

Using a Simple Colour Pre-filter to Make Cameras More Colorimetric

Graham D. Finlayson*, Yuteng Zhu*, Han Gong
 School of Computing Sciences, University of East Anglia, Norwich, UK
 ** Joint First Authors

Abstract

The idea of placing a coloured filter in front of a camera to make it more colorimetric has been previously proposed. However, this prior art approach sought to increase the dimensionality of the capture — i.e. to take an image with and without a filter — rather than to change the spectral characteristics of the sensor itself.

In this paper, we set out a new method for finding the filter that, in a least-squares sense best achieves the Luther condition. That is, the filter multiplied by the camera spectral sensitivities is ‘almost’ a linear combination from the colour matching functions of the human visual system. We show that for a given sensor set the best filter and best linear mapping can be found together by solving an alternating least-squares problem.

Experiments demonstrate that placing an optimal filter in front of a camera can result in a dramatic improvement in its ability to see the world colorimetrically.

1. Introduction

The spectral sensitivities of a camera are designed to satisfy many objectives including manufacturability and the need to minimise the conspicuity of noise when the raw RGBs are transformed for display. As a consequence, most cameras do not ‘see’ like we do. That is the RGBs a camera measure are not a linear transform from colour matching triplets. More formally, a camera ‘sees’ like we do – with a simple linear colour correction – if and only if its spectral sensitivities are a linear combination of the corresponding human visual system matching curves (or equivalently a linear combination of the XYZ colour matching functions).

In the literature, it has been proposed that to make a camera more colorimetric we could take several pictures under different coloured filters [1]. This generates a higher dimensional representation with respect to which a linear combination can better predict what we see (e.g. better predict XYZs). Alternatively, in the field of fine art, multispectral and hyperspectral imaging systems are widely used to obtain great colour fidelity [2].

The multiple-image approach is problematic because it necessitates a long capture process where the user has to change filters between image captures. If nothing else, the captured images are not in registration, which is a problem in itself. As for the multispectral approach, it is characterized by longer capture time and much more expensive imagers.

Of course, a more pragmatic way to proceed – and the one that is implemented in all camera processing pipelines – is simply, given the RGBs for a given camera do, by mathematical means, find a correction transform that brings them as close as possible to the desired XYZs (or display RGBs). The preponderance of methods take a regression approach and involve either linearly cor-

recting RGBs or linearly correcting a polynomial-type expansion of RGBs [3, 4, 5, 6, 7, 8, 9]. Other mappings include look-up-tables [10], and artificial neural networks [11, 12].

In this paper, we propose that finding the filter to make a camera more colorimetric should be chosen in tandem with the best colour correction transform taking RGBs to XYZs. Our method is built directly on top of the Luther condition [13]. If $\underline{Q}(\lambda)$ denotes the vector of red-, green- and blue-spectral sensitivities and $\underline{\chi}(\lambda)$ denotes the XYZ colour matching functions then an explicit Luther condition inspired formulation of the filter design problem can be modeled as:

$$\min_{f(\lambda), M} \| f(\lambda)[M^t \underline{Q}(\lambda)] - \underline{\chi}(\lambda) \| \quad (1)$$

where $f(\lambda)$ defines the spectral transmittance of the colour filter, and M is a fixed 3×3 linear transformation matrix.

Figure 1 illustrates Equation 1 with respect to the colour matching experiment. Here the goal of the observer is to adjust the intensities of the three primary lights (R, G and B) to make a colour that looks identical to the bipartite of the test light. In this example, R, G and B are set to respectively the intensities ρ, γ, β in order to make a visual match. In Figure 1 we seek a filter, shown in red in front of the camera, so that the camera RGB values multiplied by matrix M record the same triplets as the observer does under the same viewing geometry. That is the camera ‘sees’ mixture of the intensities a human observer would use to make a match.




Figure 1: We try and find a filter such that the RGBs a camera measures (after a 3×3 correction matrix) are the same as the colour mixtures a human observer would make to find a match. That the filter is red here is for illustration (the best filter colour is not a priori specified). The position of the observer and the camera should be the same in the real match.

