An investigation into perceptual hue-ordering

David Connah[1], Marina Bloj[2], and Graham D. Finlayson[1]; 1) Department of Computer Science, University of East Anglia, UK; 2) Bradford Optometry Colour and Lighting Lab, School of Life Sciences, University of Bradford, UK.

Abstract

There are a number of problems for which colour-vectors need to be ordered: vector morphology and colour-to-greyscale conversion are two important examples. A lexicographic ordering can perform this function, but while the luminance and saturation components have clearly defined scales with fixed origins there is no such scale for the hue dimension. In this paper we tackle this problem directly using a greyscale assignment paradigm. Observers were shown pairs of iso-luminant and iso-saturation colours, and asked to assign greyscale values to them such that the assigned greyscale difference matched their perceived colour difference. From the resulting magnitudes we construct a 1D scale of hue values using multi-dimensional scaling, and the results show a sinusoid-like ordering. We also analyse the signs of the differences (i.e. which colour was assigned a brighter greylevel), and find that different observers reliably use the same sign-assignment strategy. This result can be used to fix the phase of the sinusoidal scale resulting from the magnitude experiments. This result also supplements earlier work on the Helmholtz-Kohlrausch effect, but is not the same as we have uniquely used colors that were set as iso-luminant and iso-saturation individually for each observer. Here we visualise the effect by encoding iso-luminant and iso-saturation images with the derived greyscale ordering.

(a) A simple colour image

(b) Luminance of (a)

Figure 1. *A simple colour image and it's reproduction using the standard luminance transform.*

Introduction

Imagine the simple image in Figure 1(a); how might this image look if it were printed in greyscale? In order to maintain the basic content of the image the two sides should be assigned different greyscale values. In this case, however, the colours are iso-luminant; thus, the standard method of computing the luminance of the colours, and using this as a basis for the greyscale, results in Figure 1(b), where it is not possible to discriminate between the colours. Now, in order to assign a different grey value to each side we must make one side brighter than the other; but which should it be?

This is an ordering question: given a pair of colour vectors, can they be ordered to make one "greater than" the other? In this particular application "greater than" is synonymous with "brighter than". The ordering problem is also encountered in morphological processing, where the erosion and dilation operators are defined with relation to the greatest lower bound and lowest upper bound of a set of pixel values; clearly these quantities have no meaning without the notion of order. As a result, in order to apply morphological operators to colour images, some ordering principle must be devised.

In the field of colour morphology Comer and Delp [1] have suggested a simple *reduced* ordering strategy of projecting the 3D colour values onto a single 1D value; this can be the projection onto the first principal component of the image data, or a more perceptually relevant variable such as luminance. The problem with this approach is that there are still likely to be many boundaries in the image where adjacent colours have similar luminance values, and thus correct greyscale orderings cannot be devised [2]. The same approach is also employed by Socolinsky and Wolff [3] in colour-to greyscale conversion. Their algorithm uses a strong mathematical definition of contrast between colours, but by using the luminance image to perform the ordering the order becomes undefined at isoluminant colour boundaries.

A more involved approach, used extensively by Hanbury and Serra [4, 5], again with relation to morphological processing, is to use a lexicographic ordering (also called a *conditional* ordering). This approach is similar to that used in the dictionary to order words: ordering is first performed on a particular 1D variable, but if two colours are equal with respect to this variable, then they are compared based upon a second variable. A simple example could be to order based upon just the red-channel, but if the red value of two colours is equal, then order them based upon the green channel. Rather than use RGB space, Hanbury and Serra use more perceptually motivated colour-spaces such as HSV, HSL and CIELAB. Thus the primary ordering variable is luminance (or value), the second is saturation, and the third is hue, although they need not necessarily be applied in this order. The main problem with this approach is how to deal with the hue component. Hue is defined as an angle on the interval $[0°, 360°)$, or $[-180°, 180°)$; thus at some hue values there is an inevitable discontinuity, where two colours with almost identical hue values can be seen as polar opposites in terms of their order.

