants on the ranking of colour schemes within each set. The significance values are derived from a table lookup that does not include values of p less than 0.01. | 181 |
| 5.9 | The correlations between the four fitness function terms for randomly generated schemes with high overall fitness (method: Pearson's product moment correlation) | 188 |
| 5.10 | The effect on colour schemes resulting from the omission of each fitness function term. | 193 |
| 5.11 | The estimated term weights calculated from a regression of fitness function term scores against the user-assessed ranks two to six. | 194 |
| 5.12 | Optimised fitness function weights estimated from colour schemes ranked by evaluators at rank 2 and greater. Note: distinguishability weight is not determined from the regression, but is set equal to the lowest of the other three terms. | 195 |
| 6.1 | The demographic data from the colour scheme evaluations in the "Compare the Results" experiment. | 224 |
| 6.2 | The evaluation scores of the different methods of creating colour schemes. The participants rated the schemes using a five point scale. In the table, the scores range from 1 (worst) to 5 (best). Notable are the surprising placement of Harmoniser-based schemes above the humanly-created schemes, including those produced by artists, the improvement in scores caused by adjusting Harmoniser-based schemes, the consistency of ordering of scores for the image creation methods, and the poor scores of randomly-coloured schemes. | 226 |
| 6.3 | The significance of the differences between the mean scores of colour schemes creating using different methods, estimated using a t-test (df~700). In all cases, the mean scores orderings are: <i>random < non-artist < artist ≤ CH-Raw < CH-Adjusted</i> . Therefore, if there is a significant difference between two methods, e.g. CH-Raw and CH-Adj, it means that CH-Raw < CH-Adj, or equivalently, CH-Adj schemes scored more highly and are therefore better. | 227 |
| 6.4 | The significance of the differences between the means of the Harmoniser-created schemes <i>after human adjustment</i> and all other methods, for the four evaluation criteria. Orderings of the means: <i>random < non-artist < artist ≤ Harmoniser-Raw < Harmoniser-Adjusted</i> | 228 |
| 6.5 | Effects of demographics and past experience on the evaluations of the <i>Adjusted Harmoniser-based schemes</i> . No effects were apparent except for gender: females gave higher scores than males, across all four evaluation criteria. | 228 |
| 6.6 | The details of the gender differences in scoring the Harmoniser-adjusted colour schemes. For all but website suitability, females gave significantly higher scores. | 229 |

| | | |
|------|---|-----|
| 6.7 | The significance of the difference between the means between the Harmoniser-created schemes <i>before any human adjustment</i> and all other methods, for all four evaluation criteria (two-tailed t-test). Orderings of the means: $random < non\text{-}artist < artist \leq Harmoniser\text{-}Raw < Harmoniser\text{-}Adjusted$ | 229 |
| 6.8 | The relative placement and grouping of colour schemes created using the different methods. After human adjustment, the Colour Harmoniser-based schemes are rated more highly than all other methods except for the artistic criterion, where they are judged as comparable to those of artists and Harmoniser-raw schemes. The Harmoniser-raw schemes are comparable to those created by humans, both artists and non-artists. | 230 |
| 6.9 | Harmoniser-raw schemes: the effects of evaluator demographics and past experience on the evaluations. Unlike the Harmoniser-adjusted schemes there is no gender effect, but there are effects relating to colour sense and age. | 230 |
| 6.10 | Harmoniser-raw schemes were scored more poorly by evaluators with excellent colour sense. The difference is statistically significant. | 231 |
| 6.11 | The significance of the difference between the means between the <i>schemes created by artists</i> and all other methods, for all four evaluation criteria (two-tailed t-test). Orderings of the means: $random < non\text{-}artist < artist \leq Harmoniser\text{-}Raw < Harmoniser\text{-}Adjusted$ | 232 |
| 6.12 | Schemes by artists: there are no statistically significant demographics effects, except for the colour-sense criterion. Those with (self-assessed) excellent colour sense appear to be able to differentiate something in the schemes created by artists (from anova: $EvaluationCriterionScore \sim ColourSense$). However, this difference is not (as might be suspected) to rate the artist's schemes more highly. Surprisingly, as can be seen in Table 6.13, schemes by artists are awarded lower scores. | 232 |
| 6.13 | Schemes by artists are scored more poorly by those with excellent colour sense than by those with average colour sense. Unlike lower scores for the Harmoniser-raw schemes, the lower scores for the schemes by artists are less easy to explain. | 233 |
| 6.14 | The schemes by artists are given significantly lower scores by those with excellent colour sense than by those with average colour sense. This difference only exists for the artists schemes and for the "How Artistic" criterion of the Harmoniser-raw schemes. | 233 |
| A.1 | The full t-test results found when comparing the mean scores for the different methods of creating colour schemes from the "Compare the Results" experiment. | 258 |