

A Hardware-Independent Colour Calibration Technique

D. Karatzas¹ and S.M. Wuerger²

¹School of Electronics and Computer Science
University of Southampton, Southampton SO17 1BJ, UK

²School of Psychology
University of Liverpool, Liverpool L69 7ZA

<D.Karatzas@ecs.soton.ac.uk>

<S.M.Wuerger@liv.ac.uk>

Abstract

Most colour-normal human observers have no difficulty to adjust a coloured light such that it appears neither red nor green, or such that it appears neither yellow nor blue. Moreover, it has been shown that these hue judgements are not significantly influenced by language or age and individual differences in colour sensitivity are not reflected in the unique-hue settings. Here we show how we can use the invariance of these unique-hue judgements to develop a colour calibration technique for display devices, which eliminates the need for an external calibration standard or a measurement device.

1 Introduction

The colours of the same image when viewed on different display devices are likely to appear dissimilar due to the different viewing conditions and the individual physical characteristics and specific contrast/brightness settings of the display devices. This problem, namely to communicate colour veridically across different platforms, is often addressed by the use of specialised colour calibration hardware, by providing device profiles based on the information given by manufacturers or using standardised colour

spaces. These techniques necessitate a degree of technical knowledge and are unfriendly to the low-tech end-user. In this paper we propose a different, inexpensive and user-friendly solution which is solely based on perceptual colour judgements.

‘Unique hues’ were first mentioned by Hering [1] who proposed that any hue can be described in terms of its redness or greenness and its yellowness or blueness. Red and green are opposite hues because they cannot be elicited simultaneously by a single colour stimulus; the same is true for blue and yellow. This observation led Hering to postulate the existence of two colour-opponent channels coding red-green and yellow-blue sensations. Hurvich and Jameson [2] used a hue cancellation technique to determine the unique hues experimentally. Observers viewed a small test stimulus on a white surround and adjusted the test light until it looked ‘neither red nor green’ or ‘neither yellow nor blue’. The idea is that unique yellow and unique blue are obtained by silencing a Red-Green mechanism; and unique red and unique green are obtained by silencing a Yellow-Blue mechanism [2, 3]. Larimer and colleagues tested the linearity of these mechanisms and concluded that the Yellow-blue mechanism (yielding unique red and green) is non-linear in cone space, whereas the Red-Green mechanism (yielding unique yellow and blue) can be described as a linear function of cone space [4, 5]. Subsequently, researchers noticed that the consistent failure of linearity for unique red and unique green settings can simply be accounted for by assuming not a single but two different linear mechanisms [6, 7] (Figure 1). Furthermore, unique hue judgements are not significantly influenced by language or age [8, 9] and individual differences in colour sensitivity are not reflected in the unique hue judgements [10].

The human visual system seems to be able to calibrate itself, but how this is achieved is an open question [11]. Whatever the underlying mechanism is, we can exploit the invariance resulting from such auto-calibration to solve an important technological problem, namely the ability to communicate colour veridically across different platforms. In this paper we propose a low-tech colour calibration solution that exploits the invariance of these particular colour mechanisms to achieve device-independent calibration: the unique hues can serve as an internal standard against which any display device can be calibrated. In contrast to existing solutions the calibration technique presented here makes no assumptions about the chromaticities or luminances of the primaries of the particular display device.

Section 2 details the colour calibration process, data collection (2.1), construction of a device profile (2.2.), and derivation of the colour transformation (2.3). A calibration example along with some preliminary results is subsequently given in section 3.




Figure 1: Unique hue loci in a cone-opponent cardinal colour space [12].

2 Colour Calibration Process

The Colour Calibration Process comprises three distinct stages. First, data are collected about a specific hardware setup. During this stage, the subject (computer user) is asked to make a number of unique hue judgements that allow us to determine key colour properties of the display device. Subsequently, these data are analysed and the results are used to build a device profile for the specific hardware setup. Finally, given any two device profiles, where the first corresponds to the user display device (which is to be calibrated), and the second to the reference display device (which is to be matched), the corresponding colour transformation is calculated.