The lexicographic approach of Hanbury and Serra would also seem to be well suited to the colour-to-greyscale problem. Using luminance and saturation as the first two ordering variables seems logical since both are naturally scalable and have well-defined origins. Both also have a perceptual grounding in that increasing luminance and saturation tends to make colours seem brighter (as described by the Helmholtz-Kohlrausch effect [6], which states that more saturated colours will tend to appear brighter). The hue variable, however, does not have a natural origin and it is not clear how a scale should be built to represent it.

The concept of hue ordering, while not tackled explicitly by other colour-to-greyscale methods, is inherent in a number of them. Gooch *et al.* [7] discuss the ambiguity in hue-ordering, and employ a user-defined hue-angle parameter, where colours whose hue angle is close to the predefined angle are deemed larger than those far away. Grundland [8] *et al.*'s method, on the other hand, will tend to give colours of a certain hue brighter values if there are many other colours of the same hue that are relatively bright; i.e. the ordering is done on an image dependent basis.

In this paper we use a direct approach to investigate the problem of perceptual hue ordering. We devise an experiment where we ask observers to assign greyscale values to a pair of patches, such that the difference between them matches the difference between a pair of colour patches. Using multi-dimensional scaling we extract, from the magnitude data, a 1D scale of hue values. The distance between these scale values approximates, as closely as possible, the judged differences between the colours. In addition to this we analyse the signs assigned by observers, and find highly significant within- and between-observer agreement. Thus we find that colours of equal luminance and saturation can be reliably ordered using an approximately sinusoidal ordering, where the phase of the sinusoid can be fixed by the observers' responses.

Hue ordering experiment

The goal of the experiment is to derive a one-dimensional ordering of hues which has a perceptual grounding. The scale should, as well as possible, reflect the differences between the colours as well as maintaining brighter colours as higher values in the scale. To ensure that the scale is only related to the hue component of the colours, we use stimuli that are equal to one another in both their luminance and saturation. To generate such colours is more challenging than it may first appear, as individual differences between observers dictate that what may be isoluminant colours to one observer may not be for another [9].

As a starting point we generated 16 colours from the same chromaticity plane in CIE$L^*u^*v^*$ space, thus giving them equal luminance values. To ensure that the colours were also iso-saturation we sampled them from a circle of constant saturation, where the centre of the circle is given by the chromaticity of our monitor's white-point. To make sure that the colours were equally spaced in hue, we sampled them at equal angular increments from the iso-saturation circle.

Using the CIE $L^*u^*v^*$ space to define the colours ensures only that they will be approximately perceived as iso-luminant and iso-saturation. Now, to ensure that each observer perceived the colours to be iso-saturation and iso-luminant, we allowed each observer to adjust the 16 colours. We did this using two carefully controlled preliminary experiments; the first to enforce iso-luminance, and the second to ensure iso-saturation; see the Appendix for details of these experiments.

At the end of the preliminary adjustments each observer had a different set of 16 colours that were, as far as possible, iso-luminant and iso-saturation. In the main experiments that we describe here, only 8 of these colours were actually used and Figure 2 shows colours that represent one of the observer's settings [1].

Figure 2. *The 8 colours used in the experiment. These colours are referred to elsewhere in this paper, from left-to-right, as colours 1-8.*

In each trial a pair of square patches was displayed in the centre of the screen, where each member of the pair was filled with a different target colour. Below this pair of colours were placed two grey patches whose luminance was variable, but whose chromaticity was fixed at that of the monitor whitepoint; see Figure 3 for a representation of the experimental set-up [2]. The observers' task was then to view the pair of colour patches, and to adjust the luminance of the two grey patches until they represented the difference between the colour pair.

