White-point preserving color correction

Graham D. Finlayson
Colour Research Group,
University of Derby,
Derby DE22 3BL,
United Kingdom

Mark S. Drew
School of Computing Science,
Simon Fraser University,
Vancouver V5A 1S6,
Canada

Abstract

To characterize color values measured by color devices (e.g. scanners, color copiers and color cameras) in a device–independent fashion these values must be transformed to colorimetric tristimulus values. Often it is assumed that RGBs are approximately linearly related to XYZs and so this transformation is determined by least–squares (LS) linear regression. While the LS method is guaranteed to minimize the residual squared error it makes no *a priori* statement about which colors will be mapped well and which will be mapped poorly. However, we argue that such a statement must be made. In particular because it is important to preserve the white and the gray-scale in color reproduction, we argue that achromatic colors should be preserved in color correction. This argument led us to develop a new regression procedure: the *white-point preserving least-squares fit* (WPPLS). As the name might suggest, this method finds the linear transform which maps RGBs to XYZs such that the residual squared error is minimized *subject to the constraint that white and grays are preserved*. Of course, by definition, the WPPLS regression must, in terms of squared error, deliver poorer color correction compared with the LS procedure; but, squared error need not necessarily correlate with perceived visual error. Indeed, we present a number of simulation experiments which show that the WPPLS procedure performs as well or better than the LS regression. These results provide experimental confirmation of the privileged status of white in visual perception and color reproduction.

1. Introduction

Color sensors in scanners, color copiers and color cameras are not colorimetric, in the sense that device RGB values are not a linear transformation away from the X,Y,Z tristimulus values [1]. The transformation from RGB to XYZ forms the first step in developing a device–independent description of color for these devices [2]. Given a particular set of targets or dyes one can map from RGB to XYZ using interpolation and look-up-tables[3, 4, 5]. More elegantly, the device coordinates can be transformed to visual tristimuli using a single linear transform (a 3×3 matrix). This transform is often defined to be the best linear least–squares mapping for a particular calibration set of surface reflectances.

The least-squares regression optimally maps RGBs to XYZs so that the residual squared error for a calibration data set is minimized. However, there is no mechanism for choosing the surfaces which will be mapped well and those which will be mapped poorly. For example, it is possible that for one calibration set a white reflectance may be mapped exactly and for another it may be mapped with a high colorimetric error. Given the importance of white (and the gray-scale) in color reproduction[6], we would rather not have this variable performance. Thus, in this paper we report on the white-point preserving least-squares fit (WPPLS) procedure that we have developed. As the name suggests, WPPLS is a method for determining the best least-squares transform that takes RGBs to XYZs subject to the constraint that the RGB response induced by a white reflectance is mapped without error. Importantly, when white is mapped correctly then this implies that the gray-scale is also preserved.

Of course, in terms of squared residual error, the WPPLS regression must perform less well than a least–squares regression for any particular set of calibration reflectances that are used. However, residual squared error is a purely numerical notion and does not exactly agree with perceived visual error (e.g. in terms of CIELAB[1] or CMC[7] defined color differences). To test the colorimetric performance of the LS and WPPLS correction methods we carried out a variety of simulation experiments. RGBs were generated for a Sharp JX450 color scanner along with the corresponding XYZs. The LS and WPPLS transforms were calculated under two different statistical assumptions. First, correction transforms were derived using a real set of reflectances (e.g. the Munsells) and second, with respect a hypothetical set that contains all surface reflectances (the so-called maximum ignorance case). When a real calibration set is used both the LS and WPPLS regressions deliver comparable performance; though of course WPPLS maps achromatic colors with much smaller error. Under the maximum ignorance assumption the WPPLS regression performs much better than the LS procedure.

2. Linear Color Correction

Suppose that we measure the XYZ tristimuli values for a set of n surfaces and place them in the rows of an $n \times 3$ matrix $\mathcal{M}$. We measure the RGB response for the same surfaces using a scanner (or color camera) and place the n rgb response triplets in an $n \times 3$ matrix $\mathcal{N}$. The aim of linear color correction is to find a 3×3 matrix $\mathcal{T}$ such that:

$$\mathcal{M} \approx \mathcal{N}\mathcal{T} \qquad \text{(Linear color correction)}$$

In order to solve for $\mathcal{T}$ we must define what we mean by $\approx$ (the symbol for approximation). In the method of least-squares (LS) we wish to find the transform $\mathcal{T}$ which minimizes the residual squared error. Such a transform is easily found[8] (details are given in the Appendix). Unfortunately in the method of least-squares we do not a priori know which rows of $\mathcal{N}$ will be mapped with small error and which will be mapped with large error and we would like to know this. In particular, should the LS procedure map white (and the gray-scale) with a large error then this will lead to unacceptable color reproduction. The white-point preserving least-squares (WPPLS) procedure is a correction method which is specifically designed to circumvent this problem.