1.1 Unique Hues: Data Collection

The psychophysical data are collected using a hue selection task. In this task the observer is asked to select from an annulus of coloured patches the patch that is ‘neither yellow nor blue’ (to obtain unique red and unique green) or ‘neither red nor green’ (to obtain unique yellow and unique blue). Prior to the hue selection task, a short grey evaluation process takes place to obtain a neutral grey colour for the particular device. By neutral grey we refer to a colour which appears of no specific hue when viewed on the display device (the

assumption that a default $R = G = B$ colour is neutral grey is not necessarily correct for an uncalibrated device).

Determining the Neutral Grey Point

Since the perceived colour of a patch is partly determined by the adaptational state and the background it is presented on [13], we first need to determine a neutral grey colour for the particular device. This neutral grey is then used as the background colour in all subsequent colour judgements.

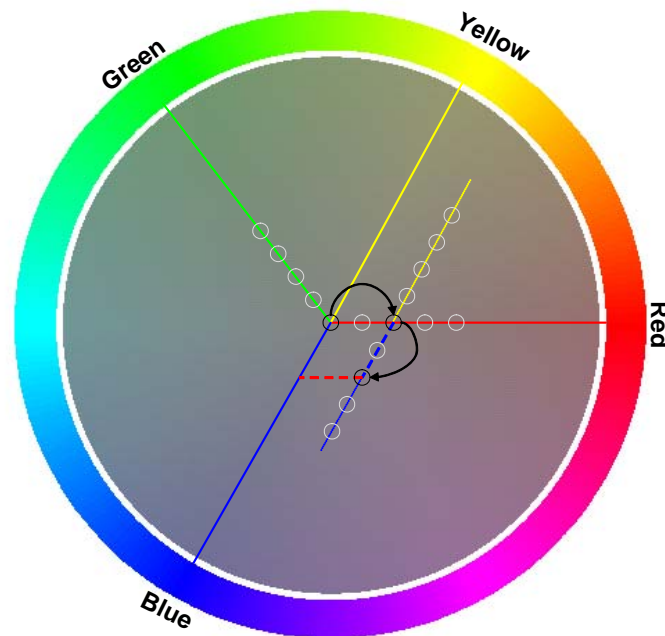


Figure 2: Visualisation of the grey evaluation process.

The iterative grey selection procedure is visualised in Figure 2 which shows a vertical slice of the HSV colour system. Saturation is defined as the distance from the centre of the disc, while Hue is defined as the angle from the red colour (red having a hue value equal to zero). The Value axis is perpendicular to the paper plane.

The centre point on the disc represents the HSV colour which is the original estimate for neutral grey. In the first iteration, this original estimate is presented to the user along with colours of increasing saturation on the directions of Red and Green hues. The colours are shown as an annulus of patches and the user is asked to select that greyish patch that appears 'neither Red nor Green'. As soon as the user makes a selection, the selected colour

is considered as the new neutral grey estimate, and a fresh set of colours are presented to the user, based on the new estimate. This time the colours of the stimuli presented are selected on the directions of Blue and Yellow hues, biased by the previous selection of the user. The user is asked to select that patch that is ‘neither yellow nor blue’. This process can be repeated, using smaller Saturation steps each time, to iteratively obtain better approximations of the neutral grey.

Collecting the Unique Hue Judgements

At each trial an annulus of patches of similar colours is presented (Figure 3). Settings for ‘unique red’ and ‘unique green’ are obtained by asking the user to select the patch that contains ‘neither yellow nor blue’. To obtain ‘unique yellow’ and ‘unique blue’ settings the user is asked to choose that patch that looks ‘neither red nor green’.




Figure 3: A unique hue (unique red in this case) assessment step.

At each trial the Value and the Saturation level are the same for all patches and the only independent variable is Hue. A specific Value level defines a horizontal plane (a cut) of the HSV colour space, while a specific Saturation level defines a cone as depicted in Figure 4. The combination of the Saturation and the Value levels associated with an individual assessment (S_{as} , V_{as}) defines a circle, and the range of Hues being tested (dH) an arc of that circle.

To accurately estimate the locus of each unique hue, the colour space must be sampled adequately. There is of course a trade-off between adequately sampling the colour space and keeping the number of assessments low, so that the end task is palatable to the user.

The number of pairs defined are at least 6 for each unique hue (for the basic test configuration), resulting in a final task with a minimum of 24 assessments.

At each trial the HSV co-ordinates of the selected hue ($h \in dH, S_{as}, V_{as}$) and their RGB representation is stored.

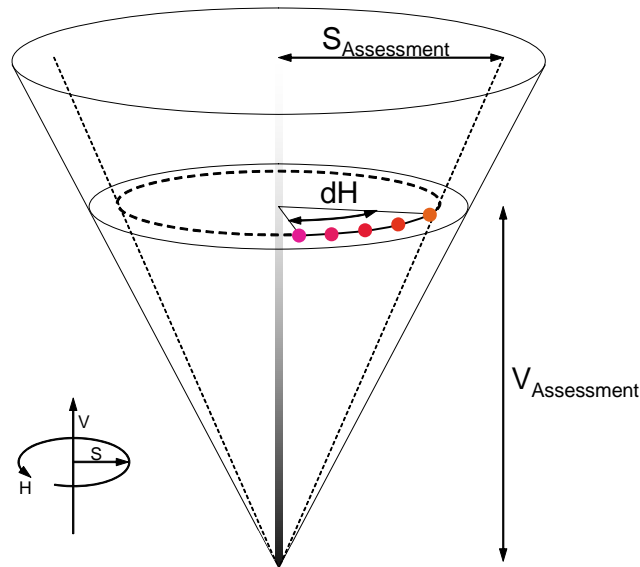


Figure 4: The colours tested during an assessment defined by ($dH, S_{Assessment}, V_{Assessment}$).

1.2 Device Profile

The data (unique hue judgements) collected for a particular display device are then used to construct a profile which characterises the colour properties of the device. The device profile contains information about the loci of the four unique hues in a device-dependent colour space which is linearly related to XYZ. The unique hue loci in such a colour space will be planes described by their normal vectors. The need to describe the colour measurements in a linear colour space is explained next while the notion of linear RGB space is introduced to facilitate the analysis.

Linear device-dependent RGB space

To a first approximation, unique hues can be described by planes in the LMS colour space [9], and therefore in any colour space which is a linear transformation of LMS (i.e. CIE XYZ). The measurements collected through the unique hue assessments are expressed in the native RGB colour system and there is, in general, no linear transformation between the native RGB colour system and LMS (or XYZ), since the applied voltage and the

luminance output are in general nonlinearly related. A conventional CRT has a power-law response to voltage with a transfer function of the general form:

$$R' = (\textit{gain} \cdot R + \textit{offset})^\gamma \quad (1)$$

where R' (or G' or B') is the measured luminance of the phosphors and R (or G or B) the pixel value indicating the voltage applied to each gun. These new set of values define a colour space (called $R'G'B'$ from now on) which is linearly related to XYZ . Various psychophysical techniques exist [e.g. 14] to estimate the non-linearity γ , the gain and the offset and here we assume that these parameters are known and that we can assume a linear, device-dependent RGB space (i.e. $R'G'B'$ space).

The Device Profile

The loci of the four unique hues are recorded as RGB coordinates and then transformed to the (device specific) $R'G'B'$ system. In the $R'G'B'$ system, each set of points can be fitted by a plane since the $R'G'B'$ space is linearly related to XYZ space. The normal vectors of the four unique hue planes are computed, and comprise along with the linearised $R'G'B'$ coordinates and the estimated parameters of eq. 1 the device profile.

Both the sets of points and the normal vectors are important, since the same device profile might be used either to describe the display device to calibrate (user display), or the reference display device. The transformation which maps the unique hue points obtained on the reference display onto the unique hue planes obtained on the user display is derived in the next section.

1.3 Derivation of the Colour Correction Transformation

From the two device profiles we can now compute the appropriate transformation so that colours that belong to a particular unique hue plane of the reference device R , to be transformed to colours that belong to the corresponding unique hue plane of the user device U (see Figure 5).

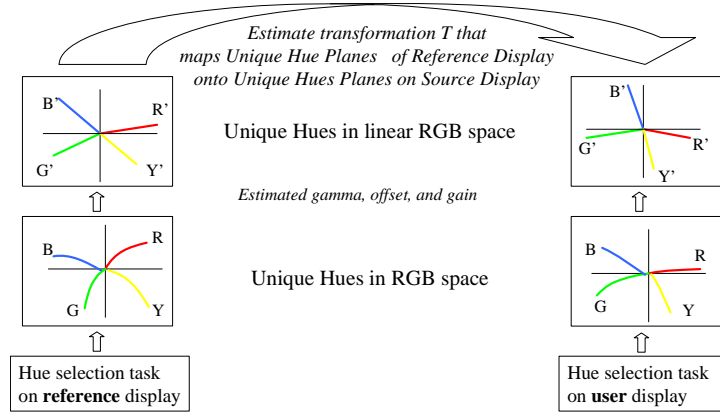


Figure 5: Transformation from reference to user display device derived from unique hue judgements.

Since both the reference- $R'G'B'$ space and the user- $R'G'B'$ space are linearly related to XYZ, the relationship between their vector bases will also be a linear one (eq. 2). If we call T the transformation matrix of eq. 2, then any colour with coordinates $[r_R \ g_R \ b_R]$ in the reference- $R'G'B'$ space can be converted to its user- $R'G'B'$ coordinates $[r_U \ g_U \ b_U]$ using the same transformation T (eq. 3).

$$\begin{bmatrix} \bar{\mathbf{r}}_R \\ \bar{\mathbf{g}}_R \\ \bar{\mathbf{b}}_R \end{bmatrix} = \begin{bmatrix} t_{11} & t_{12} & t_{13} \\ t_{21} & t_{22} & t_{23} \\ t_{31} & t_{32} & t_{33} \end{bmatrix} \cdot \begin{bmatrix} \bar{\mathbf{r}}_U \\ \bar{\mathbf{g}}_U \\ \bar{\mathbf{b}}_U \end{bmatrix} \quad (2)$$

$$[r_R \ g_R \ b_R] \cdot T = [r_U \ g_U \ b_U] \quad (3)$$

If there was a one-to-one correspondence between the colours collected on the user and the ones collected on the reference device, solving for the transformation T would entail merely solving the over-determined linear system of eq. 3. In this case though, the requirement is for each set of reference- $R'G'B'$ points to map to a given plane in the user- $R'G'B'$ space. Suppose that the normal vector that defines the plane for unique hue h on the user display device is $[\alpha_U^h \ \beta_U^h \ \gamma_U^h]^T$. Then, every colour $[r_U^h \ g_U^h \ b_U^h]$ which appears of the unique hue h when viewed on the user display device should lie on that plane. Also, every colour $[r_R^h \ g_R^h \ b_R^h]$ which appears of the unique hue h when viewed on the reference display device, when transformed should yield a colour that satisfies eq. 4, i.e. appears of the unique hue h when viewed on the user display device.

$$[r_R^h \ g_R^h \ b_R^h] \cdot T \cdot \begin{bmatrix} \alpha_U^h \\ \beta_U^h \\ \gamma_U^h \end{bmatrix} = [r_U^h \ g_U^h \ b_U^h] \cdot \begin{bmatrix} \alpha_U^h \\ \beta_U^h \\ \gamma_U^h \end{bmatrix} = 0 \quad (4)$$

Eq. 4 defines an over-determined linear system which we can solve utilising the data (sets of points and corresponding planes) contained in the two device profiles.

Implementation issues

In practice, the system of linear equations described by eq. 4 yields a set of solutions reflecting the fact that there is an infinity of ways in which a set of points can be mapped onto a plane; differing in terms of the scaling and/or rotation of the data set (including the obvious zero solution). A solution in the least-squares-sense does not necessarily yield the right transformation unless further constraints are employed. We formulated the above as a constrained minimisation problem, where the objective function is defined as the squared distance of the transformed reference-R'G'B' colour points from the given user-R'G'B' planes. The constraints employed are two; first, the sum of coefficients in each row of the transformation is required to be equal to one, and second the highest saturation value in the original set of colours is required to be preserved after the transformation.

3 Experimental Results

A controlled experiment was performed where a known transformation was used to alter the colour appearance on a display device. A user performed the unique hue assessments once with the device colours unaltered, and a second time with the device colours transformed using the aforementioned transformation. The algorithm discussed above was subsequently used to compute the transformation based solely on the information obtained during the assessments. The resulted transformation was then compared to the a priori known one. The sum of the squared differences D between the coefficients of the true transformation and the computed one was used to assess the result.

Since the calculation of the transformation is formulated as a minimisation problem, the solution arrived at depends on the initial value for the transformation matrix, and depending on the objective function the algorithm might converge at a local minimum. We run the algorithm 100 times with random initial values for the transformation matrix, and we obtained two distinct solutions. The first is practically equivalent to the a priori known transformation with a distance $D = 0.06$ from the true transformation, while the second refers to the solution which maps the original colours to their opposite ones (mirrored on line of greys) with a distance $D = 2.31$. In Figure 6 a sample image is shown along with its transformation using the true transformation and using the derived one. The average RGB difference between the two images is ~ 1.51 .




Figure 6: (a) Original image, (b) transformed using the true transformation, (c) transformed using the computed transformation.

4 Conclusions

We have shown that, from a small number of perceptual colour judgements, we can derive a sufficient amount of information to characterise the physical properties of a typical display device. From those perceptual judgements we can derive the transformation that needs to be applied to the RGB coordinates of any colour viewed on the user display to make it look identical to the colour of the same RGB code as viewed on a reference display.

The problem formulated in this paper can be approached in a multitude of ways. The constrained minimisation algorithm employed here yields promising results but further work is necessary to ensure that the correct solution is returned irrespectively of the initial search conditions. This can be achieved by either ensuring that the constrained objective function has a single global minimum or by employing certain post-processing steps to filter out or correct erroneous results.

References

- [1] E. Hering. *Outlines of the theory of the light sense*. Cambridge, Massachusetts, Harvard University Press, 1964.
- [2] L. Hurvich and D. Jameson. An opponent-process theory of color vision. *Psychological Review*, 64:384–404, 1957.
- [3] A. Valberg. Unique hues: an old problem for a new generation. *Vision Research*, 41(13):1645–1657, 2001.

- [4] J. Larimer, D. Krantz and C. Cicerone. Opponent-process additivity. I: Red/green equilibria. *Vision Research*, 14:1127–1140, 1974.
- [5] J. Larimer, D. Krantz and C. Cicerone. Opponent-process additivity. II: Yellow/blue equilibria and nonlinear models. *Vision Research*, 15:723–731, 1975.
- [6] R. Mausfeld and R. Niederee. An inquiry into relational concepts of colour, based on incremental principles of colour coding for minimal relational stimuli. *Perception*, 22(4):427–462, 1993.
- [7] B.A. Chilchilnisky and B.A. Wandell. Trichromatic opponent color classification. *Vision Research*, 39:3444–3458, 1999.
- [8] B.A.C. Saunders and J. van Brakel. Are there nontrivial constraints on colour categorization? *Behavioral and Brain Sciences*, 20(2):167–228, 1997.
- [9] B.E. Shefrin and J.S. Werner. Loci of spectral unique hues throughout the life span. *Journal of the Optical Society of America*, 7:305–311, 1990.
- [10] M. A. Webster, E. Miyahara, G. Malkoc et al. Variations in normal color vision. II Unique hues. *Journal of the Optical Society of America A*, 17(9):1545–1555, 2000.
- [11] Y. Yamauchi, D.R. Williams, D.H. Brainard et al. What determines unique yellow, L/M cone ratio or visual experience? In *Proceedings of the SPIE Vol 4421*, pages 275–278, 2002.
- [12] S.M. Wuerger, P. Atkinson and S.J. Cropper. The Cone Inputs to the Unique-Hue Mechanisms. *Vision Research*, 45:3210–3223, 2005.
- [13] S.M. Wuerger. Color appearance changes resulting from isoluminant chromatic adaptation. *Vision Research*, 36 (19):3107–3118, 1996.
- [14] H. Irtel, PXLab Manual. URL: www.uni-manheim.de/fakul/psycho/irtel/pxlab.