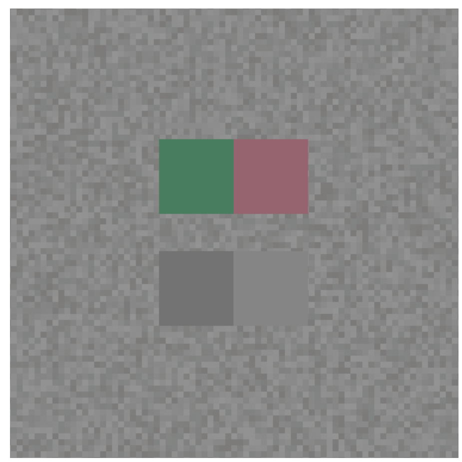

Figure 3. *A representation of the experimental set-up. The intensity of the grey patches was adjusted until they represented the difference between the colour pair. In the experiment each square patch measured 30mm along each side, and with a viewing distance of 1000mm subtended a visual angle of 1.7 degrees. The background filled the 270×380 mm monitor screen, and subtended a visual angle of approximately 20 degrees.*

At the start of each trial the grey patches were identical to one another, each having a luminance of 20 cd/m^2. Observers were able to adjust the level of each of the grey patches independently by using buttons in a response box. Adjustments were made in steps of 0.5 lightness units. Lightness is computed as CIE L^* value, and the whitepoint is that of the monitor. In each experimental run every pair of colours was shown twice; once for each left-right configuration. This made a total of 56 trials for each run. Each of our 4 observers carried out 3 experimental runs. Observers had normal colour vision and visual acuity and were naive to the purpose of the experiment. Given that each colour pair was

[1] The colours displayed here will not match exactly those shown using the calibrated colour monitor in our experiments

[2] The exact dimensions shown in the figure are not identical to those actually used

shown twice per run, this meant each pair was compared 6 times by each observer.

The experiments were performed in a dark room, and the stimuli were shown on a calibrated colour monitor. Prior to each run observers were shown an adaptation screen for 1 minute and afterimages where avoided by the presentation for 500 msec of a mask between each trial. The mask and the adaptation screen consisted of the same checkerboard pattern shown as the background in Figure 3 with a mean luminance of 20 cd/m^2 and a maximum luminance of 27 cd/m^2 and minimum of 13 cd/m^2.

Results and analysis

The data resulting from the experiment are a set of signed greyscale differences, with one greyscale difference per judgement. The magnitude of the difference is measured as the CIE L^* between the grey values, while the sign dictates which of the grey values was larger. First, we consider the magnitudes. From the experimental data we obtain Table 1. The i, jth entry of this matrix, denoted d_{ij}, gives the average difference between colour i and colour j, pooled over all observers and repetitions.

Colour	1	2	3	4	5	6	7	8
1	0.0	2.3	3.6	3.3	3.5	4.1	3.5	2.4
2	2.3	0.0	2.3	3.1	4.0	3.9	4.0	3.1
3	3.6	2.3	0.0	2.8	3.5	4.2	3.9	3.7
4	3.3	3.1	2.8	0.0	2.5	2.9	3.3	3.5
5	3.5	4.0	3.5	2.5	0.0	1.9	3.0	3.3
6	4.1	3.9	4.2	2.9	1.9	0.0	2.7	3.6
7	3.5	4.0	3.9	3.3	3.0	2.7	0.0	2.7
8	2.4	3.1	3.7	3.5	3.3	3.6	2.7	0.0

A table representing the magnitude results. The i, jth entry in this matrix corresponds to the average difference between the colour in row i and that in column j judged by the observers. The units are unsigned CIE L^* differences.

Given this matrix of differences between the colours, is it possible to arrange the colours into an ordered 1D scale? To answer this question we use Multi-Dimensional Scaling (MDS). The goal of MDS, in the context of our work, is to assign a scale value to every colour such that the distances between the scale values represent the recorded greyscale contrast assigned to that colour pair. More specifically, MDS assigns scalar values a_i and a_j to colours i and j, such that the difference between them, $x_{ij} = \sqrt{(a_i - a_j)^2}$, minimises the so-called $stress$ value, given by:

$$stress = \sqrt{\frac{\sum\sum\left(f(x_{ij}) - d_{ij}\right)^2}{\sum\sum d_{ij}}}. \quad (1)$$

There are two major variations of MDS: metric and non-metric scaling. Metric scaling aims to maintain the Euclidean distance between any pair of points; i.e. the function $f(\cdot)$ in equation 1 is the identity function. Non-metric scaling ignores the size of the differences and only looks at their rank order. In metric scaling $f(\cdot)$ can be any monotonically increasing function.