To see how the WPPLS procedure works it is useful to normalize the matrices $\mathcal{M}$ and $\mathcal{N}$ with respect to white. Specifically, let us suppose that the XYZ response to a white reflectance is (x^w, y^w, z^w) and the device RGB response to white (r^w, g^w, b^w). If (x_i, y_i, z_i) is the ith row vector of $\mathcal{M}$ then this tristimulus normalized to white is defined to be equal to $(x_i/x^w, y_i/y^w, z_i/z^w)$. Similarly, the ith row of the device response matrix N normalized to white is equal to $r_i/r^w, g_i/g^w, b_i/b^w$. Normalizing to white has the advantage that the XYZ tristimulus and RGB response for a perfect white diffuser are both be equal to (1,1,1).

The white-point preserving color correction problem is now easily stated:

$$\mathcal{M}' \approx \mathcal{N}'\mathcal{T} \;\; \textbf{AND} \;\; \underline{u}^t\mathcal{T} = \underline{u}^t$$
$$\text{(White-point preserving color correction)}$$

Here $'$ represents normalizing to white and $\underline{u}^t$ is a row vector, equal to (1,1,1), representing white (note the underscore notation is used to denotes a vector and ^{t} denotes the transpose operation). The white-point preserving least-squares procedures finds the transform $\mathcal{T}$ that minimizes the overall residual squared error and at the same time preserves white. We show in the appendix that the transform $\mathcal{T}$ is straightforward to find; that is, there is a simple closed form solution (see Appendix for details). The basic insight that is exploited is that if $\underline{u}^t\mathcal{T} = \underline{u}^t$ then the sum of the components in each column of $\mathcal{T}$ must equal 1 and this constraint is easily enforced using the method of Lagrange multipliers.

Figure 1 illustrates the LS and WPPLS procedures in action. At the top of the figure we have tabulated the RGB response of a Sharp JX450 to the 6 Macbeth checker patches (those with color names red, purple, blue, green red and yellow) and to a perfect diffusing white. All responses have been normalized to white. At the bottom of the Figure the corresponding white-normalized XYZ responses (viewing illuminant D65) are shown. In the context of this example, the goal of linear color correction is to map the 7×3 RGB matrix of responses as close as possible to the 7×3 matrix of XYZ tristimuli.

Performing a least-squares regression of the RGB data onto the XYZ tristimuli results in the table of numbers in the middle left. Notice that reasonable color correction has been achieved: the numbers in the transformed table are quite close to the desired XYZs. Importantly, the ΔE errors are reasonably small, ranging from 1.7 to 6.9 CMC[7] units[1]. Notice, however, that despite the fact that the RGB and XYZ response to white are both (100,100,100) that the LS corrected RGB response is equal to (99.8,100.0,104.1); a CMC error of 1.9 has been incurred. Perhaps a more serious problem is the fact that the LS corrected B(lue) response is equal to 104.1 is bigger than 100 and this is physically impossible since white is the most reflective surface reflectance.

The white-point preserving correction attains broadly similar performance (see table in the middle right of the Figure) for the 6 Macbeth checker patches. Notice, however, that white has been perfectly corrected and as such has 0 CMC error. The importance of this fact is borne out in more comprehensive simulation experiments reported in the next section.

3. Simulation Experiments

We generated XYZ tristimuli, for viewing illuminant D65, for 3 spectral reflectance data sets: the 462 Munsells measured by Nickerson[9], the 24 Macbeth Color checker[10] patches and the 170 real object reflectances measured by Vrhel et al[11]. Corresponding RGB responses were generated for a Sharp JX450 color scanner. We calculated the LS and WPPLS correction transforms for each of the data sets. After applying these transforms to the RGB data, we calculated the residual error using the CMC[7] color difference formulae. The results of these simulation experiments are summarized in Table 1.

[1]Broadly speaking CMC units correlate with CIELAB units. However, the CMC formula more accurately explains experimental color-difference data.

