

Improvements to Gamut Mapping Colour Constancy Algorithms

Kobus Barnard

Computer Science Division, U. C. Berkeley, Berkeley, CA, 94720
email: kobus@cs.berkeley.edu

Abstract. In his paper we introduce two improvements to the three-dimensional gamut mapping approach to computational colour constancy. This approach consist of two separate parts. First the possible solutions are constrained. This part is dependent on the diagonal model of illumination change, which in turn, is a function of the camera sensors. In this work we propose a robust method for relaxing this reliance on the diagonal model. The second part of the gamut mapping paradigm is to choose a solution from the feasible set. Currently there are two general approaches for doing so. We propose a hybrid method which embodies the benefits of both, and generally performs better than either. We provide results using both generated data and a carefully calibrated set of 321 images. In the case of the modification for diagonal model failure, we provide synthetic results using two cameras with a distinctly different degree of support for the diagonal model. Here we verify that the new method does indeed reduce error due to the diagonal model. We also verify that the new method for choosing the solution offers significant improvement, both in the case of synthetic data and with real images.

1 Introduction

The image recorded by a camera depends on three factors: The physical content of the scene, the illumination incident on the scene, and the characteristics of the camera. This leads to a problem for many applications where the main interest is in the physical content of the scene. Consider, for example, a computer vision application which identifies objects by colour. If the colours of the objects in a database are specified for tungsten illumination (reddish), then object recognition can fail when the system is used under the very blue illumination of blue sky. This is because the change in the illumination affects object colours far beyond the tolerance required for reasonable object recognition. Thus the illumination must be controlled, determined, or otherwise taken into account.

Compensating for the unknown illuminant in a computer vision context is the computational colour constancy problem. Previous work has suggested that some of the most promising methods for solving this problem are the three dimensional

gamut constraint algorithms [1-4]. In this paper we propose two methods for further improving their efficacy.

The gamut mapping algorithms consist of two stages. First, the set of possible solutions is constrained. Then a solution is chosen from the resulting feasible set. We propose improvements to each of these two stages. To improve the construction of the solution set, we introduce a method to reduce the error arising from diagonal model failure. This method is thus a robust alternative to the sensor sharpening paradigm [2, 4-6]. For example, unlike sensor sharpening, this method is applicable to the extreme diagonal model failures inherent with fluorescent surfaces. (Such surfaces are considered in the context of computational colour constancy in [7]).

To improve solution selection, we begin with an analysis of the two current approaches, namely averaging and the maximum volume heuristic. These methods are both attractive; the one which is preferred depends on the error measure, the number of surfaces, as well as other factors. Thus we propose a hybrid method which is easily adjustable to be more like the one method or the other. Importantly, the combined method usually gives a better solution than either of the two basic methods. Furthermore, we found that it was relatively easy to find a degree of hybridization which improves gamut mapping colour constancy in the circumstances of most interest. We will now describe the two modifications in more detail, beginning with the method to reduce the reliance on the diagonal model.

2 Diminishing Diagonal Model Error

We will begin with a brief review of Forsyth's gamut mapping method [1]. First we form the set of all possible RGB due to surfaces in the world under a known, "canonical" illuminant. This set is convex and is represented by its convex hull. The set of all possible RGB under the unknown illuminant is similarly represented by its convex hull. Under the diagonal assumption of illumination change, these two hulls are a unique diagonal mapping (a simple 3D stretch) of each other.

Figure 1 illustrates the situation using triangles to represent the gamuts. In the full RGB version of the algorithm, the gamuts are actually three dimensional polytopes. The upper thicker triangle represents the unknown gamut of the possible sensor responses under the unknown illuminant, and the lower thicker triangle represents the known gamut of sensor responses under the canonical illuminant. We seek the mapping between the sets, but since the one set is not known, we estimate it by the observed sensor responses, which form a subset, illustrated by the thinner triangle. Because the observed set is normally a proper subset, the mapping to the canonical is not unique, and Forsyth provides a method for effectively computing the set of possible diagonal maps. (See [1, 2, 4, 8-10] for more details on gamut mapping algorithms).




Fig. 1: Illustration of fundamentals of gamut mapping colour constancy.

We now consider the case where the diagonal model is less appropriate. Here it may be possible that an observed set of illuminants does not map into the canonical set with a single diagonal transform. This corresponds to an empty solution set. In earlier work we forced a solution by assuming that such null intersections were due to measurement error, and various error estimates were increased until a solution was found. However, this method does not give very good results in the case of extreme diagonal failures, such as those due to fluorescent surfaces. The problem is that the constraints were relaxed indiscriminately, rather than in concordance with the world. Similarly, even if the solution set is not null, if there is extreme diagonal failure, then the solution set may not be as appropriate for selecting the best solution (by averaging, say), than a set which is larger, but more faithful to the true possibilities.

