Intensity Constrained Error-Less Colour Correction

Graham D. Finlayson and Peter Morovic
School of Computing Sciences
University of East Anglia
NR4 7TJ Norwich, United Kingdom
{graham,peter}@cmp.uea.ac.uk

Abstract

Colour correction is the mapping of device-dependent RGBs to device-independent standard CIE XYZs. Due to the nature of colour image formation and the existence of metamerism, this mapping is inherently one-to-many and thus ill-posed. However, normally it is solved for through an error-minimising linear one-to-one transform.

In this paper we propose to make use of a definition of metamerism while maintaining the simplicity of a linear transform in defining an error-less colour correction. We say that a mapping is error-less if the RGB-XYZ pair put in correspondence through this mapping is such that a real, physically realisable reflectance that induces this pair exists. We show how we can solve for such a mapping using constrained linear least squares optimisation. However, since this problem is highly constrained, we introduce a notion of error in our calculations, building on a paramer set instead (the set of reflectances that map to a small uniform region in RGB space). We show that as little as 0.5% error is sufficient for a solution to exist.

We find that the metamer set constrained linear colour correction works equally well as ordinary linear least squares in terms of the mean and median CIE ΔE, however reduces the overall maximum significantly. We also show that unconstrained linear least squares is not error-less. In particular, those samples that are not error-less are saturated samples, for which the metamer set constrained method reduces the mean, median as well as maximum error dramatically.

1. Introduction

Colour imaging devices coarsely sample the colour signal through three colour filters in the long (red), medium (green) and short (blue) wavelength range, resulting in a response known as the RGB. These responses are a composite descriptor of the light, the surface and the device. Since devices vary in terms of their spectral sensitivities, and these in turn are different from the human visual system model, the CIE XYZ colour matching functions, a transformation from the raw device-dependent responses to a standard device-independent space is required. This transform is known as *colour correction* and maps RGBs to corresponding CIE XYZs.

Most methods are based on a simple 3×3 linear transform [12], however, unless device spectral sensitivities are linearly related to the CIE colour matching functions, this approach is prone to error. The reason for this error lies in the very principle of how colour is formed. Colour image formation maps continuous functions of wavelength, reflectances, to three values only, the response vector. Due to this difference in dimensions it is apparent that this is a many-to-one mapping: there are many reflectances, called *metamers*, corresponding the same RGB. This means that given an RGB we cannot uniquely recover one surface spectral reflectance that induced it. Since metamerism also depends on the spectral sensitivities, different sets of reflectances will induce identical response for different devices. It can be seen that this ambiguity translates also into the problem of colour correction. The reflectances that induce a single RGB will correspond to a set of possible XYZs. Thus also colour correction is fundamentally a one-to-many mapping.

In this paper we argue that this ambiguity is useful. While unconstrained linear transforms simply build on a set of known corresponding RGBs and XYZs and minimise error in some way, they do not take into account metamerism and thus the underlying reflectances that correspond to these responses. It then follows that such methods can and do match XYZs to RGBs which have no foundation in physical surfaces. Instead, given an RGB, we know the entire set of possible XYZs, as this is the set of metamers corresponding to the RGB projected to XYZs – we refer to this set as the *feasible set*. At the same time we need to choose a single XYZ and wish to maintain the simplicity of a linear transform. Given these assumptions, we can define an *error-less* transform. Since error in colour correction arises due to metamerism, we say that a linear transform is error-less if it maps RGBs to XYZs which are in the feasible set. Thus any RGB, XYZ pair is such that there exists a corresponding surface reflectance mapping to this pair. One example of an existing error-less transform is the *metamer set constrained colour correction*[2] approach. While based on the same principle, this method is based on an approximation of the metamer set and results in a non-linear relationship between RGBs and XYZs.

First we recapitulate the metamer set framework in the following section. Then, in Section 3, we describe the concept of error-less transforms. In Section 4 we conduct some experiments comparing a constrained error-less transform to standard unconstrained linear least squares. Finally, in Section 5, we conclude and summarise our findings.

2. The Metamer Set

Given an RGB, the spectral sensitivities of a device, the illuminant spectral power distribution and a linear model of surface reflectance of arbitrary dimension ≥ 3, it is not possible to solve for a unique reflectance that corresponds to this RGB. However, it has been shown before[9], that it is possible to solve for an infinite set of physically possible reflectances, that all induce the same RGB. Such reflectances are called metamers and form the *metamer set*.

