Bayesian color correction method for non-colorimetric digital image sensors

Xuemei Zhang**, David H. Brainard**

****Agilent Technologies Laboratories***

*****Dept. of Psychology, University of Pennsylvania***

Abstract

A Bayesian method of generating color correction matrices for digital image sensors is presented. This method was developed for sensors with poor colorimetric quality, and uses strong prior assumptions about object surface reflectance functions to improve color correction accuracy. These assumptions are expressed through linear model constraints on surface reflectance coupled with a Normal distribution over linear model weights. Results obtained with simulations and real camera images are presented. The Bayesian method works well for highly non-colorimetric sensors and has been used in industrial practice with good and stable results. For sensors with better colorimetric properties, methods employing weaker assumptions about surfaces can sometimes produce better results.

Introduction

Color correction, the process of transforming digital color pixel values captured by an image sensor into a colorimetric space (such as CIE XYZ space), is almost required in today's digital camera image processing pipeline. Color correction is especially important for highly non-colorimetric sensors, whose spectral sensitivities are poorly approximated by any linear transformation of the human color matching functions.

The general formulation for the color correction problem has been provided by many earlier papers [1, 2, 3, 4, 5] and we begin with a brief review. We assume that the spectrum arriving at the camera is formed when an illuminant reflects from a surface. We represent the object surface reflectance with a column vector R which has n elements corresponding to the object's reflectance in n sample wavelength intervals. Similarly, we represent the illuminant spectral power distribution with a column vector E, also of length n. The entries of this vector are illuminant power in each wavelength interval. For a camera with 3-sensor classes, camera spectral sensitivity is represented by an $n \times 3$ matrix S whose entries provide the sensitivity of each sensor class to light in each wavelength interval. The human color matching functions are also represented by an $n \times 3$ matrix M, whose entries provide the values of each color matching function. In this paper we will use the CIE XYZ color matching functions, but our results may be generalized to any choice of the matrix M. For an object R illuminated with E the camera responses C may be represented by a column vector of length 3 obtained as:

$$C = S^t D(E) R, \tag{1}$$

where the function $D(E)$ is a diagonal matrix with the illuminant spectral power distribution along the diagonal. To arrive at Equation 1 we assume a camera with sensors that respond linearly to light. This assumption often holds well for real cameras, and when it fails a simple correction may usually be applied to obtain linearized camera responses [6].

For the same object, the colorimetric tristimulus values XYZ, represented as a 3-vector X, are obtained as:

$$X = M^t D(E) R. \tag{2}$$

The goal of color correction is to recover the colorimetric values X from the camera responses C. Note that if the camera sensitivities S are a linear transformation of the human color matching functions M, then color correction is trivial [1]. In this special case, we can write:

$$S = MT, \tag{3}$$

where T is a 3×3 matrix. The colorimetric XYZ values X are then easily related to the sensor response values C:

$$\begin{aligned} C &= (MT)^t D(E) R \\ &= T^t (M^t D(E) R) \\ &= T^t X, \\ X &= (T^t)^{-1} C. \end{aligned} \tag{4}$$

We refer to camera sensors that satisfy Equation 3 as colorimetric sensors. In practice, camera sensitivities are never exactly colorimetric. When the deviations are small, the transformation matrix T that provides the best least-squares solution to Equation 3 may be determined and used as in Equation 4. This method of color correction has been referred to as the Maximum ignorance method, because it makes no explicit assumptions about the scene surfaces and illuminants [4]. As the sensors deviate from colorimetric, however, this approach may not produce good results (see Results section). Thus it is desirable to develop color correction methods that work better for non-colorimetric sensors. Such methods typically operate by using prior information about what surfaces and illuminants actually occur in the scenes being imaged. Here we will focus on linear color correction methods that have the same (or almost the same) form as 4 but where different criteria are used to choose the correction matrix. We also focus on the special case where the illuminant spectrum E is known up to a single scale factor k that describes the overall intensity of the illumination.[1]

We follow earlier work by using linear color-correction and assuming that surfaces are described by low dimensional linear models [2, 7, 3]. We extend the earlier work by using linear model ideas to develop an explicit prior probability distribution over the space of surface reflectance functions that might appear in the scene. Using a prior distribution of this form allows us to employ Bayes rule to guide the choice of correction matrix.

By treating color correction under the assumption that the relative spectrum of the illuminant is known, we separate the problem of color correction from the problem of color balancing. Color balancing refers to correcting the colors in an image for variation in the illumination under which the image was aquired. At the heart of the color balancing problem is the task of using the image data to estimate, either explicitly or implicitly, the spectrum of the illuminant. Algorithms for illuminant estimation and color balancing are available elsewhere (e.g. [8, 9, 10, 11]). The Bayesian color correction method we present here is not by itself a color balancing method.

Bayesian Color Correction

A Bayesian color correction method uses prior knowledge of the distribution of object surface colors in combination with the camera data to make a best estimate of the actual surface reflectances in the scene. It then uses these estimates for color correction by synthesizing the colorimetric tristimulus coordinates corresponding to these surfaces. The mathematical formulation used here is similar to that used by Brainard [12] for the demosaicing problem. For application to color correction, this formulation is coupled with a prior over surface reflectances developed by Brainard and Freeman [10] in the context of the illumination estimation problem. This prior assumes a linear model for surface reflectance functions [13, 14, 15, 7] and adds to this the assumption that a multivariate Gaussian distribution describes the distribution of weights on the basis functions of the linear model.

The linear model assumption for surface reflectances is that any surface reflectance function R in the scene can be expressed as a linear combination of a set of m surface basis functions B:

$$R = BW, \tag{5}$$

where B is an n-by-m matrix, with each column representing one surface basis function and W is a vector of length m representing the weights for each basis function. Such surface basis functions can be obtained by applying principle components analysis (PCA) on ensembles of reflectance spectra that are representative of the surfaces likely to be imaged [13, 14, 15, 7].

Suppose that the surfaces in the scene are well approximated by a linear model of dimension 3. Then we can derive

$$\begin{aligned} C &= S^t D(E) B W \\ W &= (S^t D(E) B)^{-1} C \\ X &= (M^t D(E) B)(S^t D(E) B)^{-1} C. \end{aligned} \tag{6}$$

Equation 6 provides a method of color correction that will work for arbitrary 3-color camera sensors as long as the 3 dimensional linear model constraint holds exactly.

We evaluated the quality of surface linear models using measurements of surface reflectances made by Vrhel [16]. For each surface S_i out of the 170 Vrhel surface measurements, we used the singular value decomposition on the other 169 surfaces to obtain a set of m basis functions with the m largest singular values. A linear combination of these basis functions were then found that best fit (least mean squared error) the surface reflectance S_i. Then ΔE_{94} error between the actual S_i and the fitted surface reflectance, both assumed to be illuminated by D65, was calculated. Figure 1 shows the average and maximum ΔE_{94} values over the 170 Vrhel surfaces evaluated in this way. From these results we can say with some confidence that the assumption in Equation 5 is reasonable as the number of basis functions m is chosen to be in the range 5-8, consistent with conclusions drawn by previous authors [13, 14, 15, 7]. When the number of basis functions decreases below 5, the quality of the approximation declines considerably. It is thus of interest to consider whether it is possible to develop methods that can incorporate information from linear models with more than 3 basis functions.

[1] The factor k describes uncertainty in the illumination intensity at the image plane, thus it also includes the uncertainty in overall sensitivity of the camera sensors that arises with variation in f-stop and exposure duration. We do not explicitly represent the factor k in the formulation in this paper, as it can be shown that its value has no effect on the color correction matrices we derive.

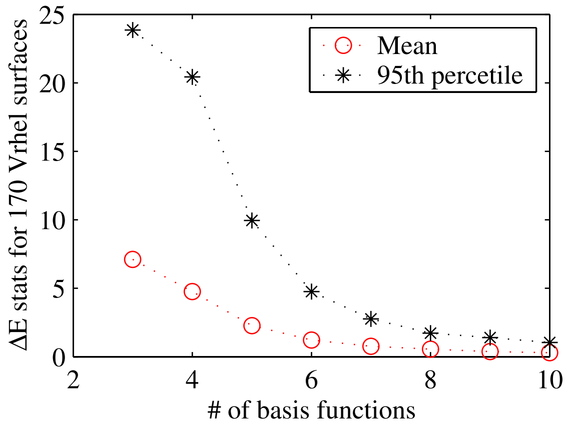

Figure 1: *Color error when using various number of linear basis functions to approximate surface reflectance functions from the Vrhel set [16]. The ΔE_{94} values were calculated as described in the text, assuming a D65 illuminant.*

The approach we take to this is to employ Bayesian methods. A Bayesian method always involves specification of a prior distribution for the parameters to be estimated, and a likelihood function for the data. The prior distribution expresses in a probabilistic fashion what is known about the scene a priori, and in our hands is used to generalize the deterministic linear model constraint expressed by Equation 5. The likelihood function expresses the image formation model used in the analysis, and is made probabilistic by incorporation of sensor noise. It is essentially a generalization of Equation 1. Given the prior and the likelihood, the Bayesian approach proceeds by computing a posterior distribution as the product of the prior and the likelihood [17, 18]. Here the posterior provides an estimate of how likely any surface reflectance is, given the camera sensor responses. From the posterior, a specific estimate of surface reflectance can be made. Here we choose this estimate as the mean of the posterior.

To express a prior over surfaces, we begin with an m dimensional linear model for surfaces. Within this model, the free parameters are the weights W. We use a multivariate Normal distribution to describe the probability that any given set of weights occurs in a surface in the scene [10]. This prior distribution is represented by a mean vector μ_w (one number for each weight) and a covariance matrix K_w:

$$P(W) \sim N(\mu_w, K_w). \tag{7}$$

Figure 2 shows the empirical distribution of weights for each basis function in a 6 dimensional linear model for the Vrhel[16] surface reflectance functions. The figure also shows the marginal distributions of a Gaussian fit

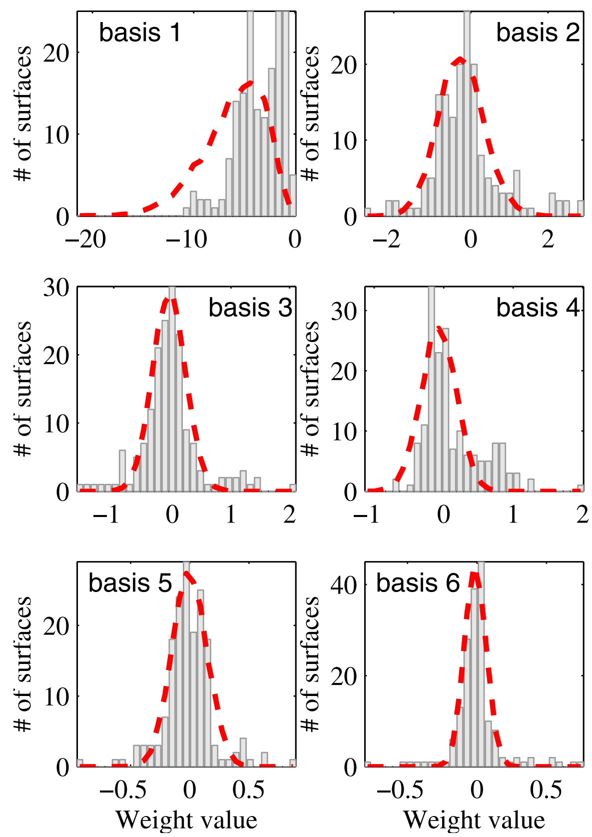

Figure 2: *Histogram of weights for the first 6 surfaces basis functions for the Vrhel set [16] (bars), with the marginals of the fitted Gaussian distribution (dashed lines).*

to the empirical distribution. The Gaussian fit was chosen such that weights drawn from it matched the empirical distribution after removal of any weights that corresponded to reflectance functions with negative values. Note that this procedure differs from simply using the sample mean and covariance matrix of the empirical distribution as the parameters of the Gaussian. The weight distributions for basis functions 2-6 are fairly well captured by the Gaussian approximation, although the empirical distributions tend to have broader tails than the Gaussian fits. The weight distribution for the first basis function is not-well described by the best-fit Gaussian. This is because our fitting procedure incorporates in a heuristic manner the fact that real surface reflectance functions are non-negative.

To express a likelihood, we start with Equation 6 and write the camera responses C as a function of the linear model weights W:

$$C = T_c W \tag{8}$$

where $T_c = S^t D(E) B$.

Camera responses are noisy. Here we assume independent additive Gaussian noise ε with a mean vector of 0 and a diagonal covariance matrix K_e:

$$\varepsilon \sim N(0, K_e). \tag{9}$$

Therefore the likelihood model for the observed camera data $C = (T_c W + \varepsilon)$, is:

$$P(C|W) \sim N(T_c W, K_e). \tag{10}$$

The prior distribution and camera data likelihood function lead directly to the posterior distribution of the surface basis weights given the camera data through Bayes rule:

$$\begin{aligned} P(W|C) &= P(C|W) * P(W)/P(C) \\ &\sim N(T_c W, K_e) N(\mu_w, K_w). \end{aligned} \tag{11}$$

The posterior distribution in the form above is Normal with mean vector μ_1 and covariance matrix K_1 [18, 12]:

$$\mu_1 = IC + i_0, \tag{12}$$
$$K_1 = (K_w^{-1} + T_c^t K_e^{-1} T_c)^{-1}, \tag{13}$$
$$I = K_w T_c (T_c K_w T_c^t + K_e)^{-1}, \tag{14}$$
$$i_0 = (\mu_w - I T_c \mu_w). \tag{15}$$

The minimum mean square error estimator for W is the posterior mean $\hat{W}$ (which is the same as the MAP estimator in the case of a Normal posterior):

$$\hat{W} = \mu_1 = IC + (\mu_w - I T_c \mu_w). \tag{16}$$

The estimated surface weights $\hat{W}$ allow us to color correct the camera responses by synthesizing the XYZ values :

$$\begin{aligned} \hat{X} &= (M^t D(E) B)\hat{W} \\ &= T_x \hat{W} \\ &= T_x (IC + i_0) \\ &= (T_x I)C + (T_x i_0). \end{aligned} \tag{17}$$

The matrix $T_x = (M^t D(E) B)$ is 3-by-m. For a 3-color sensor, the matrix I is m-by-3, and i_0 is m-by-1. Thus the Bayesian color correction process includes a 3-by-3 transformation $(T_x I)$ of the camera RGB data, plus a 1-by-3 offset term $(T_x i_0)$. For most sensors, the offset terms are very small, and thus negligible when compared to the matrix terms. We include the offset term in the evaluations in the Results section. In practice, the offset term can almost always be ignored with negligible effect on the results.

Results and Discussion

When the surfaces in the scene are drawn according to the prior distribution and when the sensor noise matches that used for the likelihood function, the Bayesian method described above is optimal in the sense that it minimizes the expected mean squared error in the estimation of the surface weights. For the application to color correction, several factors can prevent the method from performing optimally. First, it is clear that the distribution of surface reflectance weights is only approximated by a Gaussian. Second, sensor noise may not be additive and Gaussian. Finally, our interest in practice is not in minimizing the estimation error for the surface weights, but rather in minimizing the perceptual error of the estimated tristimulus values. Here we evaluate this perceptual error using the ΔE_{94} [19] metric. Since this metric is a non-linear transformation of the difference between estimated and desired tristimulus values, using the posterior mean is not guaranteed to minimize its expected value. It is therefore of interest to evaluate the efficacy of the Bayesian algorithm for conditions of interest, and to compare its performance to that of other available color correction algorithms.

Evaluation by simulation: Method

We begin with simulations. We used the Vrhel surface reflectance measurements [16] and spectral sensitivity functions of a number of image sensors with different colorimetric qualities, as measured by their different Vora values [20]. For each sensor, 3-color responses C to the Vrhel object surfaces, illuminated by a chosen illuminant, were calculated according to Equation 1. No noise was added to the simulated sensor responses.

To calculate Bayesian color correction results, we pick one surface as the target at a time, and use the remaining 169 surfaces from the Vrhel set to calculate surface basis functions and Gaussian prior over the surface weights. A color correction matrix and an offset term were calculated as in Equation 17, and the camera responses C were corrected to XYZ values accordingly for this surface. The noise level was assumed to be 0. The "true" XYZ values for this surface were calculated according to Equation 2. CIE ΔE_{94} color difference values were then calculated between the XYZ values estimated from the camera responses and the "true" XYZ values. This process was iterated, until we had color correction errors for all 170 surfaces. We also looked at the effect of estimating more than 3 weights by varying the number of surface basis functions (3 to 10) in the Bayesian color correction process.

In addition to the Bayesian color correction method, we also calculated color reproduction errors for several other methods. In order of increasingly strong assumptions made about object surfaces, the methods evaluated were:

1. Maximum ignorance method (simple pseudo-inverse method which finds matrix to minimize mean squared error on the camera spectral sensitivity multiplied with the scene illuminant).
2. Maximum ignorance with positivity assumption [4].
3. Minimal knowledge [5] with α values varying from 30 to 700.
4. Basic linear model correction assuming 3 surface basis functions (Equation 6).

For the various methods, the estimated XYZ coordinates occasionally contained small negative values. Such values were truncated to 0 before computing the ΔE_{94} error.

Evaluation by simulation: Single illuminant

Figure 3 shows results from two CMOS imager sensors with different colorimetric properties, as measured by their Vora values of 0.9 and 0.76. The results for the Basic linear model correction with 3 basis functions (Equation 6) were not measurably different from the Bayesian method implemented with 3 basis functions and are not shown explicitly in the figure.

For the sensor with better colorimetric quality (top plot), color correction results from most methods produce acceptable mean ΔE_{94}. The basic linear method with 3 basis functions and the Bayesian method with 4 bases worked best, but the Maximum ignorance with positivity constraint method and the Minimal knowledge method (with $\alpha = 30$) also worked very well. The stronger surface prior used in the Bayesian method did not make it perform better here. This is perhaps to be expected for this sensor, as when the sensor is close to colorimetric, strong prior knowledge about assumed for object surfaces should not be necessary for good performance.

An interesting features of Figure 3 is that the mean ΔE_{94} errors do not decrease as the number of basis functions used in the Bayesian method increases. This is not surprising, as the Bayesian method minimizes mean squared error on the estimated weights of surface basis functions for each surface color – which indeed decreases as number of basis functions increases in this case, but such a decrease does not always translate to a decrease in ΔE_{94} values. In general we find that the optimal number of basis functions to use with the Bayesian method depends on the sensor spectral sensitivity and the scene illuminant.

The bottom plot of Figure 3 shows results for a sensor with poor colorimetric quality (Vora value = 0.76). Here, the stronger assumptions about scene surfaces embodied by the Bayesian method provide a larger benefit. With

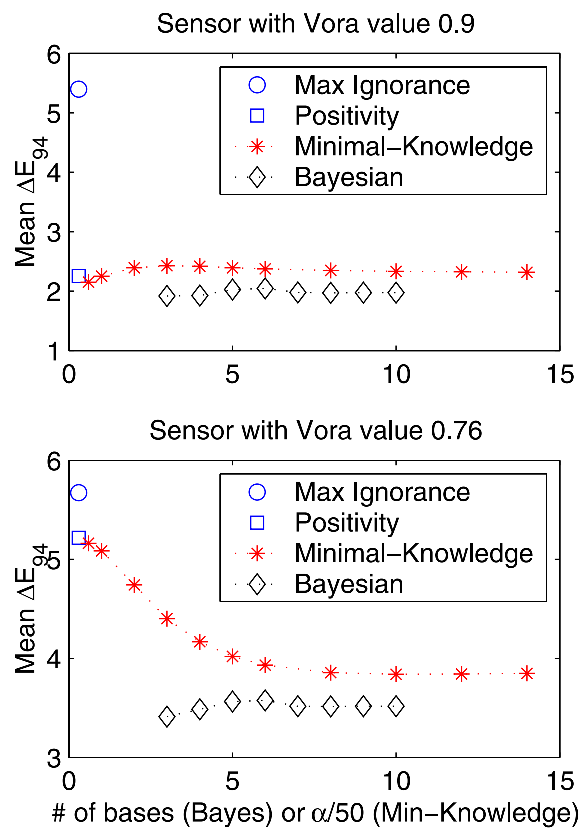

Figure 3: *Color correction results for two sensors with different colorimetric qualities, with D65 as the illuminant.*

the Bayesian color correction, mean ΔE_{94} values are lower than for any of the other methods. Methods embodying weak assumptions generally performed worse than methods embodying stronger assumptions for this sensor.

Evaluation by simulation: Many sensors and illuminants

In the previous section we evaluated color correction method for a single choice of scene illuminant and two sensors. We performed the same simulation for a large number of color sensors and different choices of scene illuminant. A total of 257 sensor spectra were generated by combining spectral sensitivity measurements from 27 imagers with 11 infra-red filters (4 of the 27 imagers have built-in IR filters, thus the total number of combinations is less than 27×11). Four different illuminants (D65, cool white fluorescent, 2277K and 2848K blackbody) were used, giving a total of 1028 sensor/illuminant combinations. The Vora values, which meausure the colorimetric quality of each sensor, were calculated and the sensors were classified into 10 Vora-value bins. For the Minimal knowledge method, the alpha value was chosen separately for each sensor/illuminant combination to minimize mean error for that combination. The number of basis functions used for the Bayesian method was set between 4 and 10 for each sensor/illuminant combination, with the setting chosen to minimize mean error. As noted previously, linear model and Bayesian correction with three bases perform similarly in the noise free case of our simulations.

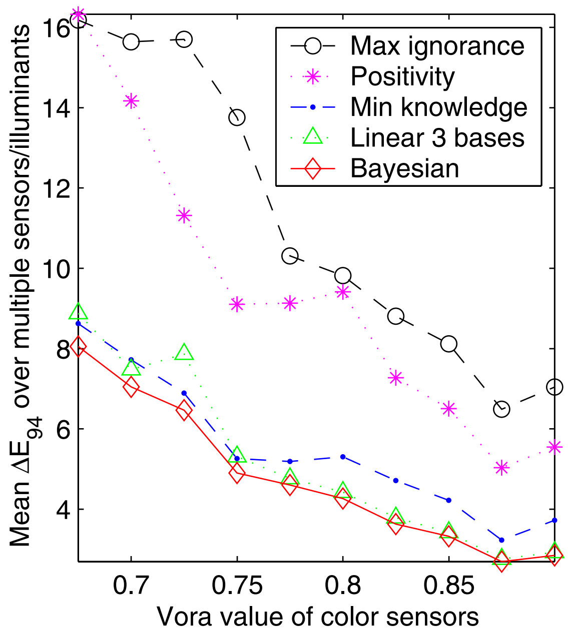

Figure 4: *Color correction results for different sensor/IR-filter/illuminant combinations as a function of the sensors' Vora values[20], which measure their colorimetric quality.*

Figure 4 shows the mean ΔE_{94} values for the 170 Vrhel surfaces, averaged over multiple sensor/illuminant combinations, as a function of the sensors' Vora values. The Bayesian model performs best on average. More generally, we see that methods which make stronger prior assumptions (Minimal knowledge, linear model with 3 bases, and Bayesian) perform better, and that the difference between these methods is small. The improvement of methods with strong assumptions versus methods with weak assumptions is largest when the colorimetric quality of the sensors is poor. The standard deviations of the error values are generally about 30-50so there is considerable variability in terms of which method is most optimal for different sensor/illuminant combinations.

Evaluation using data captured from a real camera

To check the accuracy of the color balancing process with real camera data, an image of a Macbeth color checker was taken using a Kodak DCS-460 digital camera. Scene illuminant was measured using a PhotoResearch PR650 spectroradiometer. For each color patch of the color checker, a 12 by 10 pixel area was selected from the image, and raw RGB values (before demosaicing) were extracted. For the Bayer mosaic pattern used in the Kodak DCS camera, this corresponded to 60 green values, 30 red values, and 30 blue values per color patch. These raw values were first linearized, using the built-in gamma table of the DCS camera, to yield linear raw RGB values ranging from 0 to 5445. From these values we calculated the mean camera responses (red, green, and blue) for each color patch. They were used to estimate XYZ values of the color patches and then compare to the "true" XYZ values of them calculated from measured surface reflectance functions of the macbeth color checkers and the measured scene illuminant.

Table 1: *Color correction results for 24 Gretag Macbeth color checker patches, from an image of the color checker taken with a Kodak DCS-460 camera under outdoor daylight.*

Method	Mean ΔE_{94}	95th percentile ΔE_{94}
Max ignorance	5.47	15.5
Positivity	3.48	8.8
Min knowledge)	3.26	7.6
Linear 3 bases	3.61	8.4
Bayesian	3.14	7.2

We again calculated color correction matrices using 5 methods: Maximum ignorance, Maximum ignorance with positivity constraint [4], Minimal knowledge [5], Basic 3-bases linear model [2, 7, 3], and Bayesian. For the Minimal knowledge method, we calculated the color matrix using a number of different α values and found that $\alpha = 14$ gave the best performance (lowest mean ΔE_{94} over the 24 Macbeth surfaces) for this sensor and this set of surfaces, thus only the $\alpha = 14$ results will be shown. For the Bayesian method, we used 6 basis functions and chose a noise standard deviation used to compute the estimator to minimize the error. This standard deviation corresponded to 8.6 digital values out of the linear camera response range 0-5445. The measured scene illuminant, an outdoor daylight illuminant with a correlated color temperature of 5350K, was used as the illuminant in all color matrix calculations.

Once the color correction matrices were calculated, the XYZ values were estimated from the camera mean RGB values for each of the Macbeth surface. Here the free overall scale factor was determined so that the average luminance of all 24 color checker patches matched between the color-corrected values and the "true" values. After the scales were matched, ΔE_{94} values were calculated between the color-corrected XYZ values and "true" XYZ values, using the XYZ value of the measured scene illuminant as the white point. Table 1 gives the mean and 95th percentile ΔE_{94} errors over the 24 Macbeth surfaces for each method. The Maximum ignorance method had the largest color errors. Adding the positivity constraint [4] improved color accuracy considerably. The Minimal knowledge method [5] with an optimal α improved accuracy further. The Basic linear model method gave performance similar to that of the Maximum ignorance with positivity constraint method. The Bayesian method gave the best results. Overall, the results with a real captured image of known surfaces confirmed the evaluation provided by the simulations.

Our results indicate that given our current state of understanding, choosing a preferred color correction method requires a fair amount of engineering judgment, as trade-offs must be made about robustness with respect to sensor spectral sensitivities and the range of scene conditions that will be encountered. In general, we have found it useful to test all the above methods for a new sensor, and pick the method which gives the smallest error in simulations for the most likely illuminants. For highly non-colorimetric sensors, we have obtained good results in practice using the Bayesian method presented here. For both the minimal knowledge and Bayesian methods, it is important to optimize the algorithm's parameters for the sensor spectral sensitivities, sensor noise level, and scene illuminants likely to be encountered. As an implementation detail, note that color correction matrices for use in real cameras may be precomputed, so that the complexity of the computations required to generate such matrices is not problematic.

The work reported here suggests interesting future directions. These include developing and using more accurate priors for naturally occurring surfaces, extending the method to incorporate a perceptual error metric such as ΔE_{94}, and integrating color correction with illuminant estimation to improve robustness across different scene illuminations.

Summary

We presented a Bayesian method to calculation color correction matrices for non-colorimetric sensors, and compared its performance to several well-known color correction methods. These methods make different levels of assumptions about the statistics of the object surfaces to be measured. The Bayesian method, which makes strong assumptions about object surfaces, performed very well, particularly for highly non-colorimetric sensors.

Acknowledgements

We wish to thank Yvonne Lee, Mark Butterworth, Dwight Poplin, Rick Baer, and Dieter Vook for providing the sensor spectral measurements, and for helpful discussions.

References

1. B. K. P. Horn, Exact reproduction of colored images, Computer Vision, Graphics, and Image Processing, 26, 135 (1981).
2. Brian A. Wandell, The synthesis and analysis of color images, IEEE Transactions on Pattern Analysis and Machine Intelligence, 9, 2 (1987).
3. David Hoyt Brainard, Colorimetry, McGraw-Hill, New York, pp. 26.21–26.54 (1995).
4. Graham D. Finlayson and Mark S. Drew, The maximum ignorance assumption with positivity, Proceedings of the IS&T/SID 4th Color Imaging Conference, pp. 202–204 (1996).
5. J. A. Stephen Viggiano, Minimal-knowledge assumptions in digital still camera characterization i.: Uniform distribution, toeplitz correlation, Proceedings of the IS&T/SID 9th Color Imaging Conference, pp. 332–336 (2001).
6. Poorvi L. Vora, Joyce E. Farrell, Jerome D. Tietz and David H. Brainard, Image capture: modelling and calibration of sensor responses and their synthesis from multispectral images, IEEE Transaction on Image Processing, 10(2), 307 (2001).
7. M. J. Vrhel and H. J. Trussell, Color correction using principal components, Color Research & Application, 17(5), 328 (1992).
8. Graham D. Finlayson, Paul M. Hubel and Steven Hordley, Color by correlation, Proceedings of the IS&T/SID 5th Color Imaging Conference, pp. 6–11 (1997).
9. Kobus Barnard, Graham Finlayson and Brian Funt, Colour constancy for scenes with varying illumination, Computer Vision and Image Understanding, 65(2), 311 (1997).
10. David H. Brainard and W. T. Freeman, Bayesian color constancy, Journal of the Optical Society of America A, 14, 1393 (1997).
11. Shoji Tominaga, Satoru Ebisui and Brian A. Wandell, Color temperature estimation of scene illumination, Proceedings of the IS&T/SID 7th Color Imaging Conference, pp. 42–47 (1999).
12. David H. Brainard, Bayesian method for reconstructing color images from trichromatic samples, Proceedings of the IS&T 47th Annual Meeting, pp. 375–380 (1994).
13. J. Cohen, Dependency of the spectral reflectance curves of the munsell color chips, Psychon. Sci., 1, 369 (1964).
14. L. T. Maloney, Evaluation of linear models of surface spectral reflectance with small numbers of parameters, Journal of the Optical Society of America A, 3(10), 1673 (1986).
15. T. Jaaskelainen, J. Parkkinen and S. Toyooka, A vector-subspace model for color representation, Journal of the Optical Society of America A, 7, 725 (1990).
16. M. J. Vrhel, R. Gershon and L. S. Iwan, Measurement and analysis of object reflectance spectra, Color Research & Application, 19(1), 4 (1994).
17. T. O. Berger, Statistical Decision Theory and Bayesian Analysis, Springer-Verlag, New York (1985).
18. P. M. Lee, Bayesian Statistics, Oxford University Press, London (1989).
19. CIE, Industrial colour-difference evaluation, Bureau Central de la CIE, Vienna, Publication CIE 116-95 (1995).
20. Poorvi L. Vora and H. Joel Trussell, Measure of goodness of a set of color-scanning filters, Journal of the Optical Society of America A, 10(7), 1499 (1993).