To deal with diagonal model failure, we propose the following modification: Consider the gamut of possible RGB under a single test illuminant. Call this the test illuminant gamut. Now consider the diagonal map which takes the RGB for white under the test illuminant to the RGB for white under the canonical illuminant. If we apply that diagonal map to our test illuminant gamut, then we will get a convex set similar to the canonical gamut, the degree of difference reflecting the failure of the diagonal model. If we extend the canonical gamut to include this mapping of the test set, then there will always be a diagonal mapping from the observed RGB of scenes under the test illuminant to the canonical gamut. We repeat this procedure over a

The gamuts of all possible RGB under three training illuminants.

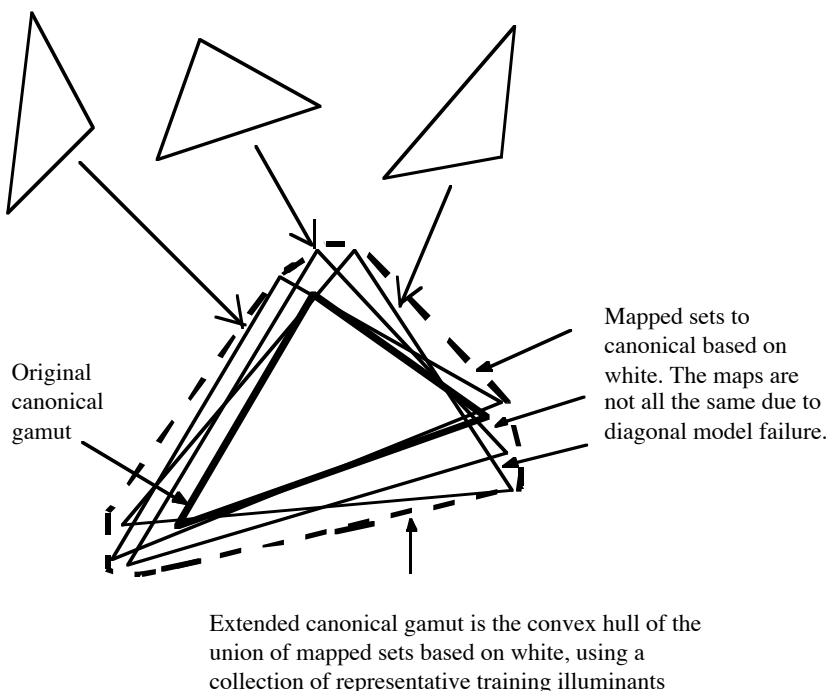


Fig. 2: Illustration of the modification to the gamut mapping method to reduce diagonal model failure.

representative set of illuminants to produce a canonical gamut which is applicable to those illuminants as well as any convex combination of them. The basic idea is illustrated in Figure 2.

3 Improving Solution Choice

Once a constraint set has been found, the second stage of the gamut mapping method is to select an appropriate solution from the constraint set. Two general methods have been used to do this. First, following Forsyth [1], we can select the mapping which maximizes the volume of the mapped set. Second, as proposed by Barnard [2], we can use the average of the possible maps. When Finlayson's illumination constraint is used, then the set of possible maps is non-convex. In [2], averaging was simplified by using the convex hull of the illuminant constraint. In [10] Monte Carlo integration was used in conjunction with the two dimensional version of the algorithm, and in [4, chapter 4] the average was estimated by numerical integration.

In [4, chapter 4], we found that both averaging and the maximum volume method have appeal. We found that the preferred method was largely a function of the error measure, with other factors such as the diversity of scene surfaces also playing a role. When the scene RGB mapping error measure is used, the average of the possible maps is a very good choice. In fact, if we are otherwise completely ignorant about the map, then it is the best choice in terms of least squares.

On the other hand, if we use an illumination estimation measure, then the original maximum volume heuristic is often the best choice. This is important because we are frequently most interested in correcting for the mismatch between the chromaticity of the unknown illuminant and the canonical illuminant. In this case, the errors based on the chromaticity of the estimated scene illuminant correlate best with our goal, and the maximum volume heuristic tends to give the best results.

In this work we will focus on estimating the chromaticity of the illuminant. Despite the success of the maximum volume heuristic, we intuitively feel that, at least in some circumstances, some form of averaging should give a more robust estimate. This intuition is strengthened by the observation that when we go from synthetic to real data, the maximum volume method loses ground to averaging (see, for example, [4, Figure 4.12]).