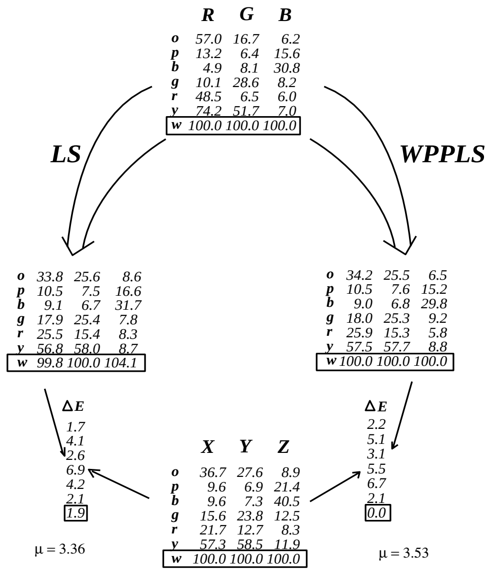

Figure 1: *A comparison of LS and WPPLS linear color correction*

	LS			WPPLS		
Data Set	μ	sd	white	μ	sd	white
Munsell	1.54	1.45	1.71	1.75	1.57	0.41
Macbeth	2.01	1.29	1.55	1.86	1.51	0.40
Object	1.53	1.63	1.51	1.60	1.72	0.40

Table 1: *Colorimetric (CMC) error in mapping scanner RGBs to XYZs using the LS and WPPLS regressions*

For the Munsell and Object reflectance data sets, the LS regression delivers slightly better performance than the new WPPLS procedure. However, the WPPLS, provides better performance for the Macbeth checker and this is probably due to the high number (25%) of achromatic colors on the chart. In the columns headed 'white' we have tabulated the colorimetric error incurred in mapping the whitest Munsell reflectance (the reflectance that is closest to the perfect white diffuser). As we would expect the WPPLS transform maps this reflectance with lower colorimetric error compared with the LS procedure. In summary, the message that Table 1 conveys is that a correction transform which preserves white (and correction transforms should preserve white) does not incur a significant colorimetric overhead compared with the unconstrained LS regression.

The data presented in Table 1 reflects very much the best case scenario for color correction. That is, we calibrate our color device for a given data set and then we test our calibration using the same data. Of course, this scenario will not often apply to the real world since it is unlikely that we shall only scan (or take pictures of) the same reflectances that we use for calibration. Indeed, it can be

	LS			WPPLS		
Data Set	μ	sd	white	μ	sd	white
MUN	6.31	2.25	12.70	3.44	2.15	0.70
MAC	6.49	2.87	12.70	4.03	3.12	0.70
OBJ	6.20	2.72	12.70	4.14	2.51	0.70

Table 2: *Colorimetric (CMC) error in mapping scanner RGBs to XYZs using the LS and WPPLS regressions*

argued that we do not want to calibrate our color device for any particular reflectance set since this might lead to poor colorimetric performance for any reflectances outside this set. To address this problem, many authors make the assumption that all reflectances occur with equal likelihood; simply put, any reflectance function that you can draw on a graph is assumed as likely as any other. Under the maximum ignorance assumption the correction transform can be derived via a numerical procedure involving the spectral sensitivity functions of the color device and the XYZ color matching curves[12]. That this is so is very useful since we side-step the whole calibration procedure (and calibration is time consuming).

In Table 2 we compare the colorimetric performance delivered by the LS and WPPLS transforms, derived under the maximum ignorance assumption, operating on the Munsell, Macbeth and Object spectral reflectance data sets. It is apparent that, in all cases, the WPPLS regression delivers substantially better performance. Indeed, it is reasonable to conclude that maximum ignorance color correction is acceptable if and only if the white-point preserving regression is used.

4. Conclusions

The popular least-squares method for color correction finds the linear transform that maps device RGBs to XYZs with minimum error. Unfortunately, this method makes no statement about which colors will be mapped with low error and we would like to make such a statement. In particular, because maintaining white, and the gray-scale, is so fundamental to color reproduction we would like to develop a correction transform that, while optimizing a global error criterion also preserves white. Such a method was reported in this paper. The **white point preserving least-squares procedure** (WPPLS) finds the linear transform that maps RGBs to XYZ tristimuli such that the residual error is minimized subject to the constraint that whites, and the gray-scale, are corrected without error. In terms of colorimetric error, the WPPLS regression always performs at least as well as unconstrained linear regression and sometimes substantially better.

Acknowledgements

This work was carried out when Graham Finlayson was a faculty member in the Department of Computer Science at the University of York.

References

[1] G. Wyszecki and W.S. Stiles. *Color Science: Concepts and Methods, Quantitative Data and Formulas*. Wiley, New York, 2nd edition, 1982.

[2] B.A. Wandell and J.E. Farrell. Water into wine: Converting scanner RGB to tristimulus XYZ. In *Device-Independent Color Imaging and Imaging Systems Integration*, pages 92–101. SPIE Vol. 1909, 1993.

[3] S. Suzuki, T. Kusunoki, and M. Mori. Color characteristic design for color scanners. *Applied Optics*, 29:5187–5192, 1990.

[4] P.C. Hung. Color rendition using three-dimensional interpolation. In *Imaging Applications in the Work World*, pages 111–115. SPIE Vol. 900, 1988.