In this article, we show how to solve the minimization of Equation 1 to find the best coloured filter to make a given camera colorimetric. To quantify the improvement in colour measurement for a given camera afforded by a colour filter we calculate the Vora-Value before and after correction [14]. Remarkably, some cameras

become ‘nearly’ colorimetric. We also evaluate their performance via a colour correction experiment using measured data e.g. ΔE_{ab}^* [15] with and without a filter.

In section 2 we review the background in colour filter design. Section 3 presents our method. In section 4 we report on our experiments. The paper concludes in section 5.

2. Background

There are three important physical variables to consider in image formation. First, there are the surface reflectance properties of the objects in the scene. Second, we must consider the illumination or illuminations under which a scene is viewed. Lastly, the spectral characteristic of the sensors is an important variable. The interaction of surface, light and sensor is elucidated in the equation:

$$\rho_k = \int_{\omega} C(\lambda) Q_k(\lambda) d\lambda, k \in \{R, G, B\} \quad (2)$$

The colour signal $C(\lambda)$ is a combination of light and surface information. An integrated response is calculated for each of three sensors $Q_k(\lambda)$ (usually, short-, medium- and long-wave sensitive mechanisms or R, G and B). The integral is taken over the visible spectrum ω .

It will be useful to recast Equation 2 in the discrete domain (matrix-vector representation). Sampling spectra allows us to write spectral quantities as vectors. We rewrite Equation 2 as

$$\underline{\rho}_k = \underline{C}^t \underline{Q}_k \quad (3)$$

where \underline{C} is the spectral colour signal spectrum sampled at measuring intervals (e.g. giving 31 values over the visible spectrum 400 to 700 nm through 10 nm interval). \underline{Q}_k is the discrete representation of the k^{th} sensor. Advantageously, in the discrete domain it is easy to write the RGB response as a single response. Equation 4 denotes the RGB and XYZ responses to the same stimulus (where Q and χ denote the 31×3 matrices of sensitivities).

$$\underline{\rho}^t = \underline{C}^t Q, \quad \underline{x}^t = \underline{C}^t \chi \quad (4)$$

A necessary and sufficient condition for a camera to measure the world colorimetrically (that is, for its RGBs to be a linear correction from XYZs) is for Q and χ to be a linear correction from one another, the so-called Luther condition:

$$if \ Q = \chi M \Rightarrow \underline{x}^t = \underline{\rho}^t M^{-1} \quad (5)$$

The Luther condition is strong. Arguably, for a fixed viewing illuminant, because the reflectance characteristics tend to be smooth, it may be possible to measure colorimetrically (to some small degree of error) without the Luther condition being met. Indeed, the colour fidelity we see on our smart phones is reasonably good. Let C denotes a set of N colour signals then the corresponding response sets of RGBs and XYZs can be written as:

$$P = C^t Q, \quad X = C^t \chi \quad (6)$$

We can treat P and X as data matrices and the best colour correction can be written as a least-squares regression.

$$\min_M \| PM - X \| \Rightarrow M = [P^t P]^{-1} P^t X = P^+ X \quad (7)$$

The ‘book formula’ shown in the right hand side of Equation 7 is called the Moore-Penrose inverse. Note the superscripts $+$ and t denote pseudoinverse (also see in Equations 10 and 11) and matrix transpose respectively.

Under the assumption that reflectances are sufficiently well modeled by a 3D linear model [16, 17], then Equation 7 yields an M that supports perfect colour correction. However, there are many examples of cameras recording the wrong colours (i.e. the reproduction formed after colour correction has the wrong colours compared to what a human observer perceives when viewing the same scene). Further, the tolerance for measurement (e.g. to decide that one surface is the same colour as another) is tighter and a linear correction does not suffice.

The quality of colour reproduction (and linear models) was the motivation for Farrell and Wandell’s [1] colour filter approach. They proposed placing a filter between the object and the measurement device (scanner or camera). Then two images are captured, once with and once without the filter. In this way six measurements are made per pixel. Mathematically, the data matrix becomes:

$$P_{N \times 6} = C^t [Q \quad Q^f] \quad (8)$$

Here the superscript f denotes the sensitivities of the camera multiplied by a colour filter $f(\lambda)$. Note that for this case the least-squares matrix M in Equation 7 is a 6×3 matrix. The ‘best’ filter was chosen empirically (by choosing the filter from a set of commercially available choices that supports the best colour measurement).

In previous work, Vora and Trussell [14] developed a simple ‘goodness’ score – a number between 0 and 1 — to assess how closely a sensor set measures colours compared with the human visual system (unity means the sensor set is colorimetric while a zero measures nothing in common). In [18, 19] they used this Vora-Value as part of an optimization scheme to find the sensor set that was most colorimetric. Perhaps counter-intuitively these sensors do not meet the Luther condition.

Rather, their problem formulation asks what are the best filters that can measure the world under multiple lights with a simple linear correction, which can map the recorded RGBs (for a given light) to corresponding measurements under a reference light. This approach is intriguing not least because we ourselves do not see in this way. Under different lights some colours do look different! In this paper we will be looking at filter design as a way of best trying to measure the same XYZs we see for a given illuminant condition.

3. Filter design

Let us rewrite Equation 1 in the discrete domain. Specifically, we propose that the filter that best matches the Luther condition can be found by minimizing:

$$\min_{D, M} \| DQM - \chi \| \quad (9)$$

Here Q and χ are respectively $N \times 3$ matrices capturing the camera spectral sensitivities and the XYZ colour matching functions. N refers to the sampling number across the visible range. M is a 3×3 correction matrix. D is an $N \times N$ diagonal matrix. Mathematically, D_{ii} , the i^{th} diagonal term in D multiplies the i^{th} row in Q . That is, the three camera sensitivities at the i^{th} wavelength are multiplied

by the same value. D in the discrete domain is a physically accurate model of $f(\lambda)$ in Equation 1.

Unfortunately, there is no closed form solution to Equation 9. Rather we solve for D and M using a technique called Alternating Least-Squares (ALS). The algorithm is shown below where $\|\bullet\|_F$ denotes the Frobenius norm (also see in Equation 11).

Algorithm 1 Filter and Linear Matrix Estimation by ALS

- 1: $i = 0, \quad Q^0 = Q;$
 - 2: **repeat**
 - 3: $i = i + 1;$
 - 4: $\min_{D^i} \|D^i Q^{i-1} - \chi\|_F;$
 - 5: $\min_{M^i} \|D^i Q^{i-1} M^i - \chi\|_F;$
 - 6: $Q^i = D^i Q^{i-1} M^i;$
 - 7: **until** $\|Q^i - Q^{i-1}\|_F < \varepsilon$
 $D = \prod_i D^i$ and $M = \prod_i M^i;$
-

In ALS we make an initial guess for the filter by ignoring the linear correction term M . See step 4. Then we hold the filter fixed and solve for M (step 5). Taken together we estimate a new ‘corrected’ (partial solution) spectral sensitivities. We repeat this process until convergence (the method is guaranteed to converge[20]).

In Figure 2, top left we show the CIE1931 2° colour matching functions. Top right shows the Canon 50D spectral sensitivities. Using the above algorithm, we solve for the filter shown in the middle of the diagram. We arrive at the comparison shown in the middle of the figure (solid, XYZ colour matching functions and dotted the approximation) by multiplying this filter (shown at the bottom) and linear fitting to the XYZ colour matching functions. The closeness of the two sets of curves is remarkable.

Cautionary Remarks

First, the reader will note the filter shape is rapidly varying especially at the long-wavelength end and so likely not easy to manufacture. Further, across most of the wavelengths the filter absorbs most of the light. Therefore, if this filter is used the resulting image would either be noisy (less light captured) or the exposure time would need to be significantly increased. The role of noise and the shape of the filter will be addressed in future research.

Second, we draw attention to the paper [21] where the cone sensitivities are discovered by modelling colour matching functions as optical pre-filtering and linear combination of the cone absorptances. Essentially, we deploy the mathematics developed in [21] here. However, in [21] it is known a priori that a good solution was possible. Further, constraints have to be placed on the shape of the pre-filtering (since much is known about the ocular media, lens etc.).

The emphasis of our own work is to solve for the best pre-filter for the application of making a camera more colorimetric. In contrast to the previous work, there is no a priori reason why good pre-filtering should exist. But, in our favour here, we need not concern ourselves with the shape of the pre-filtering filter (we do not need to consider biological plausibility).

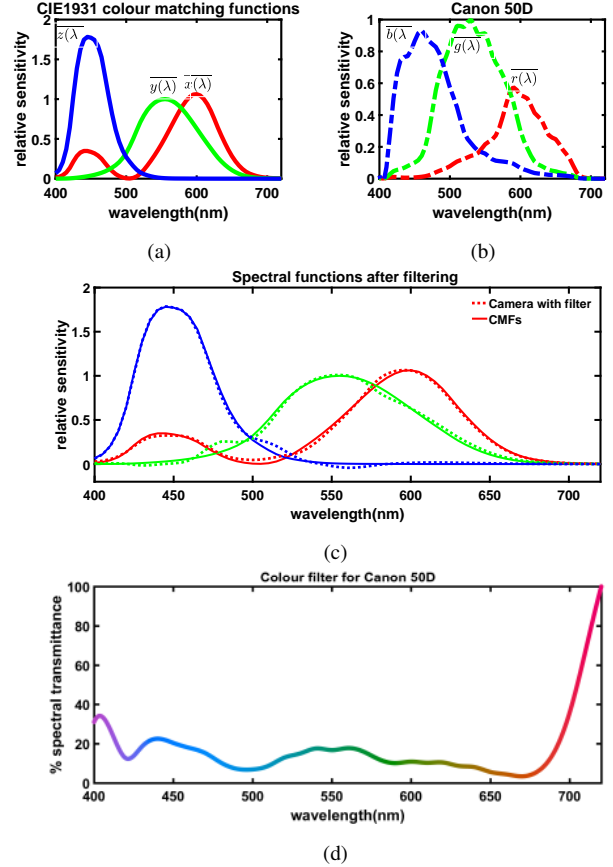


Figure 2: Spectral distributions. (a) CIE1931 CMFs of 2° standard observer, (b) camera spectral sensitivity functions of Canon 50D normalized at 560nm, (c) reconstructed camera response functions with designed filter and linear mapping (dotted lines) comparing to the reference CMFs (solid lines), (d) spectral transmittance of colour filter solved for Canon 50D over the measuring range.

4. Experiments
Spectral Evaluation

In order to measure the quality of filtered camera respect to the desired CMFs, we adopt the measures of Vora-Value [14] and normalized spectral root-mean-square error (NRMSE).

The Vora-Value is defined in Equation 10 as

$$v = \frac{\text{Trace}(\chi\chi^+ Q Q^+)}{3} \tag{10}$$

The *Trace* function sums up the terms along the diagonal of a matrix. We refer the reader to [14] where the meaning and derivation of this equation are explained.

The Vora-Value returns a number between 0 and 1 meaning respectively not colorimetric at all and 100% colorimetric. Pragmatically, a sensor system that has a Vora-Value above 0.9 captures colours that a 3×3 matrix can correct to XYZs to a tolerable perceptual error (e.g. the colours will look mostly correct and the visual error will be acceptable). A Vora-Value above 0.99 is indicative of a camera that is almost colorimetric i.e. to all practical purposes will sense the world like we do.

The normalized spectral NRMSE between corrected camera

sensitivity functions and desired CMFs is defined as:

$$NRMSE = \frac{\|\chi - QQ^+ \chi\|_F}{\|\chi\|_F} \quad (11)$$

The performance of our proposed filter design method is evaluated with respect to [22], a set of the measured spectral sensitivities of 28 commercial cameras. A large component of the set comprises 9 Canon and 10 Nikon cameras. We will pay close attention to these subsets. As a final note the data is given from 400 to 720 nm through 10 nm interval so we use 33 sample points in our calculations.

For each camera, a specific filter and a linear transformation were calculated using Algorithm 1 developed in the last section. Fig. 2(d) displays an example of colour filter calculated for Canon 50D. The relative spectral transmittance values are positive in the range between 0 and 100%. The Canon camera sensitivities multiplied by the filter with a linear combination applied are shown in dotted lines in Fig. 2(c). The filter corrected sensitivities are almost the same as the XYZ colour matching functions.

In terms of Vora-Value the Canon 50D which was originally 0.950 becomes 0.992. The NRMSE improves from 0.250 to 0.052, reduced by nearly 80%. By applying this colour filter and linear mapping, Canon 50D camera becomes almost colorimetric.

In Table 1, we look at the Vora-Value and NRMSE performance before and after a filter is added (with linear correction). The first three rows of the table show the performance of 3 cameras which post filter correction have the maximum, median and minimum Vora-Values. Then we look at the average performance of the Canon and Nikon cameras as separate subgroups. Finally, the average performance over the whole data set is summarized. We repeat this methodology for NRMSE.

Table 1: Performance of camera systems with colour filter versus the original functions using Vora-Value and NRMSE.

Camera	Vora-Value		NRMSE	
	After	Original	After	Original
Canon50D	0.992	0.950	0.052	0.250
NikonD40	0.947	0.924	0.104	0.305
NikonD90	0.941	0.922	0.113	0.317
Canon group	0.987	0.938	0.062	0.285
Nikon group	0.944	0.921	0.108	0.316
Whole data	0.961	0.918	0.091	0.316

Spectrally, all cameras are improved to some extent. About 50% of the cameras lend themselves to significant improvements in their colour measurement ability. This is an interesting result since it shows that within the range of manufacturable sensors there are sets which can be made much more colorimetric with the addition of a coloured filter. And these cameras could then be used in applications where accurate colour measurement is needed.

Taken as a group, Canon cameras can be filter-corrected to become colorimetric more readily than Nikon cameras. The corrected Canon cameras have Vora-Values of 0.972 at least, with an average value as high as 0.987 for the whole subset compared to 0.944 of Nikon subgroup.

Colour Correction Experiment

We now carry out a second experiment. Here we wish to evaluate how well the RGBs measured by a camera can be corrected

to match XYZs and how much this colour correction performance is improved when a coloured filter (designed by Algorithm 1) is placed in front of the camera. For each camera (normal and filter corrected) we calculate the RGBs under CIE D65 [23] for the SFU set of 1995 reflectances [24]. We then find the optimal least-squares linear correction to best map to ground-truth XYZs. We then calculate the average, median, and 95% quantile of ΔE_{ab}^* over the test data set.




Figure 3: Comparison the performance of colour correction with and without colour pre-filtering with linear mapping and direct linear mapping. The blue squares represent direct linear correction, while the red squares show the results by colour pre-filtering with linear correction. On the horizontal axis, there lists three indicators respectively mean, median and 95% quantile of the colour differences for the SFU reflectance data set under CIE D65. The trend of the three measures is consistently dropping showing the reduction of colour error.

The results are presented as shown in Fig. 3. For each error statistic we show the means (i.e. the mean of the means calculated over different cameras and the mean of the medians and the mean of the 95% quantiles). We also show the spread in the error defined to be 2 standard deviations. Whichever statistic is employed, there is over 20% on average boost in performance. However, when reviewing the cameras individually, a few cameras like Sony Nex5N and Nikon D40, are not greatly improved by pre-filtering.

5. Conclusion

In this paper, we develop a method to find the best coloured filter such that when placed in front of a camera its new effective spectral sensitivities best match the Luther condition (become most colorimetric). Experiments show that our method often makes a camera ‘almost’ colorimetric which means that they can, in principle, be used for precise colour measurement.

References

- [1] J. Farrell and B. Wandell, U. S. Patent No. 5479524, Washington, DC: U.S. Patent and Trademark Office, 1995.
- [2] H. Liang, “Advances in multispectral and hyperspectral imaging for archaeology and art conservation,” *Appl. Phys. A*, vol. 106, no.2, pp. 309-323, 2012.
- [3] G. D. Finlayson and M. S. Drew, “Constrained least-squares regression in color spaces,” *J. Electron. Image*, vol. 6, no. 4, pp. 484-493, 1997.
- [4] G. Hong, M. R. Luo, and P. A. Rhodes, “A study of digital camera colorimetric characterization based on polynomial modeling,” *Color Res. Appl.*, vol. 26, no. 1, pp. 76-84, 2001.
- [5] S. H. Lim and A. Silverstein, “Spatially varying color correction matrices for reduced noise,” *Imaging Systems Lab., HP Lab., Palo Alto, CA, USA, Tech. Rep. HPL-2004-99*, 2004.

- [6] G. D. Finlayson, M. Mackiewicz, and A. Hurlbert, "Color correction using root-polynomial regression," *Trans. on Image Processing*, vol. 24, no.5, pp. 1460-1470, 2015.
- [7] G. D. Finlayson, D. M. Mohammadzadeh, and M. Mackiewicz, "The alternating least squares technique for non-uniform intensity color correction," *Color Res. Appl.*, vol. 40, no.3, pp. 232-242, 2015.
- [8] M. Mackiewicz, C. F. Andersen, and G. Finlayson, "Method for hue plane preserving color correction," *J. Opt. Soc. Am. A*, vol. 33, pp. 2166-2177, 2016.
- [9] J. Vazquez-Corral, D. Connah, and M. Bertalmio, "Perceptual color characterization of cameras," *Sensors*, vol. 14, no. 12, pp.23205-23229, 2014.
- [10] P.-C. Hung, "Colorimetric calibration in electronic imaging devices using a look-up-table model and interpolations," *J. Electron. Image*, vol. 2, no. 1, pp. 53-61, 1993.
- [11] T. Cheung and S. Westland, "Colour camera characterisation using artificial neural networks," in *10th Color Imaging Conference*, vol. 4, pp. 117-120, 2002.
- [12] X. Li, "A new color correction model for based on BP neural network," *Adv. Inf. Sci. Service Sci.*, vol. 3, no. 5, pp. 72-78, 2011.
- [13] R. Luther, "Aus dem Gebiet der Farbreizmetrik," *Zeitschrift fur Technische Physik*, vol. 12, pp. 540-558, 1927.
- [14] P. L. Vora and H. J. Trussell, "Measure of goodness of a set of color scanning lters," *J. Opt. Soc. Amer. A*, vol. 10, no. 7, pp. 1499-1508, 1993.
- [15] G. Wyszecki and W. S. Stiles, *Color Science: Concepts and Methods, Quantitative Data and Formulae*, New York: Wiley, 1982.
- [16] D. H. Marimont and B. A. Wandell, "Linear models of surface and illuminant spectra," *JOSA A*, vol. 9, no.11, pp. 1905-1913, 1992.
- [17] G. D. Finlayson, H. Gong, and R. Fisher, "Color Homography: theory and applications," *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 2017.
- [18] P. L. Vora and H. J. Trussell, "Mathematical Method for Designing a Set of Colour Scanning Filters," *Proc. SPIE 1912, Color Hard Copy and Graphic Arts II*, 1993.
- [19] P. L. Vora and H. J. Trussell, "Mathematical methods for the design of color scanning lters," *IEEE Trans. Image Processing*, pp. 312-320, 1997.
- [20] T. Zhang and G. H. Golub, "Rank-one approximation to high order tensors," *SIAM Journal on Matrix Analysis and Applications*, vol. 23, no. 2, pp. 534550, 2001.
- [21] C. F. Andersen, G. D. Finlayson, and D. Connah, "Estimating individual cone fundamentals from their color-matching functions," *J. Opt. Soc. Am. A*, vol. 33, pp. 1579-1588, 2016.
- [22] J. Jiang, D. Liu, J. Gu, and S. Ssstrunk, "What is the space of spectral sensitivity functions for digital color cameras?" *Applications of Computer Vision, IEEE Workshop*, pp. 168-179, 2013.
- [23] N. Ohta and A. Robertson, *Colorimetry: fundamentals and applications*, John Wiley & Sons, 2006.
- [24] K. Barnard, L. Martin, B. Funt, and A. Coath, "A data set for color research," *Color Res. Appl.*, vol. 27, no. 3, pp. 147-151, 2002.

Author Biography

Graham D. Finlayson obtained his BSc in Computing Science from the University of Strathclyde in 1989 and his MSc and PhD degrees from Simon Fraser University in 1992 and 1995 respectively. Prof. Finlayson is interested in Computing how we see and his research spans computer science (algorithms), engineering (embedded systems) and psychophysics (visual perception).

Yuteng Zhu received her BS and MSc in Optical Engineering from Zhejiang University in 2010 and 2014 respectively. She is currently a PhD candidate supervised by Prof. Graham Finlayson in Computer Science at the University of East Anglia. Her work has focused on Computational Photography.

Han Gong received his Ph.D. from the University of Bath. He is currently a lecturer at the University of East Anglia. His research interests include Computer Vision and Computational Photography.