Root-Polynomial Colour Correction

Graham D. Finlayson*[1], *Michal Mackiewicz*[1], *Anya Hurlbert*[2]; [1] *School of Computing Sciences, University of East Anglia, Norwich, UK;* [2] *Newcastle University, Newcastle, UK

Abstract

Cameras record three colour responses (RGB) which are device dependent i.e. different cameras will produce different RGB responses for the same scene. Moreover, the RGB responses do not correspond to the device-independent tristimulus values as defined by the CIE. The most common method for mapping RGBs to XYZs is the simple 3×3 linear transform (usually derived through regression). While this mapping can work well it does sometimes map RGBs to XYZs with high error. On the plus side the linear transform is independent of camera exposure. An alternative and on the face of it more powerful, method for colour correction is polynomial regression. Here, the RGB at a pixel is augmented by polynomial terms e.g up to second order RGB maps to the 9-vector ($R,G,B,R^2,G^2,B^2,RG,RB,GB$). With respect to this polynomial expansion colour correction is a 9×3 linear transform. For a given calibration set-up polynomial regression can work very well indeed and can reduce colorimetric error by more than 50%. However, unlike linear maps the polynomial fit depends on exposure: as exposure changes the vector of polynomial components alters in a non linear way. In this paper we propose a new polynomial-type regression which we call 'Root-Polynomial Colour Correction'. Our idea is to take each term in a polynomial expansion and take its k^{th} root of each k-order term. For the 2^{nd} order polynomial expansion the corresponding independent root terms are R,G,B, $\sqrt{RG}$, $\sqrt{RB}$ and $\sqrt{GB}$ (6 independent terms instead of 9: the first roots of R, G and B equal the 2^{nd} roots of R^2, G^2 and B^2). It is easy to show terms defined in this way scale with exposure and so a 6×3 regression mapping can be used for colour correction. Encouragingly, our initial experiments demonstrate that root-polynomial colour correction enhances colour correction performance on real and synthetic data.

Introduction

The camera characterisation methods can be divided into the two main groups: (a) spectral sensitivity based (b) colour target based. The former encompass the methods, which utilise expensive equipment: a monochromator and a radiometer [1, 2] and therefore they are limited only to the well equipped laboratories. The latter encompass the methods which only require a reference target and corresponding known XYZs (for this target). The colour target is captured by the camera being characterised and also measured by a spectrophotometer (for color correction to a standard reference space such as sRGB the chart can be measured once ideally at the time of manufacture) resulting in the set of RGB values of all patches and their corresponding tristimulus values. The colour target based characterization is more widely used, as it only requires a known target and a spectrophotometer/chromameter. Many methods attempting to establish a mapping between the RGB and XYZ values have been reported in the literature and they include: three dimensional look-up tables [3], least-squares polynomial regression [4–6] and neural networks [6–8]. Hong et al. [4] studied camera characterization using variable length polynomial regression with least squares fitting and found that camera characterization accuracy is reliable when the number of training samples is at least 40-60. For that sample size, they obtained ΔE (CMC(1:1)) [9] of around 1 unit.

Despite the variety of correction transforms the venerable and simple 3×3 matrix transform has much to commend it. First, assuming reflectances are 3 dimensional (approximately the case), the mapping from RGB to XYZ has to be a 3x3 matrix. Marimont and Wandell [10] extended this idea to model 'pseudo reflectances' which comprise that part of a reflectance that can be measured by a camera (under different lights). To a first order typical lights and surfaces interact with typical cameras as if reflectances were 3 dimensional. Another plus for linear transforms is that they are, theoretically, independent of changes in the overall illumination irradiance. That is, the linear transform should correctly correct colours independently of changes in the overall light levels on a fixed scene. At a fixed camera exposure, as the overall light level is scaled up or down, the RBGs recorded by the camera in turn scale by a factor k (assuming perfectly linear sensors). Likewise, the tristimulus values of the corresponding surfaces in the scene will scale by the same factor k (assuming the light level scaling is constant across wavelength). The linear transform therefore remains the same, since the factor k cancels out on both sides of the equation. Conversely, for a fixed scene under fixed illumination, changes in camera exposure will induce linear scaling of the RGBs that map to the fixed tristimulus values, necessitating a corresponding scaling of the linear transform. This said, for typical digital camera usage we are much more interested in the former cases e.g. mapping RGBs to corresponding non-fixed; XYZs (or display RGBs). The issue of the 'scale' of the display is typically dealt with by normalising to (usually) the bright part of the image (where this normalisation can be carried out when the image is acquired by an 'exposure control' algorithm. Further mapping of RGBs to display range is called tone mapping.

Nonetheless, significant errors may still result from the linear model for some surfaces. Indeed, given a linear fit from RGBs to XYZs, errors for individual surfaces can be in excess of $10\Delta E$. While this 'model error' is high the linear map has the advantage that it works correctly as exposure changes. Consider the same surface viewed under different light levels in different parts of the scene. The correct linear map taking RGBs to XYZs (or display RGBs) is the same in both cases. This turns out to be a very important practical property which we, in this paper, treat with some care. To reduce 'model error' the simplest extension to the linear approach is to use polynomial regression. A 2^{nd} order polynomial expansion where each image RGB is mapped to the 9-vector $[R\ G\ B\ R^2\ G^2\ B^2\ RG\ RB\ GB]$. Unfortunately, if the RGB is scaled by k, the individual components of the 9-vector either scale by k or k^2. Thus, if we scale our data (physically this is the effect of changing the exposure) then the best 9×3 colour correction matrix must also change. This is a significant problem in real images. For example if we place a reference target in a bright well exposed area of a scene and then place it in shade the RGB values (no shade and shade) can be a scaling factor of 10s to 1 different. Naively used if we have a polynomial transform for our colour camera, the in-shade target will be mapped to different colours than the target not in the shade.

Figure 1 illustrates the problem. We trained the polynomial and root-polynomial models using SFU 1995 reflectance dataset [11]. Next, we selected small number of reflectances from this dataset and calculated their true xy chromaticity coordinates (marked with x). We also calculated their chromaticity coordinates according to the root-polynomial model (marked with +). And finally, we calculated their chromaticities for the polynomial model for a number of intensity scales ranging from 0.1 to 1. In Figure 1, one can see the chromatic shifts induced by the polynomial model as one scales the intensity of the light.

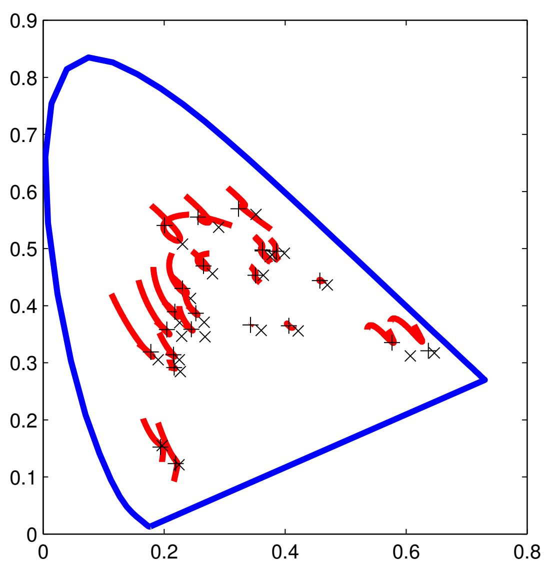

Figure 1. Selection of reflectances from the SFU 1995 dataset and their true xy coordinates (x); xy coordinates according to the root-polynomial model of the 4^{th} order (+); and chromatic shifts produced by the polynomial model of the 4^{th} order (red).

Background

Let ρ define a three element vector representing the three camera responses and q their corresponding tristimulus values. A simple 3×3 colour correction transform is written as:

$$q = M\rho \tag{1}$$

The matrix M is generally found by some sort of least-squares regression. Let us denote a set of N known XYZs for a reflectance target as Q and the corresponding camera responses as the $N\times 3$ matrix R. We find the least-squares mapping from R to Q using the Moore-Penrose inverse:

$$M = (R^T R)^{-1} R^T Q \tag{2}$$

In polynomial regression, vector ρ is extended by adding additional polynomials of increasing degree. In this study, we will be investigating the following polynomials:

- (p,1): $\rho = [r,g,b]^T$
- (p,2): $\rho = \left[r,g,b,r^2,g^2,b^2,rg,gb,rb\right]^T$
- (p,3):
$$\rho = \Big[r,g,b,r^2,g^2,b^2,rg,gb,rb, r^3,g^3,b^3,rg^2,gb^2,rb^2,gr^2,bg^2,br^2,rgb\Big]^T$$
- (p,4);
$$\rho = \Big[r,g,b,r^2,g^2,b^2,rg,gb,rb, r^3,g^3,b^3,rg^2,gb^2,rb^2,gr^2,bg^2,br^2,rgb, r^4,g^4,b^4,r^3g,r^3b,g^3r,g^3b,b^3r,b^3g, r^2g^2,g^2b^2,r^2b^2,r^2gb,g^2rb,b^2rg\Big]^T$$

Using a polynomial expansion the 3 numbers recorded in a pixel are mapped to respectively 9, 19 and 34 numbers respectively. Colour correction is now carried out by 3×9, 3×19 and 3×34 matrices. If we think of R as, in general, denoting the polynomial response of N surfaces then Equation 2 can be used to solve for the polynomial colour correction matrix. Of course high order data expansions can result in unstable (rank deficient data sets). This problem can be ameliorated by regularising the regression.

Root Polynomials

For fixed exposure, polynomial regression really can deliver significant improvements to colour correction. However, in reality the same reflectance will induce many different brightnesses (even for the same fixed exposure and viewing conditions). The act of placing a surface in different parts of the scene will markedly change its brightness. Thus, if we are to use a polynomial type regression we must progress with care.

The starting point of this paper was to ask the following question. Is there a way we can use the undoubted power of polynomial data fitting in a way that does not depend on exposure/scene irradiance? Our key observation is that the terms in any polynomial fit each have an order e.g. R, RG and R^2B are respectively order 1, 2 and 3. Multiplying each of R, G and B by a scalar k results in the terms kR, k^2RG and k^3R^2B. That is the order of the term is reflected in the power to which the exposure scaling is raised. Clearly, and this is our key insight, taking the order-root will result in terms which have the same k scalar: $(kR)^{1/1} = kR$, $(k^2RG)^{1/2} = k(RG)^{1/2}$, $(k^3R^2B)^{1/3} = k(R^2B)^{1/3}$. For a given p^{th} order polynomial expansion, we take each term and raise it to the inverse of its order. The unique individual terms that are left are what we use in Root Polynomial Colour Correction.

The root polynomials corresponding to the 2^{nd}, 3^{rd} and 4^{th} order polynomial expansions are written below:

- (r,2); $\rho = [r,g,b,\sqrt{rg},\sqrt{gb},\sqrt{rb}]^T$

- (r,3);

$$\rho = \Big[r, g, b, \sqrt{rg}, \sqrt{gb}, \sqrt{rb}, \sqrt[3]{rg^2}, \sqrt[3]{gb^2}, \sqrt[3]{rb^2}, \sqrt[3]{gr^2}, \sqrt[3]{bg^2}, \sqrt[3]{br^2}, \sqrt[3]{rgb}\Big]^T$$

- (r,4);

$$\rho = \Big[r, g, b, \sqrt{rg}, \sqrt{gb}, \sqrt{rb}, \sqrt[3]{rg^2}, \sqrt[3]{gb^2}, \sqrt[3]{rb^2}, \sqrt[3]{gr^2}, \sqrt[3]{bg^2}, \sqrt[3]{br^2}, \sqrt[3]{rgb}, \sqrt[4]{r^3g}, \sqrt[4]{r^3b}, \sqrt[4]{g^3r}, \sqrt[4]{g^3b}, \sqrt[4]{b^3r}, \sqrt[4]{b^3g}, \sqrt[4]{r^2g^2}, \sqrt[4]{g^2b^2}, \sqrt[4]{r^2b^2}, \sqrt[4]{r^2gb}, \sqrt[4]{g^2rb}, \sqrt[4]{b^2rg}\Big]^T$$

Notice that the number of terms is reduced. This is because the root operation is many to 1. For example R, R^2, R^3 and their respective 1^{st}, 2^{nd} and 3^{rd} roots are all equal to R. Clearly R can only occur once in the root polynomial regression.

What we expect from the root-polynomial model is improvement over the linear model for the regions where linearity is poor (due to the types of surfaces, sensors or lights), but also, crucially, elimination of the effects of non-linear magnification of linear changes in the overall light level. Thus, we expect the root-polynomial model to perform better than the linear model for fixed scenes, illuminations and camera exposures, as well as for a fixed scene and camera exposure, under changing illumination. (Note that these predictions all assume that the camera responses remain in the unsaturated range.)

Experiments

To compare the performance of polynomial and root-polynomial colour correction, we performed both synthetic data simulations and a real data experiment. As to the former, we used the Sony DXC-930 camera sensor sensitivities [11] to integrate the spectral data from three surface reflectance datasets. First dataset comprised 96 reflectances of the X-rite SG colour checker (border patches excluded), the second dataset contained 180 patches of the Macbeth DC colour checker (again border patches were excluded) and the last one 1995 surfaces collated at the Simon Fraser University [11]. For each dataset we performed a simulation, in which we integrated the Sony sensor sensitivities and the colour matching functions under D65 illuminant producing corresponding sets of camera responses (RGBs) and XYZs. Spectra calculations were carried out for 31 spectral channels - 400-700nm sampled every 10nm. Next, we built the regression models using previously defined polynomial and root-polynomial terms. For each of the four datasets, we performed the leave-one-out validation i.e. we built the model using all but one of the surfaces from the dataset and tested that model on the remaining patch; we repeated this for all the patches in the dataset and calculated mean ΔE in the CIELUV colour space. The results of these simulations can be seen in Table 1.

Next, for the DC and SFU datasets, we simulated an increase and decrease in the scene irradiance by multiplying the earlier calculated camera sensor responses and the ground truth XYZ values by factors 0.5 and 1.5. We used these corresponding sets of RGBs and XYZs to test the earlier learned original scene irradiance polynomial and root-polynomial models. The results of these simulations can be seen in Table 2.

Moreover, we performed a real camera characterisation. The experimental set-up was as follows. The X-rite SG colour chart was positioned in a dark room, illuminated with a tungsten bulb light and imaged with Nikon D60 camera. First, we captured an image of the colour chart and measured XYZ values of each of the 96 patches using Photo Research PR-650 spectrophotometer. Next, the level of the light was decreased by roughly the factor of two using the dimmer and the image capture and patch measurements were repeated (for the same camera settings). The linear 16-bit images were extracted from the camera raw images using DCRAW program. The dark level was captured with the lens cap on and subtracted from the camera responses. We used the data obtained for the first lighting condition to derive a set of polynomial and root-polynomial models as described in the earlier sections. The models were evaluated using the leave-one-out method in the same manner as in the earlier simulations. The results of the validation can be seen in the second column (fixed illumination) in Table 3. The third column in the same table contains the results of testing these models after the intensity of the light was decreased. In addition, we simulated an increase in the the light intensity by multiplying the original (first lighting condition) camera sensor responses and the measured XYZ values by factors 1.1 and 1.2. Again, we used the resulting RGBs and XYZs to test the earlier learned polynomial and root-polynomial models. The results of these simulations can be seen in the last two columns of Table 3.

Discussion

If we compare the results of the polynomial and root-polynomial models in a fixed illumination scenario, we can see that the root-polynomials usually outperform the polynomial models including those of higher order. However, the difference between the best high order polynomial models and the root-polynomial models is not large. In Table 1, we can see that for the DC dataset the root-polynomial model performs worse than the polynomial, but for the smaller SG dataset and the largest and most relevant SFU dataset, the situation is reversed.

The results from Tables 2-3 show that the polynomial models deteriorate under change of scene irradiance/exposure condition, whereas the root-polynomial models remain invariant. An important observation is the fact that the root-polynomial results are always better than the results obtained for the linear model (p,1) - the only polynomial model, which is invariant to the change of exposure. In the experiments we report here, we simulated only a slight increase in the scene irradiance/exposure as for the larger increases the polynomial models fail completely, which results e.g. in mapping camera responses to negative XYZs and consequently meaningless LUV errors. However, the message from Tables 2 and 3 as well as Fig. 1 is clear. If you carry out naive polynomial regression to fit data at different exposures, high error can result. Conversely, the root-polynomial colour correction works well independent of exposure.

With regard to the real camera experiment, similar trends can be observed i.e. the root-polynomial correction performs better than the polynomial correction in particular significantly better than the linear and is invariant with respect to the change of illumination intensity. Moreover, polynomial regression fails under the change of illumination condition. The smaller errors for that camera (than those for the simulated Sony sensors and the SG chart) suggest that Nikon D60 sensors are more colorimetric than

Table 1. Synthetic data simulation results for fixed illumination condition. The errors obtained for four datasets are given as the mean, median and 95 percentile error in the CIELUV colour space.

dataset	SG			DC			SFU		
model type	mean	med	95 pt.	mean	med	95 pt.	mean	med	95 pt.
p,1	4.8	2.9	20	3.8	1.9	14	2.6	1.4	7.7
p,2	4.0	2.7	11	3.1	1.8	10	2.4	1.4	7.2
p,3	3.2	2.2	7.5	2.4	1.3	7.7	1.9	1.2	6.5
p,4	3.1	2.2	8.2	**2.0**	1.2	6.9	**1.8**	1.1	6.0
r,2	2.8	1.8	8.7	2.4	1.3	8.5	2.1	1.2	7.0
r,3	**2.6**	1.6	8.1	2.2	1.3	7.5	**1.8**	1.2	6.1
r,4	**2.6**	1.6	8.3	2.2	1.3	7.7	**1.8**	1.1	6.1

Table 2. Synthetic data simulation results for the DC and SFU dataset as the light level was decreased and increased. The errors obtained for four datasets are given as the mean, median and 95 percentile error in the CIELUV colour space.

dataset	DC decr. 50%			DC incr. 50%			SFU decr. 50%			SFU incr. 50%		
model type	mean	med	95 pt.	mean	med	95 pt.	mean	med	95 pt.	mean	med	95 pt.
p,1	3.8	1.9	14	3.8	1.9	14	2.6	1.4	7.7	2.6	1.4	7.7
p,2	3.9	2.5	12	3.5	2.2	11	2.6	1.6	8.0	2.4	1.5	7.3
p,3	3.3	1.9	10	3.5	1.5	15	2.5	1.5	7.8	2.7	1.4	9.6
p,4	3.0	1.6	11	5.0	1.8	23	2.4	1.4	8.0	2.5	1.3	8.9
r,2	2.4	1.3	8.5	2.4	1.3	8.5	2.1	1.2	7.0	2.1	1.2	7.0
r,3	**2.2**	1.3	7.5	**2.2**	1.3	7.5	**1.8**	1.2	6.1	**1.8**	1.2	6.1
r,4	**2.2**	1.3	7.7	**2.2**	1.3	7.7	**1.8**	1.1	6.1	**1.8**	1.1	6.1

Table 3. Nikon D60 characterisation results. The errors are given as the mean, median, 95 percentile and maximum error in the CIELUV colour space.

	fixed illumination				illumination decreased by 50%				illumination increased by 10%				illumination increased by 20%			
model	mean	med.	95 pt.	max	mean	med.	95 pt.	max	mean	med.	95 pt.	max	mean	med.	95 pt.	max
p,1	2.8	1.9	8.4	13.4	2.7	1.8	8.6	13.3	2.8	1.9	8.4	13.4	2.8	1.9	8.4	13.4
p,2	2.3	1.9	7.0	8.4	2.6	2.3	7.0	10.4	2.3	1.9	6.5	8.3	2.4	1.8	7.0	8.1
p,3	2.2	1.8	6.0	7.9	2.6	2.3	6.0	7.4	2.4	1.9	6.0	7.6	2.6	1.8	6.8	11.8
p,4	2.4	2.0	6.3	8.6	3.1	2.5	7.2	10.4	2.8	2.2	6.7	14	3.5	2.2	10.5	21.4
r,2	**2.1**	1.5	6.6	8.3	**2.2**	1.6	7.4	8.8	**2.1**	1.5	6.6	8.3	**2.1**	1.5	6.6	8.3
r,3	**2.1**	1.5	6.4	11.2	**2.2**	1.5	6.1	8.7	**2.1**	1.5	6.4	11.2	**2.1**	1.5	6.4	11.2
r,4	2.2	1.5	6.2	12.9	2.2	1.7	6.2	9.1	2.2	1.5	6.2	12.9	2.2	1.5	6.2	12.9

the Sony DXC-930 sensors.

Another interesting observation about the root-polynomials is that, the results obtained for different orders are relatively similar. Usually the largest improvement over the linear model takes place in the second order root-polynomial by adding just three extra terms into the model. The results of the third and the fourth order root-polynomials are very similar and only slightly better than those of the second order. In Tables 1 and 3, we can see that the 2^{nd} order root-polynomial (6 terms) outperforms even 4^{th} order polynomial (34 terms) for the SG dataset and for the Nikon camera; and is comparable to the 3^{rd} order polynomial (19 terms) for the DC and SFU datasets. Thus, the smaller number of terms of root-polynomials is their yet another advantage over the polynomials.

Conclusions

'Root-Polynomial Colour Correction' builds on the earlier widely used polynomial models, but unlike its predecessors is invariant to the change of camera exposure and/or scene irradiance. The results presented in this paper show that this algorithm always outperforms linear regression and offers a significant improvement over polynomial models when the exposure/scene irradiance varies.

References

[1] F. Martínez-Verdú, J. Pujol, and P. Capilla, "Calculation of the color matching functions of digital cameras from their complete spectral sensitivities." *J. Imaging Sci. Technol.*, vol. 46, pp. 15–25, 2002.

[2] M. Brady and G. E. Legge, "Camera calibration for natural image studies and vision research," *J. Opt. Soc. Am. A*,

vol. 26, no. 1, pp. 30–42, 2009.

[3] P. Hung, “Colorimetric calibration in electronic imaging devices using a look-up tables model and interpolations,” *Journal of Electronic Imaging*, vol. 2, no. 1, pp. 53–61, 1993.

[4] G. Hong, M. R. Luo, and P. A. Rhodes, “A study of digital camera characterisation based on polynomial modelling,” *Color research and application*, vol. 26, no. 1, pp. 76–84, 2001.

[5] H. Kang, “Colour scanner calibration,” *Journal of Imaging Science and Technology*, vol. 36, pp. 162–70, 1992.

[6] T. Cheung and S. Westland, “Colour camera characterisation using artificial neural networks,” in *IS&T/SID Tenth Color Imaging Conference*, vol. 4. The Society for Imaging Science and Technology, The Society for Information Display, 2002, pp. 117–20.

[7] H. Kang and P. Anderson, “Neural network application to the color scanner and printer calibration.” *Journal of Electronic Imaging*, vol. 1, pp. 125–34, 1992.

[8] L. Xinwu, “A new color correction model based on bp neural network,” *Advances in Information Sciences and Service Sciences*, vol. 3, no. 5, pp. 72–8, June 2011.

[9] F. Clarke, R. McDonald, and B. Rigg, “Modification to the jpc 79 colour-difference formula,” *J Soc Dyers Colour*, vol. 100, pp. 128–32, 1984.

[10] D. H. Marimont and B. A. Wandell, “Linear models of surface and illuminant spectra,” *J. Opt. Soc. Am. A*, vol. 9, pp. 1905–13, 1992.

[11] K. Barnard, L. Martin, B. Funt, and A. Coath., “A dataset for color research.” *Color Research and Application*, vol. 27, no. 3, pp. 147–151, 2002.