

The Alternating Least Squares Technique for Nonuniform Intensity Color Correction

Graham D. Finlayson, Maryam Mohammadzadeh Darrodi,*
Michal Mackiewicz

School of Computing Sciences, University of East Anglia, Norwich, NR4 7TJ, United Kingdom

Received 26 November 2013; revised and accepted 9 April 2014

Abstract: Color correction involves mapping device *RGBs* to display counterparts or to corresponding *XYZs*. A popular methodology is to take an image of a color chart and then solve for the best 3×3 matrix that maps the *RGBs* to the corresponding known *XYZs*. However, this approach fails at times when the intensity of the light varies across the chart. This variation needs to be removed before estimating the correction matrix. This is typically achieved by acquiring an image of a uniform gray chart in the same location, and then dividing the color checker image by the gray-chart image. Of course, taking images of two charts doubles the complexity of color correction. In this article, we present an alternative color correction algorithm that simultaneously estimates the intensity variation and the 3×3 transformation matrix from a single image of a color chart. We show that the color correction problem, that is, finding the 3×3 correction matrix, can be solved using a simple alternating least-squares procedure. Experiments validate our approach. © 2014 Wiley Periodicals, Inc. *Col Res Appl*, 40, 232–242, 2015; Published Online 24 May 2014 in Wiley Online Library (wileyonlinelibrary.com). DOI 10.1002/col.21889

Key words: color correction; nonuniform intensity; alternating least squares; characterization; shading field

INTRODUCTION

Color correction is the process by which device *RGBs* are mapped to a device independent space such as, for example, display *RGBs* (*sRGB*¹) or as in this article, *XYZ* tristimulus values.²

Interpolation,³ look up tables,⁴ polynomial or root-polynomial regression,^{5,6} artificial neural networks,⁷ and methods based on the spectral reconstruction from a set of basis functions⁸ can be used to map the camera *RGB* responses to *XYZ* values. Some similar techniques have been applied for scanner characterization.^{8–11} A linear mapping from *RGB* to *XYZ* is achieved through a 3×3 linear transformation matrix *M*. Mathematically, linear color correction can be written as:

$$\mathbf{PM} \approx \mathbf{Q} \quad (1)$$

where *P* is a $N \times 3$ set of raw linear camera (or scanner) *RGB* responses for *N* color patches and *Q* is the corresponding $N \times 3$ matrix of corresponding *XYZ* triplets. The color correction problem consists of finding the matrix *M* such that *M*: *RGB* → *XYZ*. The best mapping to *XYZs* could be found by minimizing the following expression.

$$\text{Min}_{\mathbf{M}} \|\mathbf{PM} - \mathbf{Q}\| \quad (2)$$

While linear color correction generally works well, it turns out that finding the correct transformation matrix is not easy. Indeed, to solve (2) we need to have the *RGBs* measured in a real scene and imaged in the same lighting and viewing geometry as the reference *XYZs*. If for example one side of the *RGB* image of a reference chart was darker—due to lighting variation—than the other, then we could only solve for (2) if the measured *XYZs* showed the same shading profile. Thus, the matrix *M* would map the captured *RGBs* with the shading profile to the corresponding measured *XYZs* with the same shading profile. However, in this article we address the situation when we do not have access to the measured *XYZs* at this particular scene, but only to the *XYZs* measured by the manufacturers under some standard illuminant. In Fig. 1(a), we show an image of a color checker captured under nonuniform light. It is clear that there is a significant intensity variation. Thus, the basic premise that the *RGBs* acquired in this image correspond to the measured *XYZs* (under

*Correspondence to: Maryam Mohammadzadeh Darrodi (e-mail: m.darrodi@uea.ac.uk)

This is an open access article under the terms of the Creative Commons Attribution License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

© 2014 The Authors *Color Research & Application* Published by Wiley Periodicals, Inc.




Fig. 1. Illustration of (a): cropped Macbeth color chart image, (b): cropped image of the gray chart showing the non-uniform shading, (c): color chart in (a) after dividing out by (b).

different light geometry) does not hold. Of course, we could divide the *RGB* image by the shading field [Fig. 1(b)] to give the shading corrected image [Fig. 1(c)]. Yet, now we need two reference charts to carry out color calibration rather than one.

Recently, Funt and Bastani¹² proposed that it is not necessary to take a second image of a gray patch if the optimization problem is changed. Specifically, they propose finding a 3×3 matrix by minimizing a shading independent quantity. For example, through search we can find a 3×3 matrix that maps *RGBs* to corresponding *XYZs* such that the angular error is minimized. Significantly, they report that color correction performance can be almost twice as good when the correction matrix is found in an intensity independent way compared with using (2) directly (without dividing out the shading).

In this article, we wish to build on the Funt and Bastani result. First, we note that by minimizing the angular error (between mapped *RGBs* and corresponding *XYZs*) that of the $N \times 3$ *RGB* measurements they effectively use only $N \times 2$ degrees of freedom (the orientation of a vector—on which the angular error depends—is parameterized by two numbers).¹² Returning to Fig. 1 where the 24 patch Macbeth color checker is used then in their method the 9 terms of the 3×3 correction matrix are calculated with — effectively — only the 24×2 parameters necessary to represent the orientations of the *RGBs*. In fact the situation is, arguably, even a little more austere than this. There is no extra information to be gained from using all 6 achromatic patches. Thus, there are really only $(19) \times 2$ independent measurements.

In this article, we wish to use more of the data available or, put another way, to discard less information in solving for the correction matrix. We begin by modelling the light intensity variation directly, by writing the minimisation as follows:

$$\text{Min}_{\mathbf{D}, \mathbf{M}} \|\mathbf{DPM} - \mathbf{Q}\|. \quad (3)$$

Here, \mathbf{D} denotes an $N \times N$ diagonal matrix. Each element of the diagonal matrix models a shading correction for one of the patches of the color chart (premultiplying by a diagonal matrix scales the rows of \mathbf{P} by the reciprocal of the shading field). In spirit, (3) is like the Funt and Bastani formalism. Although, advantageously we are not thinking of the target *XYZs* as orientated vectors but as

the full 3-dimensional colors. As we will see in section “Alternating Least Squares Color Correction,” Equation (3) also admits a straightforward alternating least-squares solution strategy.

However, should we really think of the shading for one color as being independent from another? Clearly, in Fig. 1 while the shading varies across an image, proximate pixels have similar shading. Modeling shading and incorporating it into the minimization are the key contributions of this article. We show that shading can be modelled by the 2D Discrete Cosine Transform (DCT¹³) basis functions. The shading field is smooth and therefore we can suspect that it can be modelled using relatively few DCT basis. In the experimental section we show that by taking linear combinations of around 15 bases, we can model the typical shadings that occur across the color chart image. We then incorporate this basis idea into our minimization.

In section “Background,” we present the background on linear color correction. Our new alternating least squares algorithm for the joint estimation of the color correction matrix and shading is presented in section “Alternating Least Squares Color Correction.” Experiments demonstrate the utility of our method in section “Experiments and Results.” The article concludes in section “Conclusions.”

BACKGROUND

Equation (1) is an over-determined system of equations, which is usually solved by the least squares regression,¹⁴

$$\mathbf{M} = \mathbf{P}^+ \mathbf{Q}, \quad \mathbf{P}^+ = [\mathbf{P}^t \mathbf{P}]^{-1} \mathbf{P}^t \quad (4)$$

where \mathbf{P}^+ denotes the Moore-Penrose pseudo-inverse¹⁵ and t the matrix transpose. From a statistical point of view, we assume that the *XYZs* in Equation (2) are measured without error. However, Marimont and Wandell provide an alternate “total least-squares” minimization where measurement error can lie in the camera *RGBs* or measured *XYZs*.¹⁶ Vrhel and Trussel carried out the minimization indirectly.¹⁷ First, Principal component analysis is used to find a 3D basis to model spectral reflectance. With respect to this, 3D assumption color correction is exactly a linear transform. Finlayson et al. proposed a constrained least-squares regression, where the 3×3

linear transform is constrained to map one (or possibly 2) color patches exactly.¹⁸ The authors note that usually it is beneficial to map the white patch without error, hence the name of their algorithm, the white preserving color correction. Although the constrained regression in general returns a higher mean error than standard linear color correction, the authors maintain that it performs particularly well when there is incomplete training set and so called maximum ignorance training is performed.

An advantage of 3×3 linear mapping is that when the exposure changes (e.g., an *RGB* becomes darker) the resulting *RGBs* are correctly mapped to the corresponding *XYZs*.

$$\mathbf{kPM} \approx \mathbf{kQ} \quad (5)$$

In contrast, with the exception of the root-polynomial method,⁵ other nonlinear color correction schemes (in particular polynomial color correction)⁶ work less well when exposure varies (the same object color might be mapped to different *xy* chromaticities if, say, the object is viewed both in a well exposed and under exposed part of the image). Andersen and Hardeberg¹⁹ proposed an improvement to the linear color correction with a view to tackle the aforementioned nonlinear color correction exposure dependence problem. Their proposition of the mapping from *RGB* to *XYZ* uses a set of linear transforms, where each transform is applied in a different region (hue slice) of the color space. Their method improves on the linear color correction, while maintaining its major advantage, that is, (5) still holds.

Nevertheless, none of the above color correction algorithms addresses the problem of nonuniform lighting variation (unless the shading field is known) except the Funt and Bastani¹² method mentioned in the Introduction. That method determines the 3×3 color correction matrix that minimizes the sum of angles between the mapped *RGBs* and the known *XYZ* values. Thus, instead of accounting for both direction and magnitude of color vectors as is the case in standard linear color correction, they find \mathbf{M} by minimizing the following

$$\text{Min}_{\mathbf{M}} \sum_i \text{angle}([\mathbf{PM}]_i, \mathbf{Q}_i). \quad (6)$$

However, there is no closed form solution to (6). Rather, a “good”—often nonglobally optimal—solution is found by applying standard search algorithms (e.g., Matlab’s *fminsearch*²⁰). This search strategy is complex. Indeed, the complexity of the algorithm makes it an unlikely candidate for implementation in a digital camera. Moreover, minimizing the angular error effectively discards some information, which might be used to find a more precise solution.

ALTERNATING LEAST SQUARES COLOR CORRECTION

We outlined our proposition of the intensity independent color correction in the Introduction. We model the light

intensity variation directly by introducing the diagonal matrix \mathbf{D} to the color correction equation:

$$\mathbf{DPM} \approx \mathbf{Q}.$$

Recall that we are minimizing:

$$\text{Min}_{\mathbf{D}, \mathbf{M}} \|\mathbf{DPM} - \mathbf{Q}\|.$$

Unlike the standard color correction, this optimisation does not have a least-squares closed form solution for \mathbf{D} and \mathbf{M} . Rather an alternating least-squares procedure is used. Here, we solve for \mathbf{M} by keeping \mathbf{D} fixed and then solve for \mathbf{D} given \mathbf{M} . We alternate between solving for \mathbf{D} and \mathbf{M} (and at each stage we update our estimate). The principle of alternating methods is that parameters are iteratively improved in turn, until they converge to optimum good solution.^{21,22} Importantly, alternating least squares is a procedure that is guaranteed to converge. However, the alternating least square approach, like searching, is not guaranteed to find the global optimum.

A step-by-step algorithm to solve (3), in which both parameter matrices \mathbf{D} and \mathbf{T} are updated while keeping the other fixed, is given below. \mathbf{P}^k denotes matrix \mathbf{P} after iteration k and \mathbf{P}^0 is the initial matrix \mathbf{P} .

Algorithm 1:

1. Update \mathbf{D}

$$\mathbf{D} \leftarrow \begin{bmatrix} d_{11} & \cdots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \cdots & d_{NN} \end{bmatrix}, \quad (7)$$

where

$$d_{jj} = \frac{p_j \cdot q_j^T}{\|p_j\|^2}, \quad j = 1, \dots, N, \quad (8)$$

and p_j and q_j represent the j th row of \mathbf{P}^{k-1} and \mathbf{Q} (i.e., \mathbf{D} is the optimal diagonal transform in a least-squares sense between \mathbf{P}^{k-1} and \mathbf{Q}).

2. Update \mathbf{M}

$$\mathbf{M} = (\mathbf{DP}^0)^+ \mathbf{Q}. \quad (9)$$

3. Update \mathbf{P}

$$\mathbf{P}^k = \mathbf{P}^0 \mathbf{M}. \quad (10)$$

4. Repeat steps 2–4 until convergence.

Usually, about 15 iterations of the algorithm suffice to produce a stable solution for matrices \mathbf{D} and \mathbf{M} . The shading profile cannot be recovered precisely and consequently it is recovered with some error. This is the case when there is no shading and when there is. When there is no shading, applying the diagonal matrix \mathbf{D} should have no influence (ideally it should be equal to identity); and when there is shading the elements of matrix \mathbf{D} will

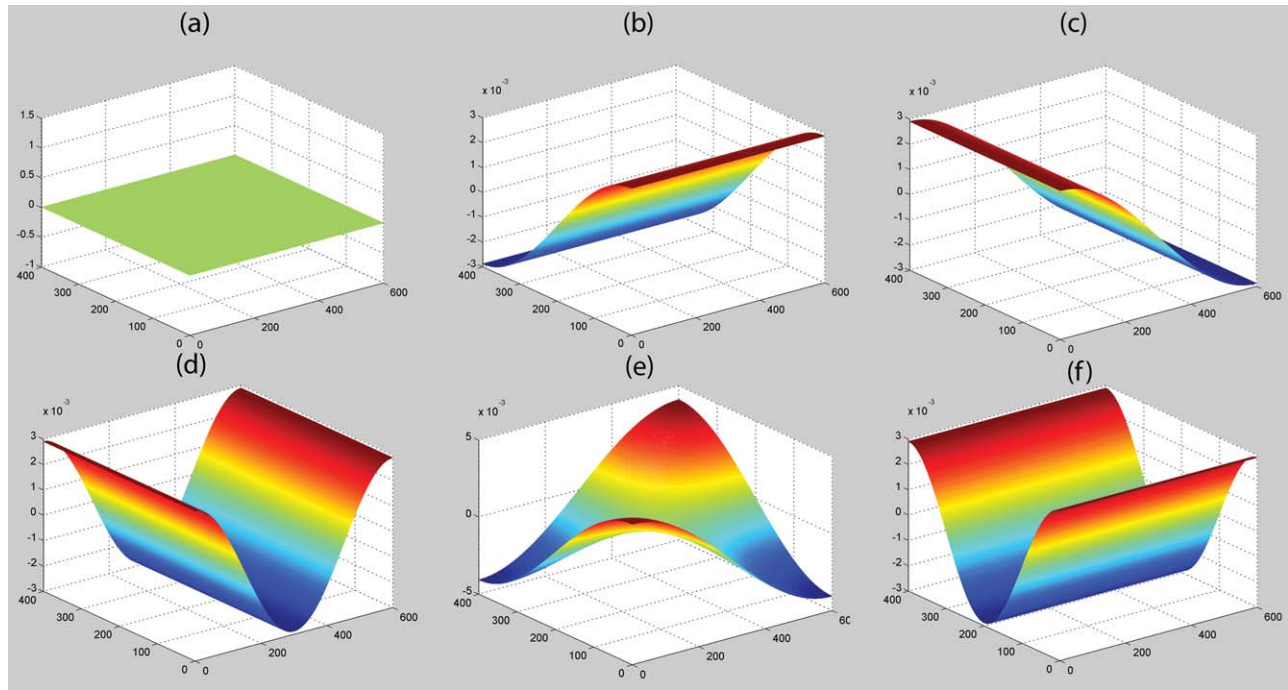


Fig. 2. Mesh plots of the first six 2-D DCT basis functions used in ALS-DCT. The weighted linear combination of these basis functions are used to approximate a shading field.

approximate shading in the corresponding points. However, the N diagonal elements of matrix \mathbf{D} do not provide the smooth shading field across the scene as \mathbf{D} approximates the intensity correction patch by patch. Below, we develop an algorithm targeted at recovering the shading field with the smoothness constraint. Our motivation is three-fold. First, there is strong correlation—for real shading affected images—between the best intensity correction terms for color patches and their close neighbours. Second, modeling the shading field with few parameters means (as we show below) that there are less unknowns to solve for (and this means our ALS minimization converges more rapidly). Finally, as we show in the experimental section, there is no advantage of shading correction per patch compared with the modeling of the shading field as a whole across the calibration target.

In this approach, we model smoothness in the shading field pixel by pixel. This is formulated as

$$\mathbf{J} = \sum_{k=1}^K w_k \mathbf{G}_k, \quad (11)$$

where \mathbf{J} is the $E \times F$ shading matrix (reciprocal of the shading field) representing the $E \times F$ image of intensity variation across the scene, K is the number of 2-D DCT basis functions (six first DCT basis can be seen in Fig. 2), G_k represents the k th, $E \times F$ DCT basis, and w_k is a scalar representing its corresponding weight. See Appendix for further details on obtaining vector \underline{w} given an image of the color chart.

Analogously to Algorithm 1, we can solve for the best transformation matrix \mathbf{M} and the shading matrix \mathbf{J} using a similar alternating least squares method as presented

below in the Algorithm 2. Each step of the Algorithm 1 has a corresponding step in the Algorithm 2.

Let R_i denote a three dimensional row vector of RGB s at i th pixel of the color chart image \mathbf{R} , J_i is the intensity value of the corresponding pixel in matrix \mathbf{J} and \mathbf{P} is a $N \times 3$ matrix containing RGB values averaged over the center area of each patch in the color chart image \mathbf{R} . As in Algorithm 1, the upper index denotes the iteration number, thus \mathbf{R}^0 denotes the initial color chart image \mathbf{R} .

Algorithm 2:

1. Update \mathbf{J} (from \mathbf{R}^{k-1} and DCT basis \mathbf{G}) by determining was described in the appendix.
2. Update \mathbf{M}
 - a. Update $\mathbf{R}, \mathbf{R}^k \leftarrow \mathbf{J} * \mathbf{R}^{k-1}$ where $*$ denotes pixel by pixel multiplication.
 - b. Update $\mathbf{P}, \mathbf{P}^k \leftarrow$ average RGB for each color patch in \mathbf{R}^k .
 - c. Update \mathbf{M} :

$$\mathbf{M} = (\mathbf{P}^k)^+ \mathbf{Q}. \quad (12)$$

3. Update \mathbf{R}

$$\mathbf{R}_i^k = \mathbf{R}_i^0 \mathbf{M}, \quad i = 1, \dots, E \times F. \quad (13)$$

4. Repeat steps 2–4 until convergence.

Empirically, we found that this algorithm converges in about 15 iterations (for 21 DCT basis) which is independent of the size of matrices \mathbf{P} and \mathbf{Q} . When the mean least squares (Eq. (3)) error difference between the iterations is less than a very small amount ($0.5 \Delta E$ values) then we indicate convergence. This is a much faster algorithm

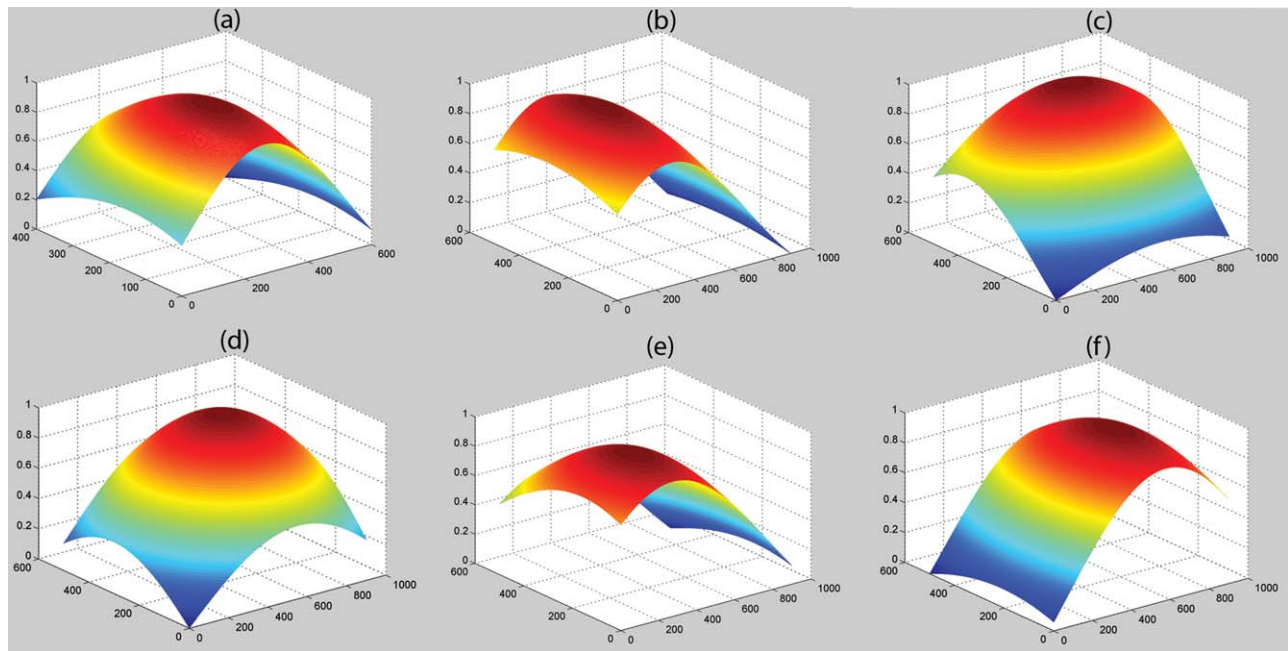


Fig. 3. Mesh plots of six synthetically generated shading fields with shading ratios of 3.3, 4.6, 3.2, 3.9, 2.1, and 6.7 for shading images (a–f).

than angle minimization. In the next section, both alternating least squares methods are applied to synthetic and real image data and are compared with other methods.

EXPERIMENTS AND RESULTS

We compared the performance of our proposed algorithms with the standard linear color correction and Funt and Bastani method in both synthetic and real data experiments. As to the former, we used the Sony DXC-930 camera sensor sensitivities²³ to integrate the spectral data from two surface reflectance datasets. The first dataset comprised 96 reflectances of the Xrite SG color

checker (border patches excluded) and the second dataset contained 180 patches of the Macbeth DC color checker (again border patches were excluded). For each dataset, we integrated the Sony sensor sensitivities and the color matching functions under D65 illuminant²⁴ producing corresponding sets of camera responses (*RGBs*) and *XYZs*. Spectra calculations were carried out for 31 spectral channels: 400–700 nm sampled every 10 nm. For each of the two datasets we created the set of 20 synthetically generated shading fields (six of these can be seen in Fig. 3). Those shading fields were generated by randomly varying a center and spread of the normal distribution. The shading fields have the ratio of the maximum pixel value to the minimum value between 1.7

TABLE I. CIELAB ΔE errors obtained after color correcting the SG and DC chart images averaged over 20 different synthetic shading recoveries.

Method	SG chart			DC chart		
	Mean ΔE_{ab}^*	Median ΔE_{ab}^*	Max ΔE_{ab}^*	Mean ΔE_{ab}^*	Median ΔE_{ab}^*	Max ΔE_{ab}^*
Least squares	5.4	3.9	38	4.3	2.5	53
Angle minimization	4.2	2.2	25	3.1	2.2	25
ALS-Diagonal matrix	3.7	2.8	17	2.9	1.9	23
ALS-DCT Basis						
1 Basis	5.4	3.9	38	4.3	2.5	53
3 Basis	4.9	3.7	27	4.0	2.2	54
6 Basis	4.1	3.0	22	3.1	1.6	25
10 Basis	4.0	2.9	19	3.0	1.6	24
15 Basis	3.7	2.7	18	2.9	1.6	23
21 Basis	3.7	2.7	18	2.9	1.6	22
28 Basis	3.7	2.7	18	2.9	1.6	22
Uniform lighting LSQ	3.6	2.6	17	2.9	1.6	22

The transformation matrices were tested on the ground truth image. Five methods are compared: Least squares, angle minimization, alternating least squares with the diagonal matrix, alternating least squares with 2-D DCT using 1, 3, 6, 10, 15, and 21 basis functions and the least squares on the ground truth image (shading field removed).

1	2	6	7	15	16	0	0
3	5	8	14	17	0	0	0
4	9	13	18	0	0	0	0
10	12	19	0	0	0	0	0
11	20	0	0	0	0	0	0
21	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0
0	0	0	0	0	0	0	0

Fig. 4. Zigzag ordering of 2-D DCT image components as used in JPEG compression¹³.

(weak shading) and 13 (strong shading). Next, we multiplied the two color checker images by the generated shading fields obtaining 2×20 images of the color checker with nonuniform lighting. Then, for each of these 2×20 color checker images, we calculated the color correction matrices using standard linear color correction, Funt and Bastani method and our two alternating least squares methods. We also calculated two additional color correction matrices from the two color checker images before any of the shading fields were applied, that is, with a uniform lighting field. These two images were also used as the ground truth in our experiments, that is, all the above methods were tested on these two uniform lighting images. For each color correction matrix \mathbf{M} , we calculated the mean, median and max ΔE in the CIELab color space²³ from all patches in the test color chart giving us the above three statistics for each of the 20 training color charts. The average of those three figures across the 20 images can be seen in the rows of Table I. The reader can see that the Algorithm 2 was tested with a varying number of DCT basis, ranging from 1 to 21 (the actual numbers of basis were chosen according to the DCT zigzag pattern of 1, 3, 6, 10, 15, and 21 derived




Fig. 6. (a) Illustration of synthetic SG chart with shading field [Fig. 3(a)] applied; (b) shading field removed using ALS-DCT with 21 basis functions.

from Fig. 4¹³). From comparing the rows of Table I, it is clear that a single DCT basis is unable to correct for the nonuniform lighting field and returns the result, which is identical to the standard linear color correction. As the number of DCT basis increases, the results improve reaching the accuracy of the diagonal alternating least squares and also the standard LCC with the uniform lighting. Both alternating least squares methods produce the similar results and both slightly outperform the Funt and Bastani method. Moreover, for the synthetic data, the Algorithm 1 produces the same error statistics for any of the 20 shading fields, that is, the mean of the statistics in the Table I is identical to the 20 separate error statistics. This is not the case for the Funt and Bastani method,¹² where the means of the three statistics given in the two tables are slightly higher than for the Algorithm 1, but the individual result error statistics can be substantially higher, the highest was 7.9, 3.6, and 42 for the mean, median and max ΔE (SG chart) and 3.7, 2.6, and 44 (DC chart). The results for the alternating least squares with the basis functions vary only slightly between different shading fields and are in fact very close to the mean results provided in the two tables, for example, for 21 DCT basis, the corresponding figures are 4.1, 3.0, and 20 (SG chart) and 2.9, 1.7, and 23 (DC chart). Figure 5 illustrates the mesh plots of the recovered shading fields for 21 DCT basis and the diagonal alternating least squares algorithms. Images of the SG chart, with shading

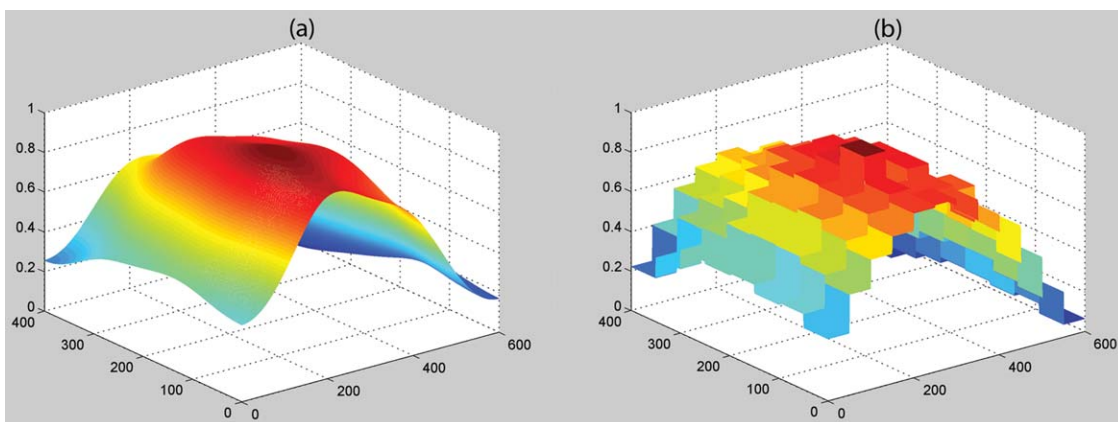


Fig. 5. (a) Mesh plot of the recovered shading field [Fig. 3(a)] for synthetic SG chart using ALS-DCT with 21 basis functions. (b) Mesh plot of the same recovered shading field using ALS with the diagonal matrix.

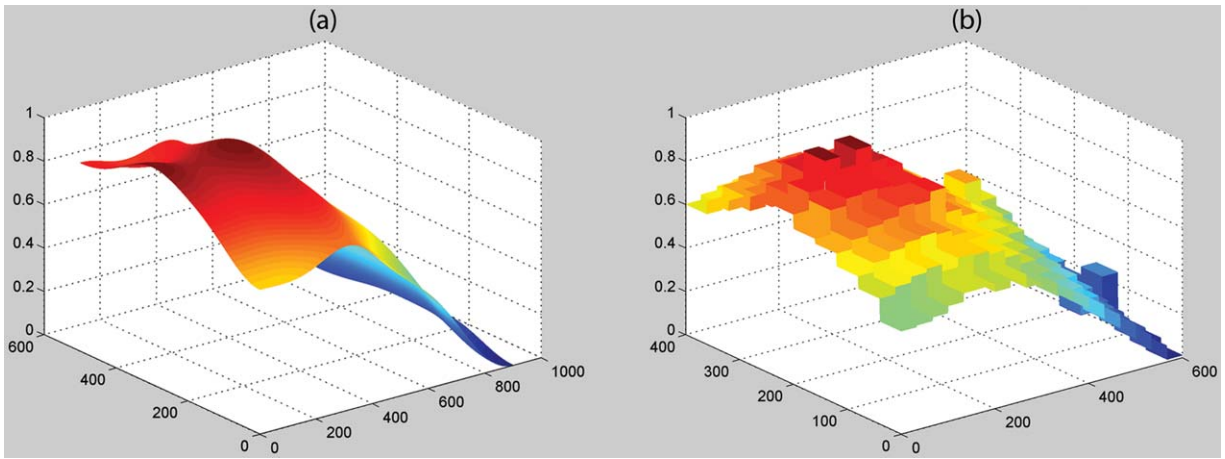


Fig. 7. (a): Mesh plot of the recovered shading field (Figure 3 (b)) for synthetic DC chart using ALS-DCT with 21 basis functions. (b): Mesh plot of the recovered shading field using ALS with the diagonal matrix.

field applied and removed after recoveries are compared in Fig. 6. For the DC chart, the same illustrations are shown in Figures 7 and 8.

Next, we performed the same experiment using the data obtained from the camera. We used Nikon D5100 camera to capture two images of the Macbeth color checker placed in the VeriVide cabinet and lit by the D65 metamer illuminant. We call the two images: test image 1 (with a strong shading gradient-shading ratio of 14 produced by partially blocking the light in the light cabinet see Fig. 9) and test image 2 (with a smoother shading, shading ratio of 2.5). Here, we ignore light fall-off around the edges of the image. For both color chart images, an additional image of the gray chart (positioned at the same location as the color checker) was also taken to represent the real intensity variation. Dividing the color chart *RGBs* by the corresponding intensity values from the gray chart produces the ground truth image (see Fig. 1 for test image 1).

The *XYZ* tristimulus values of each color patch were measured using a Gretag Macbeth SpectroEye spectrophotometer.* The camera responses were then mapped to the *XYZ* tristimulus values using the same algorithms as in the synthetic experiment. The ground truth *RGB* values were multiplied by the color correction matrix




Fig. 8. Illustration of (a): synthetic DC chart with shading field (Figure 10 (b)) applied; (b): shading field removed using ALS-DCT with 21 basis functions.

*This provides measurements under uniform illumination of D65 for each of the color patches.

\mathbf{M} , derived from each algorithm. The algorithms performance comparisons are shown in Table II. The results generally confirm the observations from the earlier experiments. Both alternating least squares methods provide consistently similar results and approach the performance obtained with the color correction matrix obtained from the ground truth. Similarity between the results of Tables I and II means both algorithms reached the same level of performance. Increasing the number of DCTs to 28 did not change the results. This suggests that more DCT functions are not necessary. However, Funt and Bastani angle minimization procedure provides less consistent results for these two difficult cases.

Figure 10 illustrates the mesh plots of the recovered shading fields for 21 DCT basis and the diagonal alternating least squares algorithms for test image1. Image of the Macbeth chart alongside the same image with the shading field removed is shown in Fig. 11. The corresponding figures for test image 2 are shown in Figs. 12, 13 and 14. The figures compare the color chart with and without shading.

In above experiments, the availability of the ground-truth data allowed us to calculate the colorimetric errors and establish that 21 DCT coefficients were enough to model the shading profile. In a real application, when no information is available on the shading profile, an alternative for testing whether enough DCT coefficients were included would be to check the change in the resulting *XYZ* or *CIELAB* values between various (increasing in number of elements) sets of DCT coefficients. Yet another approach would be to assume that the shading profile is not less smooth than in the experiments performed here. Where 21 DCT coefficients were enough for all the images including a very artificial strong gradient in test image 1, which would be highly unlikely to be created in practice. In summary, the number of basis functions needed depends on the expected shape of the shading image across the target. In our




Fig. 9. Mesh plot showing the variation in intensity in the grey chart for test image1. (b): Cropped image of the same grey chart.

TABLE II. CIELAB ΔE errors obtained after color correcting test images 1 and 2.

Method	Test image 1			Test image 2		
	Mean ΔE_{ab}^*	Median ΔE_{ab}^*	Max ΔE_{ab}^*	Mean ΔE_{ab}^*	Median ΔE_{ab}^*	Max ΔE_{ab}^*
Least squares	8.6	7.2	25	4.9	3.9	12
Angle minimization	4.4	3.9	7.3	3.8	3.6	8.9
ALS-Diagonal matrix	4.0	4.2	7.9	2.7	2.9	5.3
ALS-DCT Basis						
1 Basis	8.6	7.2	25	4.9	3.9	12
3 Basis	5.3	4.8	12	4.2	4.1	9.6
6 Basis	4.7	4.6	8.1	3.1	3.1	5.9
10 Basis	4.2	4.5	7.1	3.1	3.1	6.2
15 Basis	4.0	3.7	7.0	3.0	3.1	5.5
21 Basis	4.0	3.7	7.1	3.0	3.3	4.8
28 Basis	4.0	3.6	7.5	3.0	3.2	5.2
Uniform lighting LSQ	3.1	3.2	6.6	2.8	3.0	6.0

The same methods as in Table I were compared.




Fig. 10. (a) Mesh plot of the recovered shading field for test image1 using ALS-DCT with 21 basis functions. (b) Mesh plot of the recovered shading field using ALS with the diagonal matrix.



Fig. 11. (a) Cropped image of test image 1. (b) Shading recovered using ALS-DCT with 21 basis and removed from (a).




Fig. 14. (a) Cropped image of test image 2. (b) Shading recovered using ALS-DCT with 21 basis and removed from (a).

experiments, we deliberately tried to simulate hard cases where the gradients were strong without being unrealistic. For practical purposes, we believe 21 DCT coefficients suffice.

In Table III, the effect of the shading ratios on ΔE values are illustrated for various number of DCT basis on synthetic DC chart. We can see that as the shading profiles become steeper, the ΔE errors increase. This is not visible

for the 21 DCTs since all 20 images had their shading profiles discounted accurately irrespective of their steepness. However, as the number of DCT coefficients decrease, the errors increase, particularly for those images which were altered by steep shading profiles. However, for practical purposes the increase afforded by 21 as opposed to 10 basis functions is modest. We found 10 basis functions supported very good color correction performance.




Fig. 12. Mesh plot showing the variation in intensity in the gray chart for test image 2. (b) Cropped image of the same gray chart.

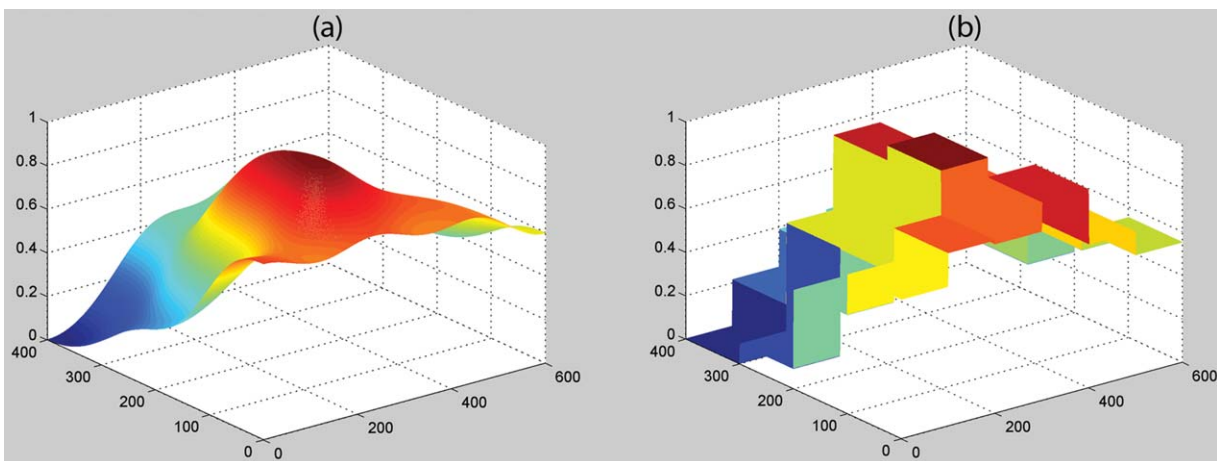


Fig. 13. (a): Mesh plot of the recovered shading field for test image2 using ALS-DCT with 21 basis functions. (b): Mesh plot of the recovered shading field using ALS with the diagonal matrix.

TABLE III. Illustration of the effect of shading ratios on CIELAB ΔE errors for 21, 10, and 3 DCT basis on synthetic DC chart

Gradient label	Shading ratio	DC chart with 21 basis		DC chart with 10 basis		DC chart with 3 basis	
		Mean ΔE_{ab}^*	Max ΔE_{ab}^*	Mean ΔE_{ab}^*	Max ΔE_{ab}^*	Mean ΔE_{ab}^*	Max ΔE_{ab}^*
13	1.8	2.9	22.4	2.9	22.2	3.2	24.4
8	2.0	2.9	22.4	2.9	22.3	3.1	21.7
5	2.1	2.9	22.4	2.9	22.2	3.2	22.9
3	2.1	2.9	22.4	2.9	22.0	3.8	48.2
11	2.1	2.9	22.4	3.0	22.2	4.0	55.1
20	2.2	2.9	22.4	2.9	22.2	3.2	24.4
2	2.6	2.9	22.4	2.9	22.4	3.2	21.8
6	2.9	2.9	22.5	3.0	23.1	3.4	22.1
9	2.9	2.9	22.3	3.1	25.5	4.9	92.4
1	3.0	2.9	22.4	3.2	27.3	5.7	122.2
15	3.1	2.9	22.4	3.0	22.2	3.3	23.9
16	3.2	2.9	22.5	3.0	24.1	4.2	52.3
12	3.3	2.9	22.4	3.0	22.5	5.3	110.6
4	3.4	2.9	22.4	3.1	25.5	4.9	90.1
18	3.6	2.9	22.4	3.2	27.5	4.7	75.2
7	3.6	2.9	22.4	3.0	23.7	5.4	115.6
10	3.9	2.9	22.5	3.0	23.0	3.4	20.9
19	4.6	2.9	22.6	3.0	23.3	3.5	20.7
17	5.0	2.9	22.4	3.1	24.8	3.9	42.8
14	6.7	2.9	22.4	3.1	26.7	4.6	69.6

CONCLUSIONS

In this article, we introduced the family of alternating least squares algorithms as a new approach for camera color correction. These algorithms are particularly useful when the reference color chart is illuminated non-uniformly as they remove the need for an additional image of the gray chart. The key to these methods is that the two unknown parameters (the shading and the color correction matrix) are solved for in turn, until convergence. Two variants of the algorithm have been proposed: first, modelling the shading with a diagonal matrix and second, decomposing the shading using 2D DCT basis functions. The latter should be the favourite as it incorporates the constraint on the intensity field smoothness and hence it requires estimating fewer parameters. We have shown that 15 DCT basis suffice to model the shading field. The diagonal matrix alternating least squares would require estimating as many parameters as there are patches in the color chart, which for the larger color charts such as DC color checker becomes a significantly higher number. Although the experiments presented here involved only linear color correction, both alternating least squares methods could be applied for root-polynomial and polynomial color corrections, given enough data are available. In particular, the basis decomposition method should be more suitable for adopting here for the same reason that was mentioned above.

ACKNOWLEDGMENT

The synthetic shading fields addressed in the article are available for download at: <https://sites.google.com/site/rankbasedspectralestimationuea/data-and-codes>.

1. IEC 61966-2-1:1999: Colour measurement and management in multi-media systems and equipment – Part 2-1: Default RGB colour space –

sRGB. Geneva, Switzerland: International Electrotechnical Commission; 1999.

2. Wyszecki G, Stiles WS. Color science, concepts and methods, quantitative data and formulas. New York: Wiley; 1967.
3. Hung PC. Color rendition using three-dimensional interpolation. Imaging Applications in the Work World. Proc SPIE 1988;900: 111–115.
4. Hung P. Colorimetric calibration in electronic imaging devices using a look-up-table model and interpolations. J Electron Imaging 1993;2:53–61.
5. Finlayson G, Mackiewicz M, Hullbert A. Root polynomial color correction. IS&T Fifth Color Imaging Conference Proceedings. Springfield VA: IS&T 1997; p 258–261.
6. Hong G, Luo MR, Rhodes PA. A study of digital camera colorimetric characterization based on polynomial modelling. Color Res Appl 2001; 26: 76–84.
7. Cheung T, Westland S. Color camera characterization using artificial neural networks. IS&T Tenth Color Imaging Conference Proceedings. Springfield VA: IS&T 2001; p 117–120.
8. Berns RS, Shyu MJ. Colorimetric characterization of a desktop drum scanner using a spectral model. J Electron Imaging 1995;4:360–372.
9. Suzuki S, Kusunoki T, Mori M. Color characteristic design for color scanners. Appl Opt 1990;29:5187–5192.
10. Hung PC. Colorimetric calibration for scanners and media. Proc SPIE 1991;1448:164–174.
11. Kang HR. Color scanner calibration. J Imaging Sci Technol 1992;36: 162–170.
12. Funt B, Bastani P. Intensity Independent RGB-to-XYZ Color Camera Calibration. AIC Conf Proc. Taipei: Taiwan: 2012;128–131.
13. Gonzalez RC, Woods RE, Eddins SL. Digital Image processing using Matlab. Pearson Prentice Hall, New Jersey; 2004.
14. Lawson CL, Hanson RJ. Solving least squares problems: Classics in applied mathematics. Philadelphia: Soc Ind Appl Math 1995;15:63–67.
15. Moore EH. On the reciprocal of the general algebraic matrix. Bull Am Math Soc 1920;26:394–395.
16. Marimont DH, Wandell BA. Linear models of surface and illuminant spectra. J Opt Soc Am 1992;9:1905–1913.
17. Vrhel MJ, Trussel HJ. Color correction using principal components. Color Res Appl 1992;17:328–338.
18. Finlayson GD, Drew M. Constrained least-squares regression in color spaces. J Electron Imaging 1997;4:484–493.
19. Andersen CF, Hardeberg JY. Colorimetric characterization of digital cameras preserving hue planes. Proceedings of the 13th Color and Imaging Conference. Springfield VA: IS&T 2005; p 141–146.

20. Lagarias JC, Reeds JA, Wright MH, Wright PE. Convergence properties of the nelder-mead simplex method in low dimensions. *SIAM J Optimization* 1998;9:112–147.
21. Kim D, Sra S, Dhillon IS. Fast Newton-type Methods for the Least Squares Nonnegative Matrix Approximation Problem. *Proceedings of the 2007 SIAM Conference on Data Mining*, SIAM, Philadelphia, 2007. p 343–354.
22. Kiers HAL. Setting up alternating least squares and iterative majorization algorithms for solving various matrix optimization problems. *Comput Stat Data Anal* 2002;41:157–170.
23. Barnard K, Martin L, Funt B, Coath A. A dataset for color research. *Color Res Appl* 2001;27:147–151.
24. Hunt R, Pointer MR. *Measuring Colour*, 4th edition. Chichester: Wiley West Sussex, UK 2011.

APPENDIX

Below, we give details on calculation of the basis weight vector w required by the Algorithm 2 (line 2). Let \mathbf{R}_j denote the j th color channel of the reference chart image where $j=1, 2, 3$, N is the number of color chart patches and \mathbf{G}_k the k -th basis image where $k=1, \dots, K$. Then, for each of the three color channels, we calculate K images as the following ratios: i.e., $\mathbf{R}_j/\mathbf{G}_1, \mathbf{R}_j/\mathbf{G}_2, \dots, \mathbf{R}_j/\mathbf{G}_K$. Next, for the above ratio images, we calculate the average pixel values at the N locations of the color

chart patches. We place these average pixel values into the $N \times K$ matrix \mathbf{H}_j . We form the $3N \times K$ matrix \mathbf{H} by stacking matrices \mathbf{H}_j for all three color channels.

$$\mathbf{H} = \begin{bmatrix} \mathbf{H}_{\text{red}} \\ \mathbf{H}_{\text{green}} \\ \mathbf{H}_{\text{blue}} \end{bmatrix}.$$

Similarly, we form the column vector \underline{u} with $3N$ elements

$$\underline{u} = \begin{bmatrix} \underline{X} \\ \underline{Y} \\ \underline{Z} \end{bmatrix}, \quad (\text{A1})$$

where $\underline{X}, \underline{Y}$, and \underline{Z} denote the column vectors of the tristimulus value matrix \mathbf{Q} . The K -vector \underline{w} can be calculated using least squares regression.

$$\underline{w} = \mathbf{H}^+ \underline{u}. \quad (\text{A2})$$