Hue Plane Preserving Colour Correction using Constrained Least Squares Regression.

Michal Mackiewicz[1], *Casper F. Andersen*[2], *Graham D. Finlayson*[1]; [1]*School of Computing Sciences, University of East Anglia, Norwich, UK;* [2]*The Norwegian Color Research Laboratory, Gjøvik University College, Gjøvik, Norway.*

Abstract

Andersen and Hardeberg proposed the Hue Plane Preserving Colour Correction (HPPCC) [1], which maps RGBs to XYZs using a set of linear transforms, where each transform is learned and applied in a subregion of colour space, defined by two adjacent hue planes. A hue plane is a geometrical half-plane defined by the neutral axis and a chromatic colour. A problem with the original HPCC method is that the selection of chromatic colors was a user defined choice (and the user might choose poorly) and the method as formed was not open to optimization. In this paper we present a flexible method of hue plane preserving colour correction which we call Hue Plane Preserving Colour Correction using Constrained Least Squares (HPPCC-CLSQ). This colorimetric characterization method is also based on a series of 3 by 3 matrices, each responsible for the transformation of a subregion, defined by two adjacent hue planes, of camera RGB values to the corresponding subregion of estimated colorimetric XYZ values. The matrices are constrained to white point preservation. In this new formulation, the subregions can flexibly be chosen in number and position in order to regularize and optimize the results, whilst constraining continuity crossing the hue planes. The method is compared to a choice of other state-of-the-art characterization methods and the results show that our method consistently gives high colorimetric accuracy for both synthetic and real camera data.

Introduction

In digital photography, camera characterization is a crucial element, since it relates the camera output values, device dependent *RGB* to standard observer colorimetric values, device independent *XYZ*. This relation is not straightforward for several reasons. Two of the most important reasons are known as metamerism, related to the spectral characteristics of the camera sensor and the spectral composition of the incoming light. Digital cameras today do not incorporate colorimetric colour filters on the sensor due to noise considerations and limitations in manufacturing. This means that a linear relation, also known as the Luther conditions [13, 17], between the CIE color matching functions [22] and the sensor filter spectral response functions does not exist. In practice, this means [21] that certain surfaces that look the same to the eye could be recorded in the camera as different colours and vice-versa.

Often the light source under which the camera images are taken (i.e. flash light, daylight etc.) is unknown and thus could be differing spectrally from the illuminant to which the colorimetric values are referred. This can in photographic practice lead to a significant inaccuracy in the estimation. With the assumption that the light source and the illuminant to which the sensor and observer responses are related are known and indeed are the same, this difficulty can be ruled out and thus not considered here.

Since no unique solution generally exist for the relation between camera *RGB* values and standard observer *XYZ* values, a need for optimized mapping rather than exact solution appear. In order to establish the optimal mapping numerous methods have been brought into existence. Methods include: best fitting methods such as least squares linear and polynomial regression [3, 4, 6, 9, 14, 16, 20], neural networks [15, 23] and look-up tables [11]

The best fitting optimization can be based on linear regression of colorimetric differences which is most common or non-linear least squares regression of perceptual differences. Whether colorimetric or perceptual accuracy is more important is case dependent.

Furthermore, the best fitting methods can be constrained by a choice of features of perceptual and /or colorimetric importance. White point preservation by which the normalized adopted white point in camera response is mapped directly to the adapted white point in colorimetric response is one such feature [8]. Perceptually, it is typically very import that scene neutral and white is mapped to estimated colorimetric neutral and white precisely in order to avoid colour casts in the rendered image. This is explored in the white point preserving 3 by 3 matrix [6].

Exposure and shading independence by which scalings of reflectances recorded by the camera device are mapped to the exact same scalings of the reflectances matched by the observer, is an inevitable priority when colorimetric correctness is desired. If the device response to the reflectance of a surface is recorded with values C_{rgb} and correspondingly C_{xyz} by the observer, then these responses would yield kC_{rgb} and kC_{xyz} if exposure time or irradiance on the surface was modified by a scalar value of k. In practice this scaling takes place when surfaces locally occur in a scene with different irradiances or globally as a function of exposure time. The 3 by 3 matrix method is inherently exposure and shading independent and will therefore map these scalings correctly. White point preservation in combination with exposure independence is explored in the HPPCC method [1] and subsequently in a modified polynomial approach known as root-polynomial colour correction [5].

Extending white point preservation and exposure independence to a more advanced constraint defined as hue plane preservation was explored in the HPPCC method. It is important to note that hue plane geometry as defined here is a colorimetric correlate of perceptual hue. The hue plane relates directly with dominant wavelength in a chromaticity diagram. Perceptual hue is not necessarily constant on a hue plane, but the colorimetric relations of the colours on it are. Hue plane preservation en-

ables an exact mapping of linear relations of colours that originate from a linear combination of a neutral reflection or a specular highlight and a diffuse reflection. If, for example, an object or a surface has a reflection which entirely (or at least approximately) can be characterized by a mixture of a neutral reflection with the colour of the illuminant and a reflection of a chromatic surface (i.e. diffuse reflection) [19], then the device response can be written $C_{rgb}(k_w,k_d) = k_w W_c + k_d D_c$. The corresponding response from the observer is $C_{xyz}(k_w,k_d) = k_w W_o + k_c D_o$, where W_c and W_o are the responses to the neutral, D_c and D_o are the responses to the chromatic surface and k_w and k_d are the scalars that determine the mixture. By varying k_w an k_d hue planes are described in device and observer space. In order for a mapping to be hue plane preserving the mapping $f(C_{rgb}(k_w,k_d)) = C_{xyz}(k_w,k_d)$ for all k_w and k_d, where $f()$ denotes the mapping transformation function. It can be easily seen that the 3 by 3 matrix mapping is hue plane preserving. But only two different hue planes can be mapped with a guarantee of zero error, limiting the matrix mapping accuracy. Furthermore, it can also readily be seen that the polynomial and root-polynomial mappings, with higher statistical accuracy, will not be hue plane preserving. They will distort the hue planes, which can lead to undesirable perceptual deviations of colour i.e hue shifts. In a photographic scene situation in which multiple objects of different diffuse reflections each describing a hue plane, are present, more degrees of freedom and higher accuracy is needed. This is where the hue plane preserving color correction finds its purpose. The idea is to map the device dependent hue planes to estimated device independent hue planes with a higher colorimetric accuracy than the 3 by 3 matrix method will allow, by introducing a larger number of degrees of freedom in the mapping. To obtain the higher colorimetric accuracy for all colours more white point preserving matrices are needed. They are arranged so that they can be applied more locally to a subset of colours that are situated in subregions of colour space, each delimited by two hue planes of choice.

Method

In the next two sections we will briefly describe the original Hue Plane Preserving Colour Correction method and then present in detail our new Hue Plane Preserving Colour Correction using Constrained Least Squares.

Throughout this paper we will always assume that we deal with the *RGB* colour responses that were white-balanced that is each *RGB* response has been divided by the respective colour response of the perfect diffuser denoted as $\overline{w} = (R_N, G_N, B_N)$.

Hue Plane Preserving Colour Correction

Andersen and Hardeberg [1] proposed to sort the *RGB* colour responses according to their hue correlates θ_i in the rg-chromaticity plane in the ascending order as:

$$\theta_i = \begin{cases} 0 & \text{if } r_i = g_i = \frac{1}{3} \\ \frac{\pi}{2} & \text{if } r_i = \frac{1}{3} \wedge g_i > \frac{1}{3} \\ \frac{3\pi}{2} & \text{if } r_i = \frac{1}{3} \wedge g_i < \frac{1}{3} \\ \arctan\frac{g_i - \frac{1}{3}}{r_i - \frac{1}{3}} + m\pi & \text{otherwise} \end{cases}.$$

where

$$m = \begin{cases} 0 & \text{if } r_i \geq \frac{1}{3} \wedge g_i \geq \frac{1}{3} \\ 1 & \text{if } g_i < \frac{1}{3} \\ 2 & \text{otherwise} \end{cases}$$

and the chromaticity values are calculated by:

$$r_i = R_i/(R_i+G_i+B_i);\ g_i = G_i/(R_i+G_i+B_i).$$

Next, they choose K chromatic *RGB* samples denoted as $\overline{q}_k$ that will define the boundaries between K hue regions. Finally, they find the 3×3 colour correction transform for each region k by inverting the following equation:

$$\left[\overline{q}_k, \overline{q}_{k+1}, \overline{w}\right] \mathbf{T}_k = \left[\overline{p}_k, \overline{p}_{k+1}, \overline{p}_w\right]. \quad (1)$$

where $\overline{p}_k$ denotes a measured XYZ corresponding to the colour response $\overline{q}_k$ and $\overline{p}_w$ a measured XYZ corresponding to the colour response of the neutral patch $\overline{w}$.

The first matrix in Eq. 1 is 3 by 3 and invertible and hence all colour samples at the boundaries of the hue regions as well as the neutral sample are mapped exactly (without error) to their corresponding XYZ values.

Hue Plane Preserving Colour Correction using Constrained Least Squares

First, let us limit the colour correction problem and assume that we are only interested in one of the colorimetric output channels e.g. the first channel of XYZ that is X.

Let us denote $N \times 3$ matrix of camera responses as $\mathbf{Q}$ and an N-vector of corresponding measured X responses as $\overline{x}$. Then, solving for linear colour correction is equivalent to solving the following optimisation:

$$\underset{\overline{t}}{\text{minimize}} ||\mathbf{Q}\overline{t} - \overline{x}|| \quad (2)$$

where $\overline{t}$ is a 3-vector to be found.

In our new method, we alter the linear colour correction optimisation by sorting N colour responses according to θ and dividing them into K subsets, each subset containing N_k samples such that $\sum_{k=1}^{K} N_k = N$. We can place each subset samples into the rows of K $N_k \times 3$ matrices denoted as $\mathbf{Q}_k$. Analogously, $\overline{x}$ is sorted and partitioned into K subsets denoted as $\overline{x}_k$.

Then, we seek a separate solution to colour correction problem (different correction transforms $\overline{t}_k$) for each subset k of the dataset. This amounts to solving the following optimisation problem:

$$\underset{\overline{t}_1 \ldots \overline{t}_K}{\text{minimize}} \sum_{k=1}^{K} ||\mathbf{Q}_k \overline{t}_k - \overline{x}_k||. \quad (3)$$

This can be also written as:

$$\underset{\overline{T}}{\text{minimize}} ||\mathbf{A}\overline{T} - \overline{X}||. \quad (4)$$

where $\mathbf{A}$ is a non-square $(N \times 3K)$ block diagonal matrix:

$$\mathbf{A} = \begin{bmatrix} \mathbf{Q}_1 & 0 & \cdots & 0 \\ 0 & \mathbf{Q}_2 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & \mathbf{Q}_K \end{bmatrix}, \tag{5}$$

$\overline{T}$ is a $3K$-vector: $\overline{T} = (\overline{t}_1^T, ..., \overline{t}_K^T,)^T$ and $\overline{X}$is an N-vector: $\overline{X} = (\overline{x}_1^T, ..., \overline{x}_K^T,)^T$.

Solving this optimisation will indeed produce different colour correction transforms $\overline{t}_k$ for each set of samples representative of a given range of θ. However, an *RGB* at the boundary of any two adjacent hue regions will be generally mapped to two different outputs by the two respective colour transforms. Furthermore, a neutral sample being at the boundary (and centre) of all hue regions will be also mapped to different outputs by the different colour transforms. Therefore, we require additional constraints to eliminate the two above possibilities i.e. we need to enforce C0 continuity at the boundaries of the hue regions.

Before we begin we must decide where precisely the boundaries of the hue regions are. Here, we assume that the last row of matrix $\mathbf{Q}_k$ (denoted as $\overline{q}_k^T$) contains a colour response that is at the boundary of regions k and $k+1$ with wraparound (the last row of $\mathbf{Q}_K$ contains a sample at the boundary of the $k = K$ and $k = 1$ regions). Then we formulate the following equality constraints:

$$\begin{array}{ll} \overline{q}_k^T \overline{t}_k = \overline{q}_k^T \overline{t}_{k+1} & k = 1, ..., K-1 \\ \overline{q}_k^T \overline{t}_k = \overline{q}_k^T \overline{t}_1 & k = K \\ \overline{w}^T \overline{t}_1 = \overline{w}^T \overline{t}_2 = ... = \overline{w}^T \overline{t}_K & \end{array} \tag{6}$$

where $\overline{w}$ is a white balanced camera response to the neutral patch.

These constraints can be written in a matrix form as:

$$\mathbf{C}\overline{T} = \overline{0}, \tag{7}$$

where $\mathbf{C}$ is a $2K-1 \times 3K$ matrix:

$$\mathbf{C} = \begin{bmatrix} \overline{q}_1^T & -\overline{q}_1^T & 0 & \cdots & 0 \\ 0 & \overline{q}_2^T & -\overline{q}_2^T & \cdots & 0 \\ \vdots & \vdots & \ddots & \ddots & \vdots \\ 0 & \cdots & 0 & \overline{q}_{K-1}^T & -\overline{q}_{K-1}^T \\ -\overline{q}_K^T & 0 & \cdots & 0 & \overline{q}_K^T \\ \overline{w}^T & -\overline{w}^T & 0 & \cdots & 0 \\ 0 & \overline{w}^T & -\overline{w}^T & \cdots & 0 \\ \vdots & \vdots & \ddots & \ddots & \vdots \\ 0 & \cdots & 0 & \overline{w}^T & -\overline{w}^T \end{bmatrix} \tag{8}$$

Hence, the colour correction problem to be solved becomes:

$$\begin{array}{ll} \underset{\overline{T}}{\text{minimize}} & ||\mathbf{A}\overline{T} - \overline{X}|| \\ \text{subject to} & \mathbf{C}\overline{T} = \overline{0}. \end{array} \tag{9}$$

This problem can be solved using the method of Lagrange multipliers, which provides the following closed form solution:

$$\begin{bmatrix} \overline{T} \\ \overline{z} \end{bmatrix} = \begin{bmatrix} 2\mathbf{A}^T\mathbf{A} & \mathbf{C}^T \\ \mathbf{C} & \mathbf{0} \end{bmatrix}^{-1} \begin{bmatrix} 2\mathbf{A}^T\overline{X} \\ \overline{0} \end{bmatrix} \tag{10}$$

where $\overline{z}$ is a vector of Lagrange multipliers.

We can write the above equation for all three XYZ channels as follows:

$$\begin{bmatrix} \mathbf{T} \\ \mathbf{Z} \end{bmatrix} = \begin{bmatrix} 2\mathbf{A}^T\mathbf{A} & \mathbf{C}^T \\ \mathbf{C} & \mathbf{0} \end{bmatrix}^{-1} \begin{bmatrix} 2\mathbf{A}^T\mathbf{X} \\ \mathbf{0} \end{bmatrix} \tag{11}$$

where $\mathbf{X}$ is an $N \times 3$ matrix of measured and sorted XYZs and $\mathbf{T}$ is a $3K \times 3$ matrix consisting of K 3×3 linear colour correction matrices.

In Fig. 1, we can see a visualisation of the colour correction transformation for one of the colorimetric channels created by the method described in this paper. It can be seen that the surface of the transform is continuous and consists of a number of planes anchored at the white point.

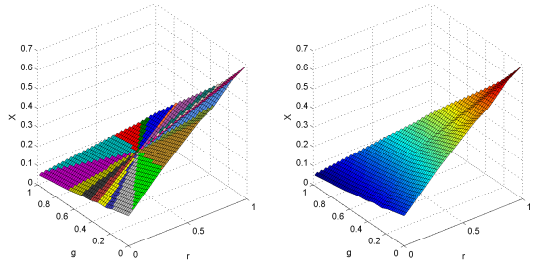

Figure 1. *A cross section along the rg chromaticity plane of the $X = f(RGB)$ colour correction hypersurface. Left: Different colours representing twenty hue slices used in this example. Right: The same figure with colouring proportional to X, making the arrangement of distinct planes at different hue slices clearly visible.*

Experiments and Discussion

To evaluate the performance of our new method, we performed both the synthetic data simulations and the real camera experiment. As to the colour correction methods we tested, we used the same suite that we employed in our earlier work [5]. These include the linear colour correction as well as the original HPPCC. For the latter, the hue circle was divided into twelve slices and the sample selection was performed based on a relative susceptibility to noise (for details see [1]).

Further, we compared the performance of the above to the polynomial and root-polynomial colour correction methods up to degree of three. An additional colour correction method that was used for comparison was the tri-linear LUT interpolation with the Graph Hessian Regularizer [7]. The size of the LUT we used was $13 \times 13 \times 13$.

Synthetic data experiments

For the synthetic experiments, we used the Sony DXC-930 camera [2] and Foveon sensor sensitivities [10]. We employed three reflectance datasets in our experiments: the X-rite SG colour

chart comprising 96 reflectances, the Macbeth DC chart comprising 180 patches, and the Simon Fraser University reflectance dataset comprising 1995 reflectances [2].

The reflectance, illuminant and sensor sensitivity spectra were sampled every 10nm from 400nm to 700nm. The above spectra were used to calculate the corresponding sets of camera responses (*RGB*s) and *XYZ*s. For each of the three datasets, we built the above colour correction models, which were then evaluated using the leave-one-out cross-validation i.e. for the dataset containing n samples, we built a model from all but one sample, which is later used for testing. This is repeated n times and the mean ΔE in the CIELUV colour space [12] was calculated.

The final matter to be decided is the number of the hue slices to be used in the HPPCC-CLSQ method. In the original HPPCC method, the authors used twelve hue slices. Here, we performed an initial experiment to select this parameter for the HPPCC-CLSQ method. We varied the number of the hue partitions from 1 to 10 and calculated the mean, median and 95 percentile ΔE errors for the three datasets and the two sensor sets. The results of this experiment can be seen in Figs 2 to 4.

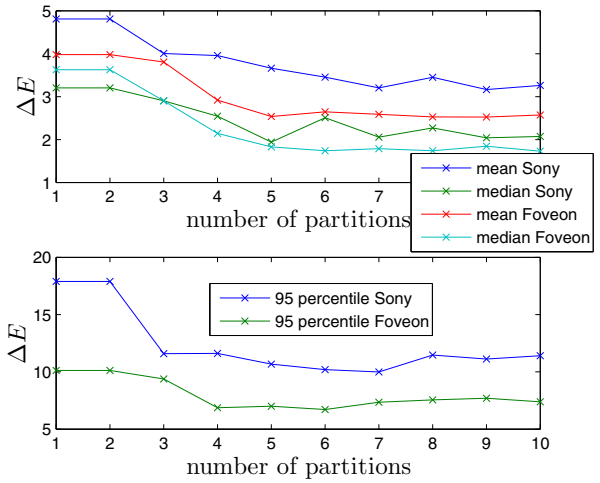

Figure 2. *Mean, median and 95 percentile ΔE errors for the increasing number of hue partitions for the SG chart dataset.*

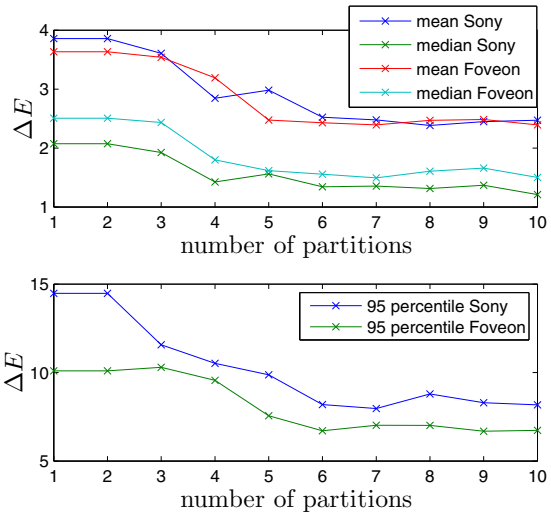

Figure 3. *As in Fig. 2, but for the DC chart.*

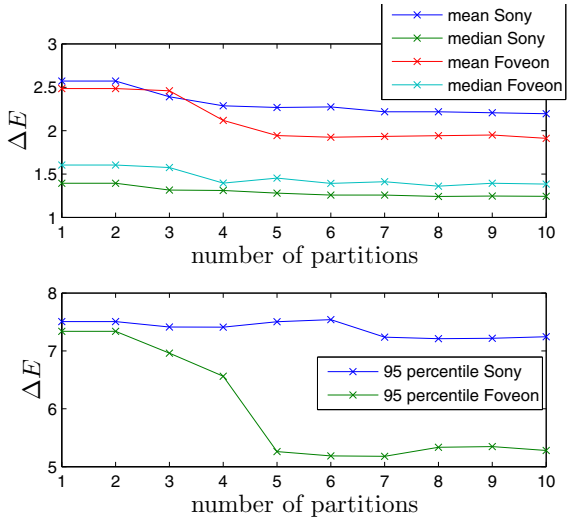

Figure 4. *As in Fig. 3, but for the SFU dataset.*

From these three figures, we can clearly see a decreasing trend of ΔE errors as the number of hue partitions increases. The strength of this trend depends on the dataset and the sensor set used. The most pronounced error reductions can be observed for the SG and DC chart datasets. As to the SFU dataset, the decrease in mean and median errors is more modest, but for this dataset these statistics are already very low leaving less scope for significant improvements. This said, we can observe a stronger decrease in the 95 percentile error at least for the Foveon sensors. This figure is important as it is usually the higher errors occurring only in few samples that produce strong visible effects.

From Figs 2 to 4, we can also see that the error trends usually level off after around six hue partitions. This is why we have chosen this number of hue partitions for further comparisons with the prior art algorithms.

These comparisons can be seen in Table 1. We present the same three statistics of ΔE errors for the two sensor sets and the three reflectance datasets.

We can make the following observations from these results. We can see that HPPCC-CLSQ-6 always performs better than the linear colour correction (LCC) and the original HPPCC. The advantage of the new method over the HPPCC is particularly visible for the largest SFU dataset as well as for all the Foveon results. This was to be expected as the HPPCC SFU dataset results expose the fact that this method was using only a small number of reflectances of this very large dataset (in fact $12+1=13$) to calculate the colour correction transform. Therefore, its results had to be worse than the results of all the other methods that utilised all (putting aside cross-validation) dataset samples. Our new method improves on this by enforcing all dataset samples to be used in the model fit while maintaining the flexibility of the HPPCC in optimising colour correction in each hue partition separately.

As to the second point regarding the Foveon sensor results, the significant improvement of HPPCC-CLSQ-6 over HPPCC

Table 1. Synthetic data characterisation results. The errors obtained 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.
Sony									
PCC,1 (LCC)	4.8	3.2	18	3.9	2.1	14	2.6	1.4	7.5
PCC,2	4.0	2.7	11	3.1	1.8	10	2.4	1.4	7.2
PCC,3	3.2	2.2	7.5	2.4	1.3	7.7	1.9	1.2	6.5
LUT	2.7	1.9	**7.3**	**1.8**	**1.0**	**6.7**	**1.5**	**1.0**	**5.0**
RPCC,2	2.8	1.8	8.7	2.4	1.3	8.5	2.1	1.2	7.0
RPCC,3	**2.6**	**1.6**	8.1	2.2	1.3	7.5	1.8	1.2	6.1
HPPCC	3.8	2.2	14	2.8	1.4	10	3.6	1.8	14
HPPCC-ConL-6	3.5	2.5	10	2.5	1.3	8.2	2.3	1.3	7.5
Foveon									
PCC,1 (LCC)	4.0	3.6	10	3.6	2.5	10	2.5	1.6	7.3
PCC,2	2.9	2.4	6.8	2.9	2.2	8.6	2.2	1.7	5.7
PCC,3	2.2	1.6	5.8	**2.1**	**1.6**	**6.3**	**1.9**	1.5	**5.0**
LUT	3.4	2.7	9.9	2.7	1.7	8.7	2.0	**1.4**	5.2
RPCC,2	3.0	2.3	7.1	2.9	1.8	7.0	2.0	**1.4**	5.5
RPCC,3	**2.0**	**1.6**	**5.1**	2.2	1.4	6.4	**1.9**	**1.4**	**5.0**
HPPCC	4.9	2.7	19	5.4	3.4	18	3.3	2.2	9.2
HPPCC-CLSQ-6	2.6	1.7	6.7	2.4	1.6	6.7	**1.9**	**1.4**	5.2

stems from the fact that the responses of this camera are very much desaturated and require more 'aggressive' colour correction. The desaturated samples that are relatively close to the achromatic centre pose a problem for the HPPCC method as for this sensor set it is more difficult than for the Sony sensor to choose the appropriate samples that would define the shape of the colour correction transform in each partition of the hue space.

The above results also show that the results of our new method are in most cases comparable to the state-of-art root-polynomial colour correction (RPCC). It is true that our new method produces slightly higher error for some high order polynomial (PCC) or the 3-D LUT method. However, these methods do exhibit certain earlier discussed drawbacks. These include the lack of exposure invariance and the lack of hue plane preservation. We note that the root-polynomial colour correction is invariant to exposure, but does not preserve hue planes.

Real data experiments

Here, we use the two real camera datasets (Nikon D70 and Sigma SD15) that we described earlier [5]. The two datasets have been captured as follows. We placed the X-rite SG colour checker in a viewing box, which was illuminated by a D65 metamer [12]. The illumination was provided by Gamma Scientific RS-5B LED illuminator [18]. The scene containing the colour checker was imaged with both cameras set to RAW image capture mode. Photo Research PR-650 spectrophotometer was used to measure the *XYZ*s of the 96 patches of the colour checker. We used DCRAW [1] (Nikon) and PROXEL X3F [2] (Sigma) to extract the 16-bit linear images. We calculated the average RGB of each manually segmented colour checker patch. The dark frames were captured with the lens cap on and then subtracted from the average camera responses. We used the average *RGB* responses and the measured *XYZ*s to colour characterise the two cameras according to all the models described in the above section. The models were also validated and evaluated in the same way as those derived from the synthetic data. The results of these experiments can be seen in Table 2.

Table 2. Nikon D70 and Sigma SD15 characterisation results. The errors statistics are given as in Table 1.

	Nikon D70			Sigma SD15		
model type	mean	med	95 pt.	mean	med	95 pt.
PCC,1 (LCC)	2.5	2.3	5.1	5.2	4.0	16
PCC,2	2.1	1.8	4.8	3.9	3.2	9.4
PCC,3	1.7	1.6	**3.1**	**3.1**	**2.3**	**7.0**
LUT	1.7	1.4	4.4	4.3	3.3	12
RPCC,2	1.9	1.7	4.0	3.8	2.8	8.7
RPCC,3	**1.6**	**1.3**	3.4	3.2	2.4	9.1
HPPCC	2.5	2.1	6.0	5.0	3.2	17
HPPCC-CLSQ-6	2.0	1.7	4.2	3.7	2.5	10

We can see that the results very closely follow the trends we could see in the synthetic data experiments. HPPCC-CLSQ-6 is always better than the linear colour correction as well as the original HPPCC. It also provides a comparable performance to the remaining methods including root-polynomial colour correction. As for the synthetic data, the improvement in performance is the most significant for the camera that initially produced the worst

[1] http://www.cybercom.net/~dcoffin/dcraw/

[2] http://www.proxel.se/x3f.html

results.

Conclusions

Hue Plane Preserving Colour Correction using Constrained Least-Squares is based on the earlier method that utilised the split of the linear colour correction transform into a set of transforms each to be applied in a different region of the colour space (hue slice). Unlike its predecessor the method is robust i.e. it provides significant improvement over linear colour correction for all tested cameras and datasets. The method benefits from being exposure invariant as well as preserving hue planes. It provides results that are comparable to state-of-art root-polynomial colour correction.

References

[1] C. F. Andersen and J. Y. Hardeberg. Colorimetric characterization of digital cameras preserving hue planes. In *Proceedings of the 13th Color and Imaging Conference (CIC)*, pages 141–146, 2005.

[2] K. Barnard, L. Martin, B. Funt, and A. Coath. A dataset for color research. *Color Research and Application*, 27(3):147–151, 2002.

[3] R. S. Berns and M. J. Shyu. Colorimetric characterization of a desktop drum scanner using a spectral model. *J. Electronic Imaging*, 4:360–372, 1995.

[4] T.L.V. Cheung and S. Westland. Colour camera characterisation using artificial neural networks. In *IS&T/SID Tenth Color Imaging Conference*, volume 4, pages 117–20. The Society for Imaging Science and Technology, The Society for Information Display, 2002.

[5] G. Finlayson, M. Mackiewicz, and A. Hurlbert. Colour correction using root-polynomial regression. *IEEE Trans. Image Processing*, 24(5):1460–1470, 2015.

[6] G. D. Finlayson and M. Drew. Constrained least-squares regression in color spaces. *Journal of Electronic Imaging*, 6(4):484–493, 1997.

[7] E. Garcia, R. Arora, and M. R. Gupta. Optimized regression for efficient function evaluation. *IEEE Trans. Image Processing*, 21(9):4128–4140, 2012.

[8] P. Green and L. W. MacDonald, editors. *Colour engineering: achieving device independent colour*. Wiley, 2002.

[9] Guowei Hong, M. Ronnier Luo, and Peter A. Rhodes. A study of digital camera characterisation based on polynomial modelling. *Color research and application*, 26(1):76–84, 2001.

[10] P. Hubel. Foveon technology and the changing landscape of digital cameras. In *Proceedings of the 13th Color and Imaging Conference (CIC)*, pages 314–317, 2005.

[11] P.C. Hung. Colorimetric calibration in electronic imaging devices using a look-up tables model and interpolations. *Journal of Electronic Imaging*, 2(1):53–61, 1993.

[12] R.W.G. Hunt and M. R. Pointer. *Measuring Colour*. Wiley, fourth edition, 2011.

[13] H. E. Ives. The transformation of color-mixture equations from one system to another. *J. Franklin Inst.*, 16:673–701, 1915.

[14] H.R. Kang. Colour scanner calibration. *Journal of Imaging Science and Technology*, 36:162–70, 1992.

[15] H.R. Kang and P.G. Anderson. Neural network application to the color scanner and printer calibration. *Journal of Electronic Imaging*, 1:125–34, 1992.

[16] S. H. Lim and A. Silverstein. Spatially varying colour correction matrices for reduced noise. Technical Report HPL-2004-99, Imaging Systems Laboratory, HP Laboratories, Palo Alto, CA, US, June 2004.

[17] R. T. D. Luther. Aus dem Gebiet der Farbreizmetric. *Zeitschrift für technische Physic*, 8:540–555, 1927.

[18] M. Mackiewicz, S. Crichton, S. Newsome, R. Gazerro, G. Finlayson, and A. Hurlbert. Spectrally tunable led illuminator for vision research. In *Proceedings of the 6th Colour in Graphics, Imaging and Vision (CGIV)*, volume 6, pages 372–377, Amsterdam, Netherlands, April 2012.

[19] Bui Tuong Phong. Illumination for computer generated pictures. *Communications of ACM*, 18.

[20] M. J. Vrhel. *Mathematical methods of color correction*. PhD thesis, North Carolina State University, Department of Electrical and Computer Engineering, 1993.

[21] B. Wandell. *Foundations of Vision*. Sinauer Associates, Inc, 1995.

[22] G. Wyszecki and W.S. Styles. *Color Science: Concepts and Methods, Quantative Data and Formulae*. John Wiley and Sons, NY, 1982.

[23] Li Xinwu. A new color correction model based on bp neural network. *Advances in Information Sciences and Service Sciences*, 3(5):72–8, June 2011.