We apply both metric and non-metric scaling to the pooled distance matrix in Table 1. The result of metric scaling, which gave a $stress$ score of 0.3985, is given in Figure 4; the non-metric results, associated with a $stress$ score of 0.2394, are given in Figure 5. In these Figures indices 1-8 for the experimental colours (see Figure 2) are plotted against the scale values. Superficially the two scales seem to be very different; with the metric MDS assigning positive scale values to the second and third colours, while non-metric MDS assigns negative values. This is not important here, however, since the scales are derived from unsigned differences, reflecting the scale values in the x-axis gives an equally good scale. We can see that the metric ordering gives a scale that looks very similar to a sinusoid, especially when we consider that the hue-values "wrap-around" on the x-axis. The non-metric scaling, on the other hand, produces a less smooth function, whose extrema occur at slightly different points.

A sinusoidal ordering is perhaps a natural ordering for the colours we have used. We have sampled points on a circle in an isoluminant plane, and then assessed the colour differences between these points (something we might have done through simulation alone, by filling the distance matrix with Euclidean distances calculated in CIE$L^*u^*v^*$ space). By forcing a 1D ordering it may be that we are projecting the points around a circle onto a line, which naturally produces scale-values that form a sinusoid.It is not clear, however, what the phase of this sinusoid should be; i.e. whether there is a natural origin for the hue-circle. If there is such an origin, then it may be provided by examining the sign of the greyscale differences established by our observers.

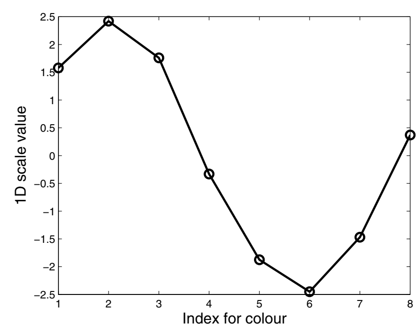

__Figure 4.__ Scale values derived using metric MDS. The x-axis represents the 8 different experimental colours (see Figure 2), while the y-axis shows the scale-value assigned to that colour.

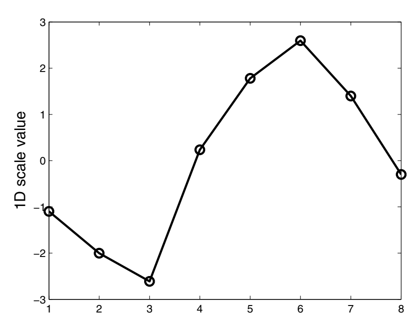

__Figure 5.__ Scale values derived using non-metric MDS. The x-axis represents the 8 different experimental colours (see Figure 2), while the y-axis shows the scale-value assigned to that colour.

Sign assignment

In the second part of our analysis we look at the sign information, i.e. which colour was made brighter and which darker. To analyse this data we begin by tabulating the results. Specifically, each time colour i is judged brighter than colour j we increment the ijth entry of the data matrix shown in Table 2, which we will denote p. With 4 observers, each repeating all judgements 6 times, the maximum entry in cell p_{ij} is 24, and $p_{ji} = 24 - p_{ij}$. If the observers are assigning signs randomly, then we expect all the entries in this matrix to be at, or close to, 12. This is clearly not the case, and there is a considerable amount of agreement; e.g. it is reliably agreed that colour 8 should be assigned a brighter greylevel than colour 1.

Colour	1	2	3	4	5	6	7	8
1	0	20	20	13	5	4	2	2
2	4	0	18	4	2	2	5	2
3	4	6	0	5	3	3	5	3
4	11	20	19	0	5	2	4	8
5	19	22	21	19	0	5	12	11
6	20	22	21	22	19	0	14	18
7	22	19	19	20	12	10	0	19
8	22	22	21	16	13	6	5	0

A table representing the sign-assignment results. The entry in cell i, j indicates how often a colour in row i is judged brighter than a colour in row j.

It is possible to quantify the agreement between observers using Kendall's coefficient of agreement[10]. To describe this statistic we firstly calculate the total number of agreements between pairs, Σ, which is given by:

$$\Sigma = \frac{1}{2} \sum_{i \neq j} p_{ij} \left(p_{ij} - 1 \right) \tag{2}$$

Kendall's coefficient of agreement u can now be defined as:

$$u = \frac{8\Sigma}{N(N-1)T(T-1)} - 1 \tag{3}$$

where N is the number of colours used and $T = 24$ is the number of observations. The closer the value of u to 1, the greater the agreement between observations. Our results show, for the pooled data, a u value of 0.367. We test the significance of this statistic, i.e. test the null hypothesis that observers do not agree with one another, by applying the χ^2 statistic [11]. This shows that a u value of 0.367 is highly significant; the probability that we should accept the null hypothesis of no agreement is less than 0.0001 (on a probability scale of 0 to 1).

As well as agreement, we can also examine observer *consistency*. The idea of consistency can be captured by a simple example: if we have three colours, A, B and C, and we compare each colour with each of the others, then the judgements are consistent if, for example A is brighter than B, B is brighter than C and A is brighter than C. This consistency can be violated if, for example, C is brighter than A. Kendall and Babington-Smith [10] provide a statistic, the coefficient of consistency Ω, to measure this property. We omit the details here, but a value Ω of 1 indicates perfect consistency, in which case the methods can be ordered reliably. In our experiments we measure the average value of Ω over all observations and find a value of 0.7875, which is high, and suggests that the colours can be ordered reliably.

An idea of the ordering that might be produced from the sign-data is given in Figure 6. Here we plot the number of times a colour is judged "brighter" as a proportion of all judgements; this scale clearly shows that colour 3 tends to be judged brighter much less often than colour 6. The similarity of this Figure to Figure 5 is clear, and the knowledge that the third colour is repeatedly assigned as "darker" suggests that we can accept the sign of the ordering given by the non-metric MDS. This result supports findings from earlier work on the hue-dependency of the Helmholtz-Kohlrausch effect (see, e.g. [12]), which finds, by direct brightness (or lightness matches), that some hues tend to appear brighter than others when saturation and lightness are kept approximately constant.

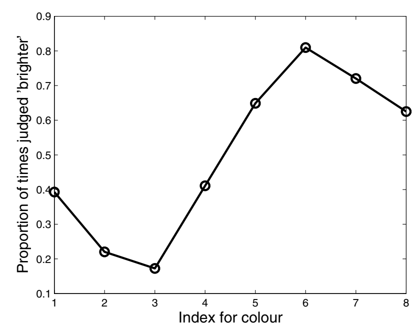

Figure 6. *A graphical representation of the sign assignments. The x-axis shows the index for each colour (see Figure 2), while the y-axis shows the number of times that colour was assigned a brighter greyscale value than another colour (expressed as a percentage of all 168 judgements).*

Visualising the ordering

In order to visualise the ordering we make use of synthetically generated iso-luminant and iso-saturation images, shown in the first column of Figure 7. To generate these images we took a set of standard RGB images and transformed them to CIE $L^*u^*v^*$ space by assuming the images were coded as sRGB [13]. We then set the L^* value to 50 for all pixels, and transformed the u^* and v^* components to lie on a circle centered on the D65 white-point, but with their hue-angle unchanged. Since hue is undefined for achromatic colours we arbitrarily assigned them an angle of 0 degrees. The images were then transformed to sRGB for display.

The hue-angle is measured on the range $[0, 360)$, where a value of 0 defines the line parallel to the positive u^* axis. In the central column of Figure 7 the hue-angle is mapped linearly onto the range $[0, 1]$ by simple scaling. In the right-hand column the scale values from the non-metric MDS are used. The hue-angle values of the 8 experimental colours are mapped directly to their respective scale-values. Intermediate hue-angles are mapped to a scale-value using cubic spline interpolation. All greyscale renderings are raised to a gamma for display.

The renderings from the MDS scale values have much lower contrast than the hue-angle since we take the magnitude of the contrast from the scale values directly. They are closer, however, to the original contrasts. In the central column we also see some cases where similar hues are assigned polar opposite (black and white) greyscale values; e.g. the two rightmost hats in the second row. This is due to the discontinuity at 0/360 degrees on the hue-angle scale. The MDS scale values are continuous, and so do not produce these effects.

Conclusions

In this paper we have investigated a perceptual approach to hue-ordering using a greyscale assignment paradigm, where observers make a greyscale settings for a pair of colour simultaneously. We have analysed both the distances between the colours, and which colours tend to be assigned brighter greyscale values; the results suggest that hue values can be ordered on a natural, perceptual, scale. The ordering is not total, in that some colours are assigned similar scale values even when the difference between them is judged to be large; this can be probed further by using more experimental colours. However, while not total, the ordering is not arbitrary: some hues are consistently assigned dark greyscale values, while others are assigned bright greyscale values; a result which agrees with earlier work using a direct brightness matching to assess brightness of colours using different hues.

Appendix: Methodology

Luminance adjustment

The luminance of a stimulus is not only observer-dependent, but is also determined by an observer's state of adaptation. In order to take account of this, we adjusted colours in the Derrington-Krauskopf-Lennie (DKL) physiological space [14], which explicitly incorporates the state of adaptation in its definition of luminance. The DKL co-ordinates are arrived at by firstly transforming the CIE $L^*u^*v^*$ values to an LMS based [15] cone-contrast space, where colours are defined in terms of their contrast from the surround. In all our experiments the background was a checkerboard of small grey patches that had a mean luminance of 20 cd/m^2 , each individual square being assigned a random deviation from this mean. Thus we defined the average surround as the grey with a luminance of 20 cd/m^2, and chromaticity of the monitor white point.

From cone-contrast space, the colours are transformed to DKL space, whose three axis represent luminance contrast, red-green contrast (L-M) and blue-yellow contrast (S-(L+M)). Since the state of adaptation is incorporated into DKL space, the experimental colours no-longer fall on an iso-luminant plane in this space. Thus we set the DKL luminance co-ordinates to zero. Then, keeping the opponent-colour co-ordinates constant, we allow the observers to adjust the luminance of each colour. To display the colours they are converted from DKL to cone-contrast space, then LMS ,and finally to monitor RGB using the monitor's calibration matrix.

We established the iso-luminant point for each colour and each observer using a heterochromatic flicker paradigm. Each colour was displayed at the centre of the screen and alternated temporally with a grey patch of constant luminance (20 cd/m^2); the two patches were alternated at a frequency of 10 Hz. The luminance of the target colour was then adjusted manually by the observer until the perception of flicker was minimised. It is this point of minimal flicker that defines the point at which the luminance of the two patches is equal for that observer [16]. The observers repeated this procedure for each colour a total of 4 times, and there were 16 colours, making 64 judgements in all. The average of the 4 trials was used as the final isoluminant colour for a given observer. After the adjustment procedure, each colour was iso-luminant with the same grey patch, and by implication all the colours were therefore isoluminant with one another for a particular observer.

Saturation adjustment

Once all the colours had been adjusted to be iso-luminant, the observers adjusted the saturation (which we referred to as colourfulness). It was important here to ensure that the observers only adjust the saturation component of the colour, without adjusting its luminance. However, since the luminance is defined on an observer-by-observer basis, this is not straightforward to achieve. We employ the assumption that if the target colour and grey background are isoluminant with one another, then any colour falling on the line in colour-space that is defined by these two colours will also be isoluminant. Colours along this same line will also have approximately the same hue as the target colour. Thus, observers changed the colour by moving it incrementally along the line joining the iso-luminant grey, and original target colour. These adjustments were performed in the monitor's linear RGB space, but could equally well be performed in any other linear transform of this space. Due to the restricted gamut of the monitor there were some colours for which the saturation could not be increased much beyond its original value; to circumvent this, observers were asked to only reduce the saturation in their adjustments.