We begin our analysis by considering solution selection by averaging. Here we will assume that we are ignorant of the prior likelihood of the various possibilities, and thus averaging the possibilities corresponds to integrating corresponding volumes. The algorithms published so far integrate in the space of diagonal maps, which is not quite the same as the space of illuminants. Under the diagonal model, the illuminant RGB is proportional to the element-wise reciprocal of the diagonal maps. Thus we see that for an illumination oriented error measure, we may be averaging in the wrong space, as intuitively, we want to average possible illuminants.

However, averaging the possible illuminants has some difficulties. As we go towards the origin in the space of possible diagonal maps, the corresponding proposed illuminant becomes infinitely bright. The origin is included in the constraint set because we assume that surfaces can be arbitrarily dark. Although it is rare for a physical surface to have a reflectivity of less than 3%, surfaces can behave as though they are arbitrarily dark due to shading. Thus we always maintain the possibility that the illuminant is very bright. Specifically, if (R,G,B) is a possible illuminant colour, then (kR,kG,kB) is also a possible illuminant for all $k>1$. Put differently, a priori the set of RGB all possible illuminants is considered to be a cone in *illuminant* RGB space [11]. When we add the surface constraints, then the cone becomes truncated. As soon as we see anything but black, we know that the origin is excluded, and specific observed sensor responses lead to specific slices being taken out of the cone.

The above discussion underscores the idea that when we average illuminants, we should ignore magnitude. However, since the work presented in [4] demonstrates that the three dimensional algorithms outperform their chromaticity counterparts, we do not want to completely throw away the brightness information. Considering the truncated cone again, we posit that the nature of the truncations matter. The problem is then how to average the possible illuminants.

Consider, for a moment, the success of the three-dimensional gamut mapping algorithms. In the space of maps, each direction corresponds to a illuminant

chromaticity. Loosely speaking, the chromaticity implied by an RGB solution, chosen in some manner, is the average of the possible chromaticities, weighted by an appropriate function. For example, the maximum volume estimate simply puts all the weight in the direction of the maximum coordinate product. Similarly, the average estimate weights the chromaticities by the volume of the cone in the corresponding direction.

Given this analogy, we can consider alternative methods of choosing a chromaticity solution. Since the maximum volume method tends to give better chromaticity estimates, especially when specularities are present, we wish to consider averages which put the bulk of the weight on solutions near the maximum volume direction. Now, one possible outcome of doing so would be the discovery that the maximum volume weighting worked the best. Interestingly, this proved not to be the case. Specifically we were able to find compromises which worked better.

We now present the weighting function developed for this work. Consider the solution set in mapping space. Then, each illuminant direction intersects the possible solution set at the origin, and at some other point. For an illuminant direction, i , let that other point be $(d_r^{(i)}, d_g^{(i)}, d_b^{(i)})$. Then, the functions we use to moderate the above weighting are powers of the geometric mean of coordinates of that mapping. Formally, we have parameterized functions f_N given by:

$$f_N(i) = (d_r^{(i)} d_g^{(i)} d_b^{(i)})^{N/3} \quad (1)$$

We note that the solution provided by the average of the mappings is roughly f_3 . The correspondence is not exact because the averaging is done over illuminant directions, not mapping directions. Similarly, as N becomes very large, the new method should approach the maximum volume method.

In order to use the above weighting function, we integrate numerically in polar coordinates. We discretize the polar coordinates of the illuminant directions inside a rectangular cone bounding the possible illuminant directions. We then test each illuminant direction as to whether it is a possible solution given the surface and illumination constraints. If it is, we compute the weighting function, and further multiply the result by the polar coordinate foreshortening, $\sin(\theta)$. We sum the results over the possible directions, and divide the total by the total weight to obtain the weighted average.

Finally we note that the above numerical integration, while clearly more computationally intensive than some of the precursor algorithms, still takes only a few seconds on a modern work station (even with using a very conservative discretization volume).

4 Experiments

We first consider the results for the method introduced to deal with diagonal model failure. Since the efficacy of the diagonal model is known to be a function of the camera sensors [1, 4, 5, 8, 12, 13], we provide results for two cameras with distinctly different degrees of support for the diagonal model. Our Sony DXC-930 video camera has quite sharp sensors, and with this camera, the changes in sensor responses to illumination changes can normally be well approximated with the

diagonal model. On the other hand, the Kodak DCS-200 digital camera [14] has less sharp sensors, and the diagonal model is less appropriate [6].

In the first experiment, we generated synthetic scenes with 4, 8, 16, 32, 65, 128, 256, 512, and 1024 surfaