Finding a Colour Filter to Make
a Camera Colorimetric by Optimisation

Graham D. Finlayson[(✉)] and Yuteng Zhu[(✉)]

University of East Anglia, Norwich NR4 7TJ, UK
{g.finlayson,yuteng.zhu}@uea.ac.uk

Abstract. The Luther condition states that a camera is colorimetric if its spectral sensitivities are a linear transform from the XYZ colour matching functions. Recently, a method has been proposed for finding the optimal coloured filter that when placed in front of a camera, results in effective sensitivities that satisfy the Luther condition. The advantage of this method is that it finds the best filter for all possible physical capture conditions. The disadvantage is that the statistical information of typical scenes are not taken into account.

In this paper we set forth a method for finding the optimal filter given a set of typical surfaces and lights. The problem is formulated as a bilinear least-squares estimation problem (linear both in the filter and the colour correction). This is solved using Alternating Least-Squares (ALS) technique. For a range of cameras we show that it is possible to find an optimal colour correction filter with respect to which the cameras are *almost* colorimetric.

Keywords: Digital camera · Filter design · Colorimetry

1   Introduction

Mapping the raw RGBs measured by a camera to either display coordinates (such as sRGB) or the XYZ tristimuli - a human vision system referenced colour space - is called colour correction. Colour correction is an essential procedure in the camera pipeline since cameras do not "see" the world as humans do. Fundamentally, this is because the relationship between the spectral sensitivities of a camera and human visual matching functions is not a linear mapping. Explicitly, the Luther condition is not satisfied [1].

Many different algorithms have been developed for solving the colour correction problem. The most common method is to apply a linear correction transform mapping RGBs to XYZs (or display RGBs). While linear correction generally works well, it can still fail in a large number of cases, especially for saturated colours. In order to address this issue, polynomial regression methods [2–5],

look-up-table methods [6] and artificial neural networks have been proposed [7]. However, most non-linear methods are not invariant to exposure change. In Finlayson *et al.* [8] a root-polynomial method is developed that is exposure invariant. Indeed, given two input RGBs $\underline{p}$ and $k\underline{p}$, the outputs of colour correction are $\underline{q}$ and $k\underline{q}$ if it is exposure invariant.

Another way of achieving better colour fidelity is to make more than three measurements [9]. However, when more sensors are used, the acquisition is generally more complex and suffers from problems such as reduced resolution (when sensor filter mosaic is used) or registration problem (when multiple pictures are captured). Multispectral camera systems are much more expensive than conventional cameras and are not widely deployed.

An alternate approach to increase the dimensionality of a camera system is to take two pictures of every scene with and without a coloured filter [10]. Of course this approach requires the two images to be registered (a far from easy problem to solve). Finlayson *et al.* [11] also proposed a prefiltering solution but the aim here was not to increase the dimensionality of capture. Rather a filter was found such that the device sensitivities multiplied by the filter and then linearly transformed by a 3×3 matrix were as close as possible to the XYZ colour matching functions. We call this method **spectral-based colorimetric filter design**. Surprisingly, for some cameras there exists a filter that makes them *almost* colorimetric. In [11] a filter is sought that will allow a camera to capture colorimetric data for all possible spectra. Yet, we know that the spectra we measure in the world are not arbitrary. In particular, surface reflectances are smooth and can be well approximated by lower dimensions [12].

In this paper we wish to find the filter that makes a camera as colorimetric as possible for a given set of measured lights and surfaces. Figure 1 illustrates our approach. Here we see a standard D65 illuminant lighting a colour target with known reflectances. Given these spectra and the spectral sensitivities of the camera and XYZ colour maching functions, we can calculate the camera RGB responses and XYZ triplets respectively. In our optimisation we seek to find a colour filter (red ellipse in the Figure) in combination with a 3×3 colour correction transform.

Mathematically, we will show that the simultaneous calculation of the colour filter and colour correction matrix is a bilinear optimisation problem. We show that this can be solved using Alternating Least-Squares (ALS). We regulate the optimisation to allow us to control the shape of the filter (e.g. its transmittance property).

Experiments validate our approach. For a large corpus of data we solve for the best filter for a large range of cameras. We show a filter can always be found, with which the camera system becomes much more colorimetric.

The paper is organized as follows. In Sect. 2, we present the background on image formation and linear colour correction. Section 3 discusses the formulation and calculation of a filter and a transform matrix. The colorimetric performance is evaluated in comparing with two other methods in Sect. 4. The paper concludes in Sect. 5.
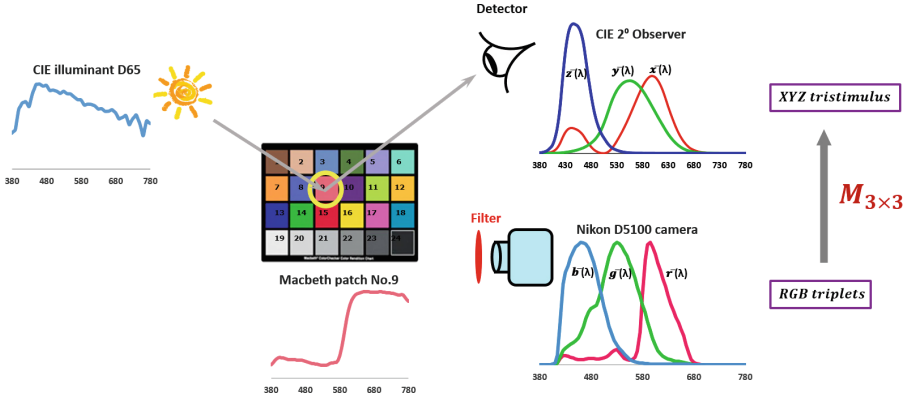
Fig. 1. Schematic diagram of colour measurement for an object viewed under a given illuminant. We try to determine a filter (placed in front) for a corresponding camera such that the RGB outputs after a linear mapping become the same as perceptual XYZ tristimulus results. Note that the human eye and camera system should be placed at the same viewing geometry in practice. (Color figure online)

2   Background

Suppose a light $E(\lambda)$ strikes a surface $S(\lambda)$ then, under the Lambertian model of image formation, the reflected light $C(\lambda)$ is proportional to $E(\lambda)S(\lambda)$. Given a set of three spectral sensitivity functions, $\underline{Q}(\lambda)$, then the sensor response is defined as:

$$\underline{\rho} = \int_{\omega} C(\lambda)\underline{Q}(\lambda)d\lambda \tag{1}$$

where the integral is taken over the visible spectrum ω. Similarly, the colour response of human visual system can be defined as

$$\underline{x} = \int_{\omega} C(\lambda)\underline{\chi}(\lambda)d\lambda \tag{2}$$

where $\chi(\lambda)$ represents the observer colour matching functions (including long-, medium- and short- wavelengths).

In practice, the spectral data is measured through sampling across the visible spectrum, i.e. typically from $400\,\mathrm{nm}$ to $700\,\mathrm{nm}$ at a $10\,\mathrm{nm}$ interval. Given a discrete representation of our data, the integrals shown above can be replaced by vector-matrix multiplication.

$$\underline{\rho} = Q^t \underline{c} \tag{3}$$

$$\underline{x} = \chi^t \underline{c} \tag{4}$$

$\underline{c}$ denotes one colour signal spectrum as a 31×1 vector. Q and χ are 31×3 matrices. The 3-vector camera response $\underline{\rho}$ and visual system response $\underline{x}$ are 3×1 vectors.

Given a $31 \times N$ matrix C of colour signal spectra (one spectrum per column) then, respectively, the camera responses and XYZ tristimuli are $N \times 3$ matrices written as

$$P = C^t Q \tag{5}$$

$$X = C^t \chi \tag{6}$$

In linear colour correction we solve for the best 3×3 matrix M that best maps camera RGBs to XYZ tristimuli. Therefore, we minimize:

$$\min_{M} \parallel PM - X \parallel \tag{7}$$

The matrix M can be solved for in closed form (using the Moore-Penrose inverse)

$$M = P^+ X = [P^t P]^{-1} P^t X \tag{8}$$

where the superscript $^+$ and ^{t} denote the pseudo-inverse and transpose operation respectively.

Finally, in the next section, we are interested in designing a filter that makes a camera more colorimetric. How then can we model the effect of a filter given the linear algebra formulation of color formation we have been developing in this section? Suppose $f(\lambda)$ denotes a transmissive filter and $C(\lambda)$ a colour signal spectrum. Physically, the light passing though a filter is equal to the product of the spectra $f(\lambda)C(\lambda)$. In the discrete domain our spectral functions are now represented by the 31-vectors $\underline{f}$ and $\underline{c}$. Unfortunately, component-wise multiplication of vectors do not exist in linear algebra. Rather we must re-express $\underline{f}$ as a diagonal matrix:

$$D(\underline{f}) = diag(\underline{f}) \begin{cases} D(\underline{f})_{ij} = 0 & \text{if } i \neq j \\ D(\underline{f})_{ij} = f_i & \text{otherwise} \end{cases} \tag{9}$$

Now, $D(\underline{f})\underline{c}$ equals the component-wise multiplication of $\underline{f}$ and $\underline{c}$.

3   Optimisation-Based Filter Design

Let us return to Fig. 1. For a given set of measured colour signal spectra and camera sensitivities, we can calculate the camera RGBs and the corresponding tristimuli. Now we wish to find a transmissive filter - that we can place in front of the camera - that will allow the RGBs to be corrected more accurately. That is, when we carry out a least-squares regression of the filtered RGBs we are closer to the ground-truth XYZs.

A high-level mathematical formulation of the optimisation - for finding the optimal filter supporting colour correction - can be addressed as:

$$\min_{\underline{f},M} \parallel C^t diag(\underline{f}) QM - C^t \chi \parallel \tag{10}$$

As before C denotes a set of N combinations of colour signal spectra. Respectively, Q and χ are the 31×3 matrices encoding the spectral sensitivities of the camera and XYZ colour matching functions. The colour filter is denoted by the 31×1 vector $\underline{f}$. Remember the function $diag()$ turns a vector into a diagonal matrix where the filter components are mapped to the diagonal of the matrix (see the end of Background Section) for a description. Finally, M denotes a colour correction matrix.

The form of Eq. 10 is *bilinear*. That is to say we are solving for $\underline{f}$ and M, and if one (or the other) is held fixed the problem becomes a simple linear optimisation. We exploit this insight to solve for the overall optimisation problem. See Algorithm 1 below for the details.

Algorithm 1. Alternative Least-squares Regression Algorithm for filter and linear matrix calculation

---

1: $i = 0, M_{3 \times 3}^0 = I, R^0 = C^t Q$
2: **repeat**
3:    $i = i + 1$
4:    $\min_{\underline{f}^i} \parallel C^t diag(\underline{f}^i) Q M^{i-1} - C^t \chi \parallel$, subject to $0 \leq \underline{f}^i \leq 1$
5:    $\min_{M^i} \parallel C^t diag(\underline{f}^i) Q M^i - C^t \chi \parallel$
6:    $R^i = C^t diag(\underline{f}^i) Q M^i$
7: **until** $\parallel R^i - R^{i-1} \parallel < \epsilon$
8: **return** $\underline{f} = \underline{f}^i$ and $M = M^i$

Most of the optimisation shown in Algorithm 1 is straightforward. Particularly in Step 5, we are solving a normal linear regression (and can use the Moore-Penrose inverse). However, solving for the filter $\underline{f}$ is more complex. It is still linear, ultimately, and can be solved using the Moore-Penrose inverse but there is some 'book-keeping' (equation rearranging) to be done.

First, let us rewrite $diag(\underline{f})$ in the following way

$$diag(\underline{f}) = f_1 \begin{pmatrix} 1 & & & \\ & 0 & & \\ & & \ddots & \\ & & & 0 \end{pmatrix} + f_2 \begin{pmatrix} 0 & & & \\ & 1 & & \\ & & \ddots & \\ & & & 0 \end{pmatrix} + ... + f_{31} \begin{pmatrix} 0 & & & \\ & 0 & & \\ & & \ddots & \\ & & & 1 \end{pmatrix}$$

$$= f_1 D_1 + f_2 D_2 + ... + f_3 D_{31} \qquad (11)$$

where matrix D_i is a sparse matrix having one non-zero value in the i^{th} diagonal. Based on this property, the calculation of $C^t diag(\underline{f}) Q M$ can be expressed as follows

$$C^t diag(\underline{f}) Q M = f_1 C^t D_1 Q M + f_2 C^t D_2 Q M + ... + f_{31} C^t D_{31} Q M \qquad (12)$$

Let us define a vector $\underline{V}_i = vec(C^t D_i Q M)$ where the $vec()$ function strips out a matrix into a vector. A new matrix $V = [V_1, V_2, ..., V_{31}]$ can now be constructed accordingly (the vector $\underline{V}_1$ is placed in the first column followed by the second $\underline{V}_2$ and the third V_3 etc.). Similarly, we vectorise on the human response and define $X = vec(C^t \chi)$. Step 4 of our algorithm can now be reformulated as:

$$\min_{\underline{f}} \parallel V\underline{f} - X \parallel \tag{13}$$

The same vectorisation is operated on the matrix X. Apparently, the filter can now be easily solved by least-squares regression as $f = V^+ X$.

Using this newly calculated filter $\underline{f}$, next we can solve the mapping matrix M as

$$M = (C^t diag(\underline{f})Q)^+ (C^t \chi) \tag{14}$$

Because solving for the filter or the colour correction matrix is to solve a least-squares problem, then the error reduces at each stage in the optimisation process. Further it is well known that Alternating Least-Squares problems (of which Algorithm 1 is a particular case) also converge [13].

It is important to note that from a physical perspective, the transmittance of the filter must be within the range [0, 100%]. Therefore Eq. 13 is solved subject to $0 \leq f \leq 1$. This linear constraint condition can be achieved using Quadratic programming (least-squares problem as in Eq. 13 can be easily converted into quadratic problem) where we apply the upper and lower constraints upon the parameters [14].

4   Experimental Results

In this work, we find the best colour pre-filter for a set of 28 digital cameras [15]. The colour signal tested here is a combination of the CIE standard illuminant D65 [16] with SFU-1995 reflectance data set [17]. The filter optimization is based on the best mapping between RGBs (after filtering and linear correction) and reference XYZs as formulated in Eq. 10.

4.1   Spectral Transmission of Filter

The filter and the corresponding transform matrix for each camera device with given testing colour signal inputs are calculated through Algorithm 1. Note that in order to simulate a physically reliable filter, we constrain its parameters in the range of [0, 100%] and for the current method, the experimental results presented here (in Table 1) are based on this constraint. For Canon D50 we show the filter found by bilinear least-squares at the top of Fig. 2a representing the transmittance within [0, 100%]. Actually, by using the Quadratic programming technique, the boundaries for filter parameters can be easily adjusted. In Fig. 2b we also show a variant where the filter transmittance is higher constrained to be between 50% and 100% (it can be regarded as a high-transparent filter which can result in less noise issues).
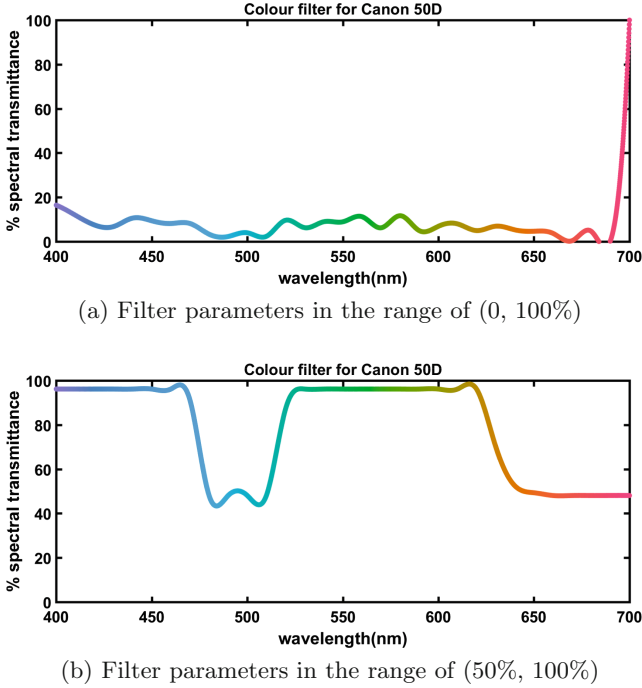
(a) Filter parameters in the range of (0, 100%)

(b) Filter parameters in the range of (50%, 100%)

Fig. 2. Filters results for Canon D50 under different constraints

4.2 Colour Evaluation

We evaluate our method compared to simple least-squares and the previous "spectral-based" colorimetric filter design method as in [11]. Our results are summarized for all 28 cameras in Table 1.

In the first three columns of the table, we record the mean, median and 95 percentile of colour difference errors in terms of CIELAB ΔE_{ab}^* for the SFU-1995 reflectance data set viewed under a CIE D65 illuminant. In the second set of three columns we record the performance of the prior filter design method (that tries to find a filter so a camera best matches the Luther Condition). Finally, in the last three columns we record the colour correction performance by our new method. The overall colour correction performance is drawn in Fig. 3 listing the results by these three methods (from left to right).

The current filtering method can achieve as small error as $0.98 \pm 0.28 \Delta E_{ab}^*$ by averaging the whole camera set. The overall medium error is even smaller, reaching $0.59 \pm 0.17 \Delta E_{ab}^*$. Clearly, our new method finds filters which support a step change in our ability to correct camera colour responses. Compared to the linear colour correction and according to the mean, median and 95% percentile error measures, the recorded error by current method is much less. Previously, we

Table 1. Comparison of colour correction results between different methods

Camera name	Linear correction			Spectral-based method			Current method		
	Mean	Median	95 pct	Mean	Median	95 pct	Mean	Median	95 pct
Canon50D	1.03	0.64	2.82	0.47	0.29	1.47	0.68	0.42	2.20
Canon60D	1.08	0.65	3.09	0.51	0.32	1.55	0.88	0.55	2.74
Canon500D	1.09	0.65	3.14	0.48	0.29	1.44	0.91	0.58	2.80
NikonD40	1.11	0.75	3.30	1.61	1.08	4.97	0.81	0.51	2.50
Sony_Nex5N	1.12	0.64	3.55	1.49	1.00	4.63	0.47	0.31	1.39
Canon600D	1.28	0.78	3.69	0.49	0.31	1.50	0.51	0.32	1.54
Canon300D	1.30	0.69	4.13	1.02	0.60	3.29	0.81	0.51	2.50
Canon1D Mark III	1.33	0.72	4.33	0.48	0.24	1.61	0.64	0.41	1.93
PentaxQ	1.36	0.90	4.06	0.85	0.45	3.03	0.61	0.39	1.82
PentaxK5	1.39	0.83	4.35	1.56	1.06	4.81	1.68	1.13	5.14
NikonD700	1.40	0.80	4.59	1.62	1.07	4.98	1.05	0.61	3.42
NokiaN900	1.41	1.01	3.86	0.60	0.35	1.90	1.01	0.58	3.30
NikonD50	1.43	0.97	4.45	1.73	1.14	5.35	0.99	0.58	3.16
NikonD3	1.45	0.84	4.63	1.60	1.06	4.92	1.09	0.65	3.51
Nikon3dx	1.46	0.83	4.76	1.65	1.11	5.06	0.95	0.56	3.07
NikonD200	1.50	0.91	4.82	1.66	1.10	5.12	1.03	0.59	3.25
NikonD90	1.54	0.87	5.05	1.66	1.11	5.11	1.06	0.63	3.41
NikonD5100	1.61	0.91	5.12	1.59	1.06	4.89	0.99	0.58	3.15
Canon40D	1.65	1.03	4.92	0.46	0.25	1.47	1.03	0.59	3.45
Canon5D Mark II	1.65	1.03	4.92	0.46	0.25	1.47	1.05	0.62	3.31
NikonD300s	1.70	0.94	5.48	1.59	1.06	4.86	1.15	0.72	3.59
Olympus_EPL2	1.77	1.15	5.56	1.33	0.85	4.15	1.22	0.70	3.83
NikonD80	1.91	1.14	6.19	1.62	1.07	4.91	1.03	0.60	3.31
Phase_One	1.95	1.18	6.24	0.74	0.43	2.37	1.00	0.55	3.32
Canon20D	2.01	1.13	6.60	0.68	0.43	2.15	0.85	0.51	2.76
PointGreyG	2.45	1.74	7.28	1.43	0.98	4.38	1.44	0.87	4.52
Hasselblad	2.65	1.65	8.45	0.89	0.57	2.78	1.61	0.94	5.13
PointGreyG2	3.03	2.07	9.38	1.60	1.08	4.90	0.98	0.57	3.06
Mean	1.60	0.98	4.96	1.14	0.74	3.54	0.98	0.59	3.11

proposed the filter design for colorimetric purpose based on Luther condition [11]. Comparing to the former spectral-based colorimetric filter design, the current method outperforms overall, especially for the Nikon cameras which illustrate a significant improvement. Among the camera set, Sony Nex5N provides the best results which are all under Just Noticeable Difference [18].
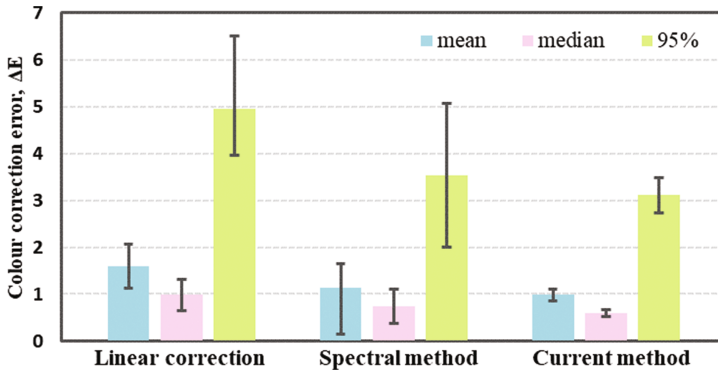
Fig. 3. Overall colour correction performance in terms of mean, median, 95 percentile colour differences with error bars.

5   Conclusion

In this article, we develop a method to find the optimal filter (to be placed in front of a camera) to make a device most colorimetric via optimisation. Experiments show that this method provides dramatic improvement over direct linear correction operating on raw unfiltered RGBs. Compared to normal linear correction the errors (calculated as mean, median or 95% ΔE_{ab}^*) are reduced by 20% to 70% on average.

References

1. Luther, R.: Aus dem Gebiet der Farbreizmetrik. Zeitschrift fur Technische Physik **8**, 540–558 (1927)
2. Hong, G., Luo, M.R., Rhodes, P.A.: A study of digital camera colorimetric characterization based on polynomial modeling. Color Res. Appl. **26**(1), 76–84 (2001)
3. Finlayson, G.D., Mohammadzadeh, D.M., Mackiewicz, M.: The alternating least squares technique for non-uniform intensity color correction. Color Res. Appl. **40**(3), 232–242 (2015)
4. Finlayson, G.D., Drew, M.S.: Constrained least-squares regression in color spaces. J. Electron. Image **6**(4), 484–493 (1997)
5. Vazquez-Corral, J., Connah, D., Bertalmio, M.: Perceptual color characterization of cameras. Sensors **14**(12), 23205–23229 (2014)
6. Hung, P.-C.: Colorimetric calibration in electronic imaging devices using a look-up-table model and interpolations. J. Electron. Image **2**(1), 53–61 (1993)
7. Li, X.: A new color correction model for based on BP neural network. Adv. Inf. Sci. Serv. Sci. **3**(5), 72–78 (2011)
8. Finlayson, G.D., Mackiewicz, M., Hurlbert, A.: Color correction using root-polynomial regression. Trans. Image Process. **24**(5), 1460–1470 (2015)

9. Liang, H.: Advances in multispectral and hyperspectral imaging for archaeology and art conservation. Appl. Phys. A **106**(2), 309–323 (2012)
10. Farrell, J., Wandell, B.: U.S. Patent No. 5479524, U.S. Patent and Trademark Office, Washington, DC (1995)
11. Finlayson, G.D., Zhu, Y., Gong, H.: Using a simple colour pre-filter to make cameras more colorimetric. In: 26th Color Imaging Conference (2018)
12. Marimont, D.H., Wandell, B.A.: Linear models of surface and illuminant spectra. J. Opt. Soc. Am. A **9**(11), 1905–1913 (1992)
13. Zhang, T., Golub, G.H.: Rank-one approximation to high order tensors. SIAM J. Matrix Anal. Appl. **23**(2), 534–550 (2001)
14. Nocedal, J., Wright, S.J.: Numerical Optimization, 2nd edn. Springer, New York (2006). https://doi.org/10.1007/978-0-387-40065-5
15. Jiang, J., Liu, D., Gu, J., Süsstrunk, S.: What is the space of spectral sensitivity functions for digital color cameras? In: IEEE Workshop on Applications of Computer Vision, pp. 168–179 (2013)
16. Ohta, N., Robertson, A.: Colorimetry: Fundamentals and Applications. Wiley, New York (2005)
17. Barnard, K., Martin, L., Funt, B., Coath, A.: A data set for color research. Color Res. Appl. **27**(3), 147–151 (2002)
18. Wyszecki, G., Stiles, W.S.: Color Science: Concepts and Methods, Quantitative Data and Formulae, 2nd edn. Wiley, New York (1982)