[5] P.C. Hung. Colorimetric calibration for scanners and media. In *Proceedings of the 1991 Electronic Imaging Meeting*, San Jose, CA, Feb, 1991.

[6] R.W.G. Hunt. *The Reproduction of Color*. Fountain Press, 5th edition, 1995.

[7] M.H. Brill. Suggested modification of cmc formula for acceptability. *Color Research and Application*, 17(6):402–404, 1994.

[8] G.H. Golub and C.F. van Loan. *Matrix Computations*. John Hopkins U. Press, 1983.

[9] Newhall, Nickerson, and D.B. Judd. Final report of the OSA subcommittee on the spacing of the munsell colors. *J. Opt. Soc. Am. A*, 33:385–418, 1943.

[10] C.S. McCamy, H. Marcus, and J.G. Davidson. A color-rendition chart. *J. App. Photog. Eng.*, pages 95–99, 1976.

[11] M.J. Vrhel, R. Gershon, and L.S. Iwan. Measurement and analysis of object reflectance spectra. *Color Research and Application*, 19:4–9, 1994.

[12] P.L. Vora and H.J. Trussell. Measure of goodness of a set of color-scanning filters. *J. Opt. Soc. Amer. A*, 10:1499–1508, 1993.

Appendix

Let $\mathcal{N}$, $\underline{v}$ and $\underline{c}$ be an $n \times 3$ data matrix, an $n \times 1$ data vector and a 3×1 coefficient vector (which is to be solved for) respectively. Let

$$I \;=\; |\mathcal{N}'\underline{c} - \underline{v}| \tag{1}$$

where $|.|$ is the L2 norm (i.e., vector length). It follows that I^2 can be written as:

$$J \;=\; I^2 \;=\; \underline{c}^t\mathcal{N}^t\mathcal{N}\underline{c} - 2\underline{c}^t\mathcal{N}^t\underline{v} + \underline{v}^t\underline{v} \tag{2}$$

We can find the $\underline{c}$ which minimizes (2) by differentiating J and equating to 0:

$$\frac{\delta J}{\delta \underline{c}} \;=\; 2\mathcal{N}^t\mathcal{N}\underline{c} - 2\mathcal{N}^t\underline{v} \;= 0 \tag{3}$$

It follows that

$$\underline{c} \;=\; [\mathcal{N}^t\mathcal{N}]^{-1}\mathcal{N}^t\underline{v} \tag{4}$$

Equation (4) is the solution to the least-squares (LS) color correction problem. To find the i^{th} column vector of the 3×3 LS transform $\mathcal{T}$ (discussed in the text) simply substitute the i^{th} column of $\mathcal{M}$ for $\underline{v}$ in (4).

Let us augment the minimization of (2) with a Lagrange multiplier term:

$$J \;=\; \underline{c}^t\mathcal{N}^t\mathcal{N}\underline{c} - 2\underline{c}^2\mathcal{N}^t\underline{v} + \underline{v}^t\underline{v} + \lambda(\underline{c}^t\underline{u} - 1) \tag{5}$$

where $\underline{u}^t \;=\; (1,1,1)$. Differentiating with respect to the Lagrange multiplier λ and equating to 0 we have:

$$\frac{\partial J}{\partial \lambda} \;=\; \underline{c}^t\underline{u} - 1 = 0 \;\Rightarrow\; \underline{c}^t\underline{u} = 1 \tag{6}$$

Relation (6) tells us that when we find the stationary point of J we must have $\underline{c}^t\underline{u} = 1$. That is, the sum of the components of $\underline{c}$ equals 1. This is precisely the condition that we need for the white-point preserving minimization.

Differentiating (5) with respect to $\underline{c}$ and equating to 0:

$$\frac{\partial J}{\partial \underline{c}} \;=\; 2\mathcal{N}^t\mathcal{N}\underline{c} - 2\mathcal{N}^t\underline{v} + \lambda\underline{u} \;= 0 \tag{7}$$

Taking (6) and (7) together and applying some algebraic manipulation it can be shown that:

$$\underline{c} \;=\; [\mathcal{N}^t\mathcal{N}]^{-1}\mathcal{N}^t\underline{v} + \frac{(1 - \underline{v}^t\mathcal{N}[\mathcal{N}^t\mathcal{N}]^{-1}\underline{u})}{(\underline{u}^t[\mathcal{N}^t\mathcal{N}]^{-1}\underline{u})}[\mathcal{N}^t\mathcal{N}]^{-1}\underline{u} \tag{8}$$

Equation (8) provides a solution to the white-point preserving least-squares color correction problem. To find the ith column vector of the WPPLS transform $\mathcal{T}$ we substitute the ith column of $\mathcal{M}$ for $\underline{v}$ in (8).