

# Extended Linear Color Correction

Graham D. Finlayson<sup>1</sup>, Garrett M. Johnson<sup>2</sup>; <sup>1</sup> The University of East Anglia, Norwich, UK; <sup>2</sup> Apple, Cupertino, USA

## Abstract

Color correction is often posed as a linear regression problem either from camera RGB to XYZ - where the aim is use a camera for color measurement - or to a display color space such as sRGB for image reproduction. While linear regression is simple and also ensures exposure independence, the mapping found through regressing RGB to XYZ is not optimal in terms of perceived color.

In this paper, we begin by observing that the best linear transform for mapping RGB to XYZ to minimize a color difference metric, such as CIE LAB, is not separable. In particular we show that the best fitted Y channel should be different depending on whether L\*, a\* or b\* error is minimized. Consequently, we develop an extended linear regression framework for CIE LAB where we solve for Y - we map RGB onto Y - three times - once for L\*, once for a\* and once for b\*. As in conventional regression we solve for X and Z only once.

Experiments demonstrate that compared to our new extended linear regression method the mean, 95% quantile and CIE LAB error afforded by simple linear least-squares is respectively 30%, 50% and 70% larger. Extended linear regression delivers leading color correction performance with fewer parameters than competing methods.

## Introduction

Camera sensors do not sample light like our own visual system. Not only are the spectral sensitivities of the R, G and B sensors different from our own cone sensors they are not linearly related, they do not meet the Luther conditions[15]. In direct consequence, there are colors that we see as the same that look different to the camera and vice versa. Metamerism is a real problem, so a camera cannot be used as an exact proxy for our own visual system.

There is however a practical and sustained research interest in *color correcting* measurements made by a camera to corresponding human visual system referred color coordinates. Examples of targets for color correction include cone responses[7], XYZ color triplets, e.g. [25], and perceptual coordinates such as CIE LAB, e.g. [4]. In computer vision, color correction is a key component of systems that seek to replicate the color visual judgements made by human observers, e.g. [5, 6]. Color correction for image reproduction involves mapping the RGBs measured by the cameras to the corresponding RGB coordinates that drive a display[2] or, in the case of hard copy, drive a printer. In image reproduction we aim to make the displayed image a facsimile of the scene as we ourselves would see it (with the caveat that

preference is also an important factor[22]).

Abstractly, the goal of all color correction algorithms is to map colors from one 3-dimensional coordinate frame to another. Here, we will consider how an  $(R, G, B)$  triplet,  $\underline{\rho}$  is mapped to the corresponding  $(X, Y, Z)$  coordinate,  $\underline{x}$ , and how this in turn is mapped to the CIE LAB [28] triplet:  $(L^*, a^*, b^*)$ , denoted  $\underline{L}$ . The CIE LAB color space has the property that the Euclidean distance between Lab triplets - called CIE LAB Delta E, or  $\Delta E$  - correlates tolerably with perceived color difference[28].

The color correction problem is often simply formulated as finding the *best* least-squares  $3 \times 3$  matrix  $M$  that maps a camera RGB  $\underline{\rho}$  to its corresponding XYZ,  $\underline{x}$  [10, 26]:

$$\hat{\underline{x}} = M\underline{\rho} \approx \underline{x} \quad (1)$$

There is a focus on linear color correction, in part, because if  $M$  delivers good color correction for a given RGB  $\underline{\rho}$  it continues to perform well when exposure changes and  $\underline{\rho} \rightarrow k\underline{\rho}$  (e.g.  $k = 0.5$  when the exposure time is halved). Further linear color correction is an easy problem to solve (see the background section) in that we can find the correction matrix in closed form and with little training data.

In the top row of Figure 1 we show, left, an input raw image with the corresponding sRGB output, right. These images are created through numerical integration. The underlying reflectance image is drawn from the Foster and Nascimento spectral image database[12], the illuminant is approximately D75 (hence the bluish cast in the raw image) and the camera spectral sensitivities were randomly selected from the 28 sensitivity sets in the RIT spectral sensitivity dataset [16] (a Point Grey camera was used). In panel d) we show the output of a least-squares calibration. The results are rather good drawing attention to the fact that simple linear color correction often works well. We zoom in on the region delimited by the white square in 2a). The corresponding crop for the least-squares correction is shown in 2b).

In Figure 1c), we also show the CIE Delta E error map for the difference between the least-squares fit image, 1d) and the ground truth output 1b). Here the colors Blue and Red respectively correspond to 0 and 20 Delta E (the in between colors linearly index Delta E, green is about 10). Significantly, it is accepted that small Delta Es are not readily perceptible in complex images and colors need to be reproduced with 3, 4 or 5 Delta E error to be noticeable[21, 23].

In this paper, we revisit linear color correction but add the constraint that we wish not only to well map RGBs to XYZs but also to find a good mapping with respect to the

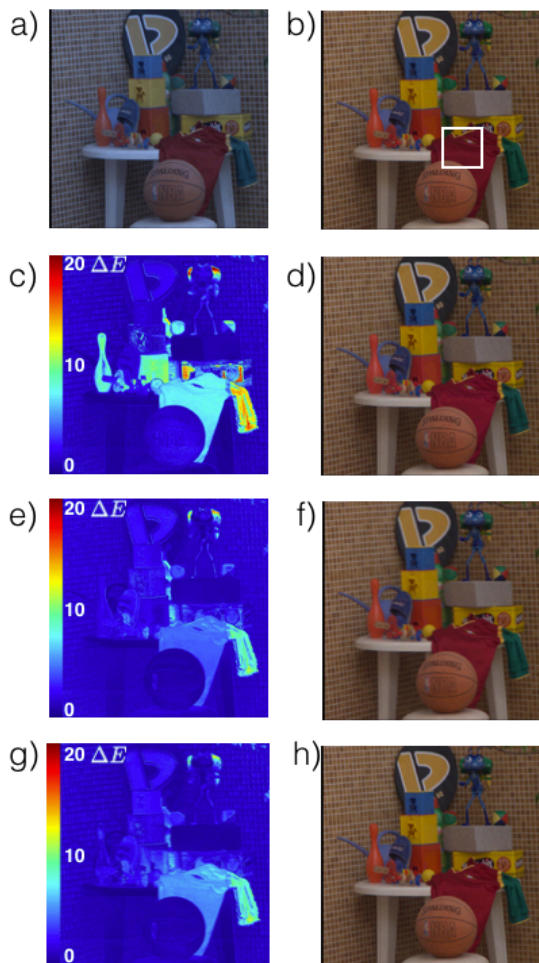


Figure 1. In a) a raw image (with a gamma of 0.5 applied to make it appear brighter) is shown next to the rendered counterpart b). Images d), f) and h) respectively show the output of a least-squares regression, the extended least-squares method developed in this paper and root-polynomial regression. Panels c), e) and g) show a pseudo color representation of CIE Delta E fitting error for imaged d), f) and h) compared with the ground truth b).




Figure 2. In a) we show the cropped region delimited by the white square from the ground truth output in Figure 1b). In b) we show the same crop for linear regression, the color is clearly a little in error. Finally, c) shows the output of our extended linear regression which has improved color fidelity.

CIE LAB color space. Our work begins with the following thought experiment. Suppose, for some reason, we are particularly interested in how well RGBs can be mapped to predict the corresponding X coordinates. We might mathematically implement this desire by by weighting the minimization to, in effect, penalize the fitting error in X (more than Y and Z). However, when we do this, it turns out that we arrive at the same least-squares color correction matrix  $M$ . This is true because linear least-squares is *separable*. The 3 coefficients in the first row of  $M$ , in Equation 1, map the RGB vector to the corresponding X coordinates. These coefficients are independent of - and can be solved independently from - those for predicting Y and Z.

If we wish to solve for  $M$  that minimizes CIE LAB color error then this separability property no longer holds. Indeed, in the CIE LAB formulae  $a^*$  depends on both of the X and Y coordinates and  $b^*$  depends on both of the Y and Z values.  $L^*$  depends only on Y. These dependencies dictate that the best linear regression formulated for CIE LAB we should singularly approximate X and Z but we must solve for Y three times (because  $L^*$ ,  $a^*$  and  $b^*$  all depend on Y). Counting parameters, our new *extended linear color correction* method has  $5 \times 3 = 15$  terms which compares to the 9 parameters in conventional color correction. However, linearity and exposure invariance is preserved.

Later, we report on a variety of experiments using real and synthetic data for many lights and cameras. In all cases moving from a simple to extended linear formulation of color correction delivers substantial benefits.

## Background

Often the mapping from RGB to XYZ is posed as a simple linear regression problem where we seek the  $3 \times 3$  correction matrix  $M$  that minimizes:

$$\min_M \|M\mathcal{P} - \mathcal{X}\| \quad (2)$$

where  $\mathcal{P}$  and  $\mathcal{X}$  respectively denote  $3 \times N$  camera RGB and XYZ color measurements for  $N$  surfaces. In (2),  $\|\cdot\|$  denotes the Frobenius norm (the square root of the sum of squares). The best  $3 \times 3$  matrix  $M$  is found in closed-form using the Moore-Penrose inverse,

$$M = \mathcal{X}\mathcal{P}^t[\mathcal{P}\mathcal{P}^t]^{-1} \quad (3)$$

One property this *least-squares* fit is separability. If we are interested in the linear combination of RGB that best approximates  $\mathcal{X}_i$  (the  $i$ th the row of the XYZ response matrix) then  $M_i$  (the  $i$ th row of  $M$ ) is simply equal to:  $\mathcal{X}_i\mathcal{P}^t[\mathcal{P}\mathcal{P}^t]^{-1}$ . In mathematical optimization this kind of separability is rare which is actually one of the reasons that linear least-squares is so commonly used.

A linear color correction matrix can also be derived using arguments based on the physical dimensionality of the

light and surfaces[18, 27, 19] or by adopting prior information about the assumed covariance structures of color spectral signals in the natural world[9]. Of direct relevance to this paper, it is also possible to formulate a minimization to find the  $M$  that minimizes error in a perceptually referenced colorspace such as CIE LAB.

We, abstractly, denote the mapping from XYZ to CIE LAB is a vector function  $\underline{f} : \mathbb{R}^3 \rightarrow \mathbb{R}^3$ . So, we formulate the linear regression problem to minimize Delta E error as:

$$\min_M \sum_{i=1}^N ||\underline{f}(M\rho_i) - \underline{f}(\underline{x}_i)||^2 \quad (4)$$

where, respectively,  $\rho_i$  and  $\underline{x}_i$  denote the  $i$ th columns of  $\mathcal{P}$  and  $\mathcal{X}$ . The matrix  $M$  can be found through a gradient descent type search[4] (local minimum) or Eq. (4) can, in a colorspace sense, be *locally linearised* and a global minimum found[8].

Another useful property of linear color correction for photography applications is that it is independent of exposure: if  $M$  optimally maps  $\mathcal{P}$  to  $\mathcal{X}$  then  $M$  also optimally maps  $k\mathcal{P}$  to  $k\mathcal{X}$  where  $k$  is a scalar modelling brightness change. When color correction does not scale with brightness the same physical object color can be mapped to significantly different image colors e.g. if the same surface is viewed in a dark and light part of the same scene the corresponding image colors can have different hue and saturation. Further, the magnitude of the color shift can, on occasion, be surprisingly large[11]. Exposure dependent methods of color correction which includes polynomial regression[14], lattice regression[13] and neural net based methods[24] are not considered further in this paper.

Extensions to linear color correction which **do** scale with intensity have been a recent focus of development in color correction research. Andersen[1] divided chromaticity, centered at the whitepoint into  $k$  unbounded triangular regions. Per region, the chromaticities at the boundary and the white-point uniquely defined a 3x3 correction matrix. Further, by construction, the method implements a continuous color mapping. The method is extended in [17] to allow the bounding chromaticities to be chosen through optimization. The idea of optimizing chromaticity and intensity separately is studied in [20] where the chromaticity mapping is implemented as a general lookup-table. While polynomial regression is not intensity preserving a variant, called root-polynomial regression, was developed that does scale with exposure[11].

## Extended Linear Color Correction

In calculating CIE LAB, the XYZ values are first divided by the XYZ for a white surface. The CIE LAB equation is recapitulated in Equation (5) below however to ease the notation, we assume that the XYZ coordinate has already

been divided by white (and so are numbers between 0 and 1).

$$\begin{bmatrix} L^* \\ a^* \\ b^* \end{bmatrix} = \begin{bmatrix} 0 & 116 & 0 \\ -500 & 500 & 0 \\ 0 & 200 & -200 \end{bmatrix} \begin{bmatrix} \gamma(X) \\ \gamma(Y) \\ \gamma(Z) \end{bmatrix} - \begin{bmatrix} 16 \\ 0 \\ 0 \end{bmatrix} \quad (5)$$

In Equation (5), the gamma function  $\gamma()$ , which is approximately the cube root power, maps the X, Y or Z to brighter counterparts. Then the matrix mixes the XYZs in generating  $(L^*, a^*, b^*)$  coordinates. There is a final offset subtraction for  $L^*$ .

Let us suppose that we are only interested in  $a^*$  error and to minimize this error we have can choose a bespoke color correction matrix  $M^a$ . Because  $a^* = 500\gamma(Y) - 500\gamma(X)$ , it has no dependency on Z, so  $M^a$  has the form:

$$M^a = \begin{bmatrix} M_{11}^a & M_{12}^a & M_{13}^a \\ M_{21}^a & M_{22}^a & M_{23}^a \\ \times & \times & \times \end{bmatrix} \quad (6)$$

where the symbol  $\times$  indicates 'don't care' (since Z has no bearing on the calculation for  $a^*$ ). Analogous correction matrices,  $M^b$  and  $M^L$ , for  $b^*$  and  $L^*$  are respectively defined:

$$M^b = \begin{bmatrix} \times & \times & \times \\ M_{21}^b & M_{22}^b & M_{23}^b \\ M_{31}^b & M_{32}^b & M_{33}^b \end{bmatrix} \quad (7)$$

$$M^L = \begin{bmatrix} \times & \times & \times \\ M_{21}^L & M_{22}^L & M_{23}^L \\ \times & \times & \times \end{bmatrix} \quad (8)$$

That is,  $b^*$  depends on Y and Z and  $L^*$  only on Y.

It follows, and this is the key insight behind our extended linear color correction method, that to map RGB to XYZ to minimize the  $L^*$ ,  $a^*$  and  $b^*$  error (minimize Delta E error) we should solve three separate minimizations. Extended linear color correction method minimizes the following expression:

$$\min_{M^c} \sum_{i=1}^N ||f_c(M^c \rho_i) - f_c(\underline{x}_i)||^2, \quad c \in \{L^*, a^*, b^*\} \quad (9)$$

where  $f_c()$  is the scalar function returns the a single particular CIE LAB value,  $c \in \{L^*, a^*, b^*\}$ .

Clearly, we need to estimate X from the RGB,  $\rho$ , only once and similarly Z is also estimated uniquely. We respectively solve for  $M_{1i}^a$  and  $M_{3i}^b$ ,  $i = 1, 2, 3$ . In contrast we estimate Y three times. Unlike linear regression we cannot find  $M^a$ ,  $M^b$  and  $M^L$  in closed-form. Rather, we adopt a gradient-descent type minimizer to solve for the correction matrices. This said, we can only be sure of finding locally optimal answers and so we acknowledge that the promising

results presented in the Experiments section might be able to be improved still further.

At this point it is informative to compare our new method to finding the single  $3 \times 3$  matrix that minimizes a Delta E error. There, we solve for X, Y and Z together and at the same time given a single  $3 \times 3$  matrix transform. That is we minimized Delta E by optimising 9 unknown parameters (the components of the matrix). In our extended framework we have 15 degrees of freedom (count the variables in Eqs. (6), (7) and (8)). These extra degrees of freedom explains why we are able to achieve an improvement in fitting performance (see next section).

Finally, we note that although the above derivation is targeted toward minimizing CIE LAB a similar methodology could be followed to minimize CIE LUV or perceptual correlated in CIE CAM.

### Visualizing the result of extended linear color correction

As presented, our new method is a way of better estimating CIE LAB values from a single estimate of X and Z and three separate estimates of Y. Clearly, unlike conventional least-squares regression, we cannot in a single image represent these 5 quantities. Instead if we wish to *see* the image we must instead map the estimated CIE LAB values to corresponding XYZs and from there to sRGB. The result of inverting the CIE LAB transform for the extended regression result for mapping Figure 1a) to 1b) - raw to rendered - is shown in Figure 1f). The corresponding Delta E error map is shown in Figure 1e). Notice there is a large improvement in color correction (the error is about 2/3 less compared with least-squares regression). In Figure 2c) we show the crop of the jersey. Visually, the extended linear correction supports a closer color match to the ground-truth, 2a), compared with simple linear regression, 2b).

## Experiments

We evaluate our new extended linear color correction, denoted **ExLCC**, against 3 competing methods. These comprise simple least-squares, **LsqCC** where we find the  $3 \times 3$  matrix best mapping RGBs to XYZs. We also test against the best  $3 \times 3$  matrix optimized for CIE LAB error, **LabCC**. As commented previously, linear methods are often used in color correction because they are exposure invariant. The non-linear Root-Polynomial color correction, **RpCC**, is also designed to be exposure invariant. In **RpCC** the terms  $\sqrt{RG}$ ,  $\sqrt{RB}$  and  $\sqrt{GB}$  are added to  $R$ ,  $G$  and  $B$  and the corresponding regression solves for a  $3 \times 6$  color correction matrix. The output generated for **RpCC** mapping Figure 1a) to 1b) is shown in 1h) and the corresponding Delta E error map shown in 1g). For this example **RpCC** delivers significantly improved color correction compared

Algorithm	Mean	Median	95% quantile	Max
<b>LsqCC</b>	1.93	1.18	5.96	22.63
<b>LabCC</b>	1.85	1.36	5.00	14.31
<b>ExLCC</b>	<b>1.48</b>	1.09	<b>3.91</b>	<b>13.32</b>
<b>RpCC</b>	1.49	<b>0.96</b>	4.35	17.15

Table 1. Mean, median, 95% quantile and max CIE  $\Delta E$  color correction performance averaged over 28 cameras and 102 lights, boldface indicates the leading algorithm per column

Algorithm	Mean	Median	95% quantile	Max
<b>ExLCC</b>	<b>64%</b>	0%	<b>96%</b>	<b>64%</b>
<b>RpCC</b>	36%	<b>100%</b>	4%	7%

Table 2. The % of times a given algorithm, per statistic, delivers the minimum error, the best performance, over the set of 28 cameras, boldface indicates the leading algorithm per column

with **LsqCC**, see 1c) and 1d). Note that the average Delta E error for **RpCC** and **ExLCC** is about the same.

We now carry out a second experimental test using synthetic data and numerical integration. For surface reflectances we take the composite set of 1995 measurements from [3]. From the same database we use the 102 measured illuminant spectra. Finally we use the 28 sets of spectral sensitivities for the cameras measured at RIT[16].

For the  $i$ th light and the  $j$ th camera we numerically integrate the 1995 RGBs and also the corresponding 1995 XYZs for this light. We now wish to evaluate in terms of CIE LAB Delta E error how well each of our 4 correction methods work. We evaluate each method on a 3-fold cross validation basis. That is to say we divide the RGBs, randomly, into 3 fold sets of roughly equal sizes. Then for each set we train our regression method on the complement of each fold set (the RGBs not in the set) and test on the RGBs on the fold set itself. For each *fold* we calculate 4 summary statistics the *mean*, *median*, *95% quantile* and *max* Delta E errors. Over all 28 cameras and all 102 illuminants we then average these average statistics. The results are summarised in Table 1 where overall our new **ExLCC** is shown to deliver the best color correction performance.

Now, we calculate per statistical measure over all 28 cameras the % of times an individual color correction method delivered the best results. These percentages are tabulated in Table 2 for **ExLCC** and **RpCC**. Note, only for, the max statistic the %s do not sum to 100% (here **LabCC** delivered leading performance 29% of the time). Excepting the median statistic, the new **ExLCC** method delivers the best performance for most cameras.



Figure 3. A Macbeth color checker

Algorithm	Mean	Median	95% quantile	Max
<b>LsqCC</b>	2.75	2.41	5.73	6.88
<b>LabCC</b>	2.52	2.23	5.74	7.10
<b>ExLCC</b>	<b>2.36</b>	<b>1.75</b>	<b>5.65</b>	6.99
<b>RpCC</b>	2.65	2.29	6.09	<b>6.86</b>

Table 3. Mean, median, 95% quantile and max CIE  $\Delta E$  color correction performance for the images of 3 Macbeth color checkers, boldface indicates the leading algorithm per column

In our third experiment, we validate our experimental result on real data. We took a picture, using a Nikon D70 camera, of a standard 24 patch Macbeth color Checker in 3 positions around a historical site near our lab, see Figure 3 for one of the pictures. Because we only have 24 patches (and 6 of these are achromatic patches) the color diversity is small and this might effect the efficacy of learning the correct transform. So, we carry out a leave-one out cross validation. Here we train on 23 RGBs and test on the 24th. We do this 24 times where each patch in the checker is left out once. For each Macbeth image we as before calculate the mean, median, 95% quantile and max errors as before and then average these averages over the 3 images. The results are tabulated in Table 3.

Compared to simple linear correction we see a significant uplift in performance for the mean and median errors (Linear least-squares has respectively 17% and 38% higher error). However, the performance step is smaller for the synthetic test (and this is also true for the Root-Polynomial method).

We posit that the difference in performance is due to the fact that the test on real color checker images was carried out under daylight conditions where it is known that a linear colour correction can (and does, as evidenced in Table 3) deliver good colorimetric accuracy. In future work we will repeat this experiment for a large corpus of real reflectances and lights and we predict the performance increment will correspond with the synthetic results reported in Table 1.

## Conclusion

In this paper we demonstrated the result that the best linear color correction transform that minimizes  $a^*$  is not the same as the best transform for  $b^*$  or  $L^*$ . This led to the development of what we call extended linear color correction where we solve for three different linear corrections (one of  $L^*$ ,  $a^*$  and  $b^*$ ). Due to the structure of the CIE LAB equations, these 3 correction transforms are parameterized by  $5 \times 3$  unique linear coefficients. Compared with our new extended linear approach, simple linear least-squares, performs between 20% and 70% worse, depending on the data set and the type of error being measured (e.g. mean, 95% quantile and max). The extended linear approach was also shown to deliver leading correction performance compared with the Root-Polynomial color correction method which had been shown to improve upon simple linear correction.

## Acknowledgements

Graham Finlayson is grateful for the support of EPSRC grant J005223 and Apple Inc.

## References

- [1] C. F. Andersen and J. Y. Hardeberg. Colorimetric characterization of digital cameras preserving hue planes. In *13th Color and Imaging Conference*, pages 141–146, 2005.
- [2] M. Anderson, R. Motta, S. Chandrasekar, and M. Stokes. Proposal for a standard default color space for the internet:srgb. In *IS&T and SID's 4th Color Imaging Conference*, pages 238–245. 1995.
- [3] K. Barnard, V. C. Cardei, and B. V. Funt. A comparison of computational color constancy algorithms. i: Methodology and experiments with synthesized data. *IEEE Transactions on Image Processing*, 11(9):972–984, 2002.
- [4] B. Bastani, R. Ghaffari, and B. V. Funt. Optimal linear rgb-to-xyz mapping for color display calibration. In *The Twelfth Color Imaging Conference*, pages 223–227, 2004.
- [5] R. Bolle, J. Connell, N. Haas, R. Mohan, and G. Taubin. Veggievision: A produce recognition system. In *Workshop on the applications of computer vision*, pages 35–38, 1996.
- [6] C. Boukouvalas, J. Kittler, R. Marik, and M. Petrou. Color grading of randomly textured ceramic tiles using color histograms. *IEEE Transactions on Industrial Electronics*, 46(1):219–226, Feb 1999.
- [7] R. Brown. The estimation of natural reflectances by a cone-based linear model. In *1st Colour Imaging Conference*, page 117, 1993.
- [8] P. D. Burns and R. S. Berns. Image noise and colorimetric precision in multispectral image capture. In *Proc. 6th IS&T/SID Color Imaging Conference, IS&T*, pages 83–85, 1998.
- [9] M. Drew and B.V. Funt. Natural metamers. *CVGIP:Image Understanding*, 56:139–151, 1992.

- [10] G. Finlayson and M. S. Drew. Constrained least-squares regression in color spaces. *Journal of Electronic Imaging*, 1997.
- [11] G. D. Finlayson, M. Mackiewicz, and A. Hurlbert. Color correction using root-polynomial regression. *IEEE Transactions on Image Processing*, 24(5):1460–1470, 2015.
- [12] D. H. Foster, K. Amano, S. M. C. Nascimento, and M. J. Foster. Frequency of metamerism in natural scenes. *J. Opt. Soc. Am. A*, 23(10):2359–2372, Oct 2006.
- [13] S. K. H. Ting, L. Zheng and M. Brown. Nonuniform lattice regression for modeling the camera imaging pipeline. In *12th European Conference on Computer Vision*, pages 556–568. 2012.
- [14] G. Hong, M. R. Luo, and P. A. Rhodes. A study of digital camera colorimetric characterization based on polynomial modeling. *Color Research & Application*, 26(1):76–84, 2001.
- [15] H. Ives. The transformation of color mixture equations from one system to another. *Journal of the Franklin Institute*, pages 673–710, 1915.
- [16] J. Jiang, D. Liu, J. Gu, and S. Ssstrunk. In *Workshop on the Applications of Computer Vision*, pages 168–179. IEEE Computer Society, 2013.
- [17] M. Mackiewicz, C. F. Andersen, and G. Finlayson. Hue plane preserving colour correction using constrained least squares regression. In *23rd Color and Imaging Conference*, pages 18–23. 2015.
- [18] L. Maloney. Evaluation of linear models of surface spectral reflectance with small numbers of parameters. *J. Opt. Soc. Am. A*, 3:1673–1683, 1986.
- [19] D. Marimont and B. Wandell. Linear models of surface and illuminant spectra. *J. Opt. Soc. Am. A*, 9(11):1905–1913, 92.
- [20] J. S. McElvain and W. Gish. Camera color correction using two-dimensional transforms. In *21st Color and Imaging Conference*, pages 250–256, 2013.
- [21] G. W. Meyer, H. E. Rushmeier, M. F. Cohen, D. P. Greenberg, and K. E. Torrance. An experimental evaluation of computer graphics imagery. *ACM Transactions on Graphics*, 5:30–50, 1986.
- [22] R. Ramanath, W. Snyder, Y. Yoo, and M. Drew. Color image processing pipeline. *Signal Processing Magazine, IEEE*, 22(1):34–43, 2005.
- [23] M. Stokes, M. Fairchild, and R. Berns. Precision requirements for digital color reproduction. *ACM Transactions on Graphics*, 11(4):406–422, october 1992.
- [24] D. C. V.Cheung, S. Wetland and C. Ripamonti. A comparative study of the characterisation of colour cameras by means of neural networks and polynomial transforms. *Coloration technology*, pages 19–25, 2004.
- [25] M. Vrhel and H. Trussel. Physical device illumination correction. In *Device-Independent Color Imaging and Imaging Systems Integration*, volume 1909, pages 84–91. SPIE, 1993.
- [26] M. J. Vrhel and H. J. Trussell. The mathematics of color calibration. In *ICIP (1)*, pages 181–185, 1998.
- [27] B. Wandell. The synthesis and analysis of color images. *IEEE Trans. Patt. Anal. and Mach. Intell.*, PAMI-9:2–13, 1987.
- [28] G. Wyszecki and W. Stiles. *Color Science: Concepts and Methods, Quantitative Data and Formulas*. Wiley, New York, 2nd edition, 1982.

## Author Biography

Graham Finlayson, fellow of the IS&T, directs the Colour Lab at the University of East Anglia. Garrett Jonhson is a senior scientist at Apple Inc.