The 16 colours were arranged in a 16×16 grid of squares. The observer could then go from patch to patch adjusting the saturation of the colours until all were perceptually equal. This procedure was repeated 4 times for each observer, and the average of these four trials was used as the iso-saturation setting in the final experiment. We also note here that observers were not aware of the ultimate goal of the experiment when they were performing the luminance and saturation adjustments.

References

[1] M. Comer and E. Delp. Morphological operations for color image processing (1999).

[2] C. L. Novak and S.A. Shafer. Color edge detection. In *Proceedings DARPA Image Understanding Workshop, Vol. 1.*, pp. 35–37, Los Angeles, CA, USA (1987).

[3] D. A. Socolinsky and L. B. Wolff. Multispectral image visualization through first-order fusion. *IEEE Transactions on Image Processing*, **11**(8), 923–931 (2002).

[4] A. G. Hanbury and J. Serra. Mathematical morphology in the hls colour space. In *Proceedings of the 12 British Machine Vision Conference*, pp. 451–460, Manchester, England (2001).

[5] A. G. Hanbury and J. Serra. Mathematical morphology in the cielab space. *Image Anal Stereol*, **21**, 201–206 (2002).

[6] R. W. G. Hunt. *Measuring Colour, Third Edition*. Fountain Press, Kingston-upon-Thames, England (1998).

[7] A. A. Gooch, S. C. Olsen, J. Tumblin, and Bruce Gooch. Color2gray: salience-preserving color removal. In *SIGGRAPH '05: ACM SIGGRAPH 2005 Papers*, pp. 634–639, New York, NY, USA (2005). ACM Press.

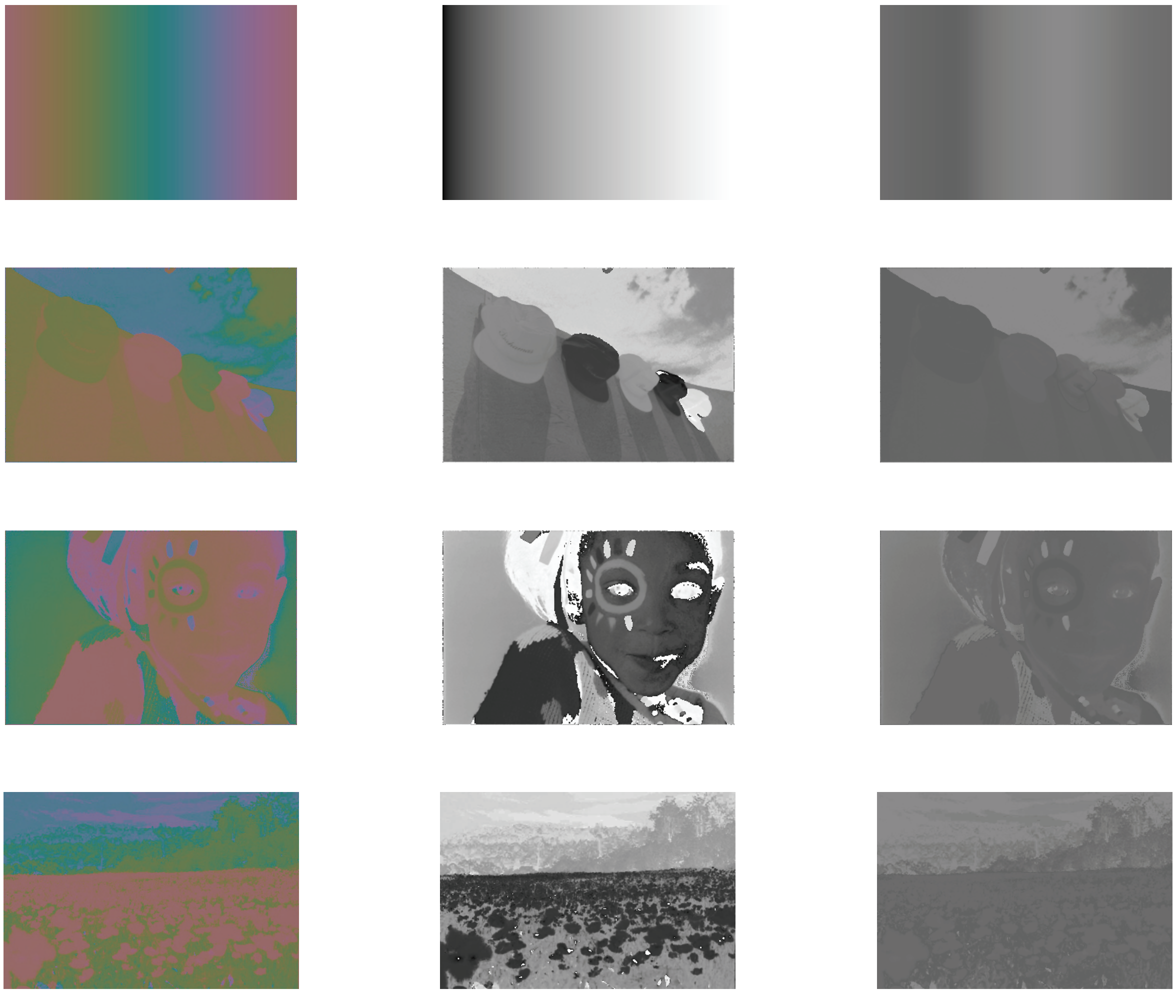

Figure 7. Greyscale renderings of iso-luminant, iso-saturation images. The left hand column shows the colour original (transformed from standard colour images into iso-saturation and iso-luminant images); the central column shows a greyscale rendering using hue-angle scaled between 0 and 1; the right column uses a greyscale rendering using the experimentally derived hue-ordering. The hats and girl images are copyright of the Eastman Kodak Company, while the poppies image was downloaded from http://www.slackerhack.com/current_affairs/index.html.

[8] M. Grundland and N. A. Dodgson. The decolorize algorithm for contrast enhancing, color to grayscale conversion. Technical Report UCAM-CL-TR-649, Computer Laboratory, University of Cambridge (2005).

[9] K.T. Mullen. The contrast sensitivity of human colour vision to red/green and blue/yellow chromatic gratings. *Journal of Physiology*, **359**, 381–400 (1985).

[10] M. G. Kendall and B. Babington-Smith. On the method of paired comparisons. *Biometrika*, **31**, 324–345 (1940).

[11] P. Ledda, A. Chalmers, T.Troscianko, and H. Seetzen. Evaluation of tone mapping operators using a high dynamic range display. In *Proceedings ACM SIGGRAPH 2005, Annual Conference on Computer Graphics*, pp. 640–648, Los Angeles, CA, USA (2005).

[12] M. D. Fairchild and E. Pirrotta. Predicting the lightness of chromatic object colors using cielab. *Color Research and Application*, **16**, 385–393 (1991).

[13] M.Stokes, M. Anderson, S. Chandrasekar, and R. Motta. A standard default colour space for the internet: srgb, http://www.w3.org/graphics/color/srgb.html.

[14] A. M. Derrington, J. Krauskopf, and P. Lennie. Chromatic mechanisms in lateral geniculate nucleus of macaque. *Journal of Physiology (London)*, **357**, 241–265 (1984).

[15] V.C. Smith and J. Pokorny. Spectral sensitivity of the foveal cone photopigments between 400 and 500 nm. *Vision Research*, **15**, 161–171 (1975).

[16] B. Lee, P. Martin, and A. Valberg. The physiological basis of heterochromatic flicker photometry demonstrated in the ganglion cells of the macaque retina. *Journal of Physiology*, **404**, 323347 (1988).