Let us adopt a linear algebra representation and define colour image formation under the Mondrian world assumptions[6]. Each spectral function is written as a discrete sampling of the signal at $q = 31$ sample points, covering the interval $[400, 700]$nm at 10nm steps[11]. We denote $\mathbf{r}$ the $q \times 1$ vector of surface reflectance, $\mathbf{e}$ the $q \times 1$ vector of the illuminant spectral power distribution, $\mathbf{R}$ the $q \times 3$ matrix of device spectral sensitivities and ρ the 3×1 RGB response vector. Colour image formation is written as:

$$\rho = \mathbf{R}^T D(\mathbf{e})\mathbf{r} \qquad (1)$$

where $D()$ is a function transforming a vector into a diagonal matrix.

From a mathematical perspective, Eq. 1 represents a set of three linear equations (one per spectral sensitivity and corresponding response) of q unknowns (the reflectance $\mathbf{r}$). Such systems are known as *under-determined*, and have in general a whole set of solutions instead of a single one. Assuming the three equations are linearly independent, the linear system has $q-3$ degrees of freedom.

One way to reduce the degrees of freedom is to make use of linear models of surface reflectances. It has been found that it is possible to represent surface reflectances with vanishingly small error within a 5 to 8 dimensional basis, found by statistical analysis of measured reflectance data sets [7]. Let $\mathbf{B}$ be the $q \times n$ matrix containing n basis vectors, such that any reflectance $\mathbf{r}$ is represented by the $n \times 1$ vector σ within this basis: $\mathbf{r} = \mathbf{B}\sigma$. Using this linear model, Eq. 1 changes to:

$$\rho = \mathbf{R}^T D(\mathbf{e})\mathbf{B}\sigma \qquad (2)$$

Eq. 2 represents a set of 3 equations of n and has therefore $n - 3$ degrees of freedom. Since $n << q$ we have significantly reduced the under-determined nature of Eq. 1. Denoting the fixed part of colour formation as the $3 \times n$ matrix $\mathbf{\Lambda} = \mathbf{R}^T D(\mathbf{e})\mathbf{B}$ we thus have:

$$\rho = \mathbf{\Lambda}\sigma \qquad (3)$$

where $\mathbf{\Lambda}$ is sometimes referred to as the *lighting matrix* [8].

From standard linear algebra we learn that the solutions to Eq. 3 can be written in the form of a decomposition [4]. One part, denoted as σ_ρ, is referred to as the particular solution (or fundamental [1]), and the other part, denoted as σ_0, is referred to as black. The particular solution accounts for the actual response ρ and lies in the row space of $\mathbf{\Lambda}$, while the black part accounts for the spectral variation, results in a zero response and is hence in the orthogonal complement of $\mathbf{\Lambda}$. The number of degrees of freedom then corresponds to the number of linearly independent black solutions to Eq. 3. Since these solutions are metameric too, they are called *metameric blacks*. Thus we can write:

$$\rho = \mathbf{\Lambda}(\sigma_\rho + \sigma_0) \qquad (4)$$

While the particular solution is the so-called fixed-point of the decomposition, the metameric blacks are by definition arbitrarily scalable, thus if σ_0 is a solution, then so is $a \times \sigma_0$. Eq. 4 characterises all solutions that are mathematically correct, however our aim is to solve for reflectances and not arbitrary vectors.

Reflectances are functions of wavelength representing the proportion of light reflected from the surface they describe. It follows that they are bound between 0 (no light is reflected) and 1 (all light is reflected). This constraint, referred to as *physical realisability*, can be written as a pair of linear inequalities per sampled wavelength – in total $2q$ inequalities. Furthermore we can constrain σ weights to mimic the nature of known measured reflectance data sets. We thus call a reflectance *natural* if it can be written as a convex combination (i.e. a linear combination with non-negative weights summing to one) of some reflectances we consider representative of the nature of our data. This constraint too can be written in the form of linear inequalities, as it represents a convex set. Both physical realisability and naturalness thus delimit a convex region of reflectance space, which we refer to as *feasible reflectances*. Eq. 4 furthermore represents a $n-3$ dimensional plane of σ weights that result in a particular ρ. The intersection of the set of feasible reflectances with the plane of solutions therefore defines the metamer set $\mathcal{M}(\rho)$:

$$\mathcal{M}(\rho) = \{\sigma_i | \rho = \mathbf{\Lambda}\sigma_i\} \qquad (5)$$

The metamer set is therefore convex (an intersection of two convex sets), bounded and of infinite cardinality [9].

3. Error-less Mapping

Assuming the Luther conditions [14] do not hold, i.e. device spectral sensitivities are not linearly related to CIE colour matching functions, it follows that for a given ρ, it's corresponding metamer set $\mathcal{M}(\rho)$ becomes a bounded convex set of χ's in the space of CIE XYZs. Let us refer to this set of χ's as the *feasible region* defined by each ρ, written as:

$$\mathcal{M}(\rho) \times \mathbf{X} = \mathcal{F}(\rho; \mathbf{R} \rightarrow \mathbf{X}) \tag{6}$$

where $\times$ is a multiplication of each member of the set with a matrix and $\mathcal{F}(\rho; \mathbf{R} \rightarrow \mathbf{X})$ the set of feasible χ's corresponding to reflectances in the metamer set $\mathcal{M}(\rho)$ for a change from $\mathbf{R}$ to $\mathbf{X}$. We then say that a 3×3 linear transform matrix $\mathbf{T}$ is *error-less* if and only if it maps each RGB ρ_i to an XYZ χ_i inside it's corresponding feasible set:

$$\rho_i \mathbf{T} \in \mathcal{F}(\rho_\mathbf{i}; \mathbf{R} \rightarrow \mathbf{X}). \tag{7}$$

In previous work[3] we presented a solution to a less constrained problem. We derived a method that solves for the linear least squares (LSQ) mapping between chromaticities corresponding to RGBs and XYZs, that is error-less (maps always inside the metamer set) while minimising the least squares error between true XYZs and mapped estimates. The reasoning behind solving such a transform is justified by considering intensity variation in a scene, that does not correspond to surface variation. Thus in the case of non-uniform illumination, this is the correct approach. At the same time it is also weaker than solving for an error-less LLSQ between RGBs and XYZs themselves. Such a transform would be relevant under more controlled conditions, where any variation in intensity would correspond to variation in the corresponding surfaces as well.

Such a constrained formulation of the problem is very strict. In fact, it can happen that a solution is infeasible due to a large number of saturated samples in the set examined for example. However we can solve this problem simply by considering a small level of noise and solving for a *paramer set*[5], instead of the metamer set. Due to the convexity of the metamer set, this turns out to be simple.

When solving for the metamer set, we start with a single ρ and find all reflectances that map to it. In solving for the paramer set instead we start with a small region in RGB space. In order to consider a uniform region, we find S uniformly sampling vertices of a sphere around ρ, of a diameter ε corresponding to the level of noise we wish to take into account. Each of these extreme vertices ρ_i has a corresponding metamer set $\mathcal{M}(\rho_i)$. Since colour image formation preserves convex combinations, i.e. a convex combination of reflectances $\alpha r_i + (1-\alpha) r_j$ results in the same convex combination of responses $\alpha \rho_i + (1-\alpha)\rho_j$, it suffices to find the S metamer sets, and their convex hull defines the paramer set, written as:

$$\mathcal{P}(\rho; \varepsilon) = CH(\mathcal{M}(\rho_1), \ldots, \mathcal{M}(\rho_S)) \tag{8}$$

where $CH()$ is a function computing the convex hull. $\mathcal{P}(\rho; \varepsilon)$ then corresponds to the volume of reflectances that map inside a sphere centred at ρ of diameter ε. The feasible region of corresponding χ's in the space of CIE XYZ space is written as:

$$\mathcal{P}(\rho; \varepsilon) \times \mathbf{X} = \mathcal{F}^*(\rho; \mathbf{R} \rightarrow \mathbf{X}) \tag{9}$$

and finding the linear least-squares error-less colour correction, with added noise, amounts to finding a linear transform minimising Eq. **??** subject to the constraint:

$$\rho_i \mathbf{T} \in \mathcal{F}^*(\rho; \mathbf{R} \rightarrow \mathbf{X}) \tag{10}$$

Thus, an error-less linear least squares transform is a matrix $\mathbf{T}$ that minimises least squares error between a set of RGBs in the $M \times 3$ matrix $\mathbf{G}$ and a set of known corresponding XYZs in the $M \times 3$ matrix $\mathbf{Y}$, written as:

$$\min_{\mathbf{T}} \parallel \mathbf{GT} - \mathbf{Y} \parallel_2 \tag{11}$$

where $\| \cdot \|_2$ is the L2 Euclidean distance metric, subject to the constraint that each RGB is mapped to an XYZ from the feasible region (Eq. 10).

4. Experiments

To test the error-less colour correction, we conducted a set of experiments comparing it to the standard linear least squares method. We created synthetic RGBs by way of the colour image formation equations in Eq. 1, using four reflectance data sets: the Dupont set of 120 samples [], the Westland set of 404 samples [13], a set of 134 saturated reflectances and the 426 Munsell data set [10]. To calculate RGBs for each of these sets, we used CIE illuminant D65, a near-neutral daylight and a set of camera spectral sensitivities, both plotted in Fig. 1. The dimension of the linear model used for each of the data sets is different, and such that it represents the reflectance data sets with high accuracy ($> 99\%$), and is the smallest dimension that is sufficient for the existence of the error-less transform. For the Dupont and Munsell set this is 5D, the saturated set 6D and for the natural reflectances of the Westland data set this is 7D.

We evaluated both colour correction mappings in terms of the mean, median and maximum CIE Lab ΔE, an error

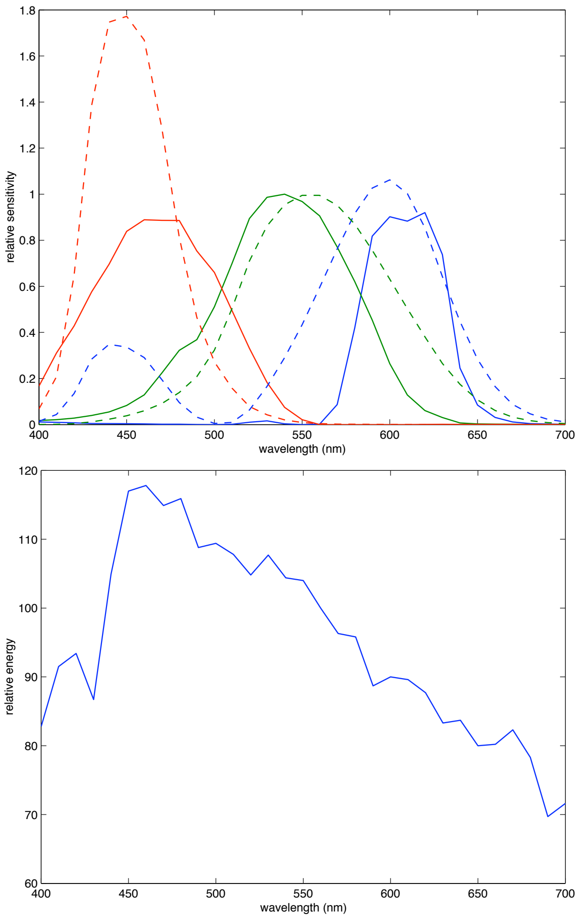

Figure 1: The illuminant spectral power distribution of CIE illuminant D65, the spectral sensitivities of a camera (solid) and the CIE colour matching functions (dashed).

measure in the near-uniform CIE Lab colour space, such that 1 ΔE corresponds approximately to a just noticeable difference. Taking all samples from each of the sets, there does not exist an error-less colour correction if the exact metamer sets are used in the optimisation. Instead we used the paramer set definition from Eq. **??**, using 0.5% error.

Tab. 1, summarises the results over all samples for each of the four sets in turn, in terms of the mean, median and maximum error.

	ELCC			LSQ		
	mean	med	max	mean	med	max
dup	3.8	2.9	17.0	4.1	1.8	19.9
sat	4.1	3.2	14.8	4.8	3.7	17.4
wst	2.6	2.1	16.6	2.8	2.0	24.4
mun	1.9	1.2	14.1	1.9	1.2	19.2

Table 1: Overall colour correction results in terms of CIE ΔE statistics.

From Tab. 1 we see that ELCC is at least as good as LSQ in terms of mean and maximum statistics, and is at least as good in all but one case in terms of median statistics. We learn from these results straightaway that error less colour correction improves the actual correction error over all samples, and most importantly reduces the maximum error. Next however, we look at only those samples which are infeasible in the LSQ case, in other words, those samples which render LSQ not error-less.

		ELCC			LSQ		
	%	mean	med	max	mean	med	max
dup	14.5	7.8	8.6	11.1	11.7	11.7	19.9
sat	10.5	8.0	8.3	10.4	13.2	13.5	17.4
wst	2.8	9.8	9.9	16.6	13.9	13.5	24.4
mun	0.9	9.9	8.8	14.1	11.5	9.2	19.2

Table 2: Colour correction results in terms of CIE ΔE statistics of the infeasible samples.

Tab. 2 shows more significantly the difference that error-less colour correction makes. In particular we can see that infeasible samples are in general samples which are mapped with higher error compared to the average case in Tab. 1. At the same time we also learn that for these, saturated samples, ELCC results in radically smaller error compared to LSQ.

While LSQ is an optimal method in terms of error (note: neither LSQ nor ELCC minimise error in terms of CIE ΔE), it does not have a notion of feasibility. The addition of the feasibility constraints to LSQ and thus solution of the ELCC transform may seem as a constraint on the performance of LSQ too, however since by adding the constraints we are adding more information to the system, we in fact see an improvement of performance. The overall similarity of the two methods is also coherent, as in a large proportion (see Tab. 2), the LSQ method is error-less too.

5. Conclusions

In this paper we looked at the problem of colour correction. Due to the nature of colour image formation and the fact that device spectral sensitivities are rarely linearly related to CIE colour matching functions, metamerism is present. Traditionally, colour correction methods avoid solving for metamerism and instead minimise error instead through a simple linear transform.

Instead, here we solved for a colour correction method that is built on the framework of metamerism, while maintaining the simplicity of a linear transform and minimising error in the least squares sense. Since error in colour correction arises due to metamerism we defined the concept of error-less transforms. We call a transform error-less if the RGB-XYZ pair it puts in correspondence is such that a physically realisable surface reflectance exists that maps to this pair. We have shown how to solve for this error-less transform.

In a set of experiments we compared the traditional linear least squares method to our novel approach and found that on average our approach is at least as good as least squares and importantly reduces maximum error. Furthermore, looking at the most problematic samples, saturated colours, our approach reduces both average and maximum error dramatically.

6. Acknowledgements

This research has been supported by the EPSRC grant number GR/R65978/01.

References

[1] Jozef B. Cohen and Wiliam E. Kappauf. Metameric color stimuli, fundamental metamers, and Wyszecki's metameric blacks. *American Journal of Psychology*, 95:537–564, 1982.

[2] Graham D. Finlayson and Peter Morovic. Metamer constrained colour correction. In *IS&T/SID 7th Color Imaging Conference, Scottsdale, Arizona*, 1999.

[3] Graham D. Finlayson and Peter Morovic. Error-less colour correction. In *17th International Conference on Pattern Recognition, Cambridge, UK*, pages ??–??, 2004.

[4] John B. Fraleigh and Raymond A. Beauregard. *Linear Algebra*. Addison-Wesley Publishing Company, 1990.

[5] Rolf G. Kuehni. Metamerism, exact and approximate. *Color Research and Application*, (8):192, 1983.

[6] Edwin H. Land. The retinex theory of color vision. *Scientific American*, 1977.

[7] Lawrence T. Maloney. Evaluation of linear models of surface spectral reflectance with small numbers of parameters. *Journal of the Optical Society of America A*, 3(10):1673–1683, October 1986.

[8] Lawrence T. Maloney and Brian A. Wandell. Color constancy: a method for recovering surface spectral reflectance. *Journal of the Optical Society of America A*, 3:29–33, 1986.

[9] Peter Morovic. *Metamer Sets*. PhD. Thesis, University of East Anglia, School of Information Systems, 2002.

[10] J. P. S. Parkkinen, J. Hallikanen, and T. Jaaskelainen. Characteristic spectra of munsell colors. *Journal of the Optical Society of America A*, 6:318–322, 1989.

[11] B. Smith, C. Spiekermann, and R. Sember. Numerical methods for colorimetric calculations: Sampling density requirements. *Color Research and Application*, 17(6):394–401, 1992.

[12] Michael J. Vrhel. *Mathematical Methods of Color Correction*. PhD. Thesis, North Carolina State University, Department of Electrical and Computer Engineering, 1993.

[13] Stephen Westland, Julian Shaw, and Huw Owens. Colour statistics of natural and man-made surfaces. *Sensor Review*, 20(1):50–55, 2000.

[14] Günther Wyszecki and Wiliam S. Stiles. *Color Science: concepts and methods, quantitative data and formulae. (2nd ed.)*. John Wiley New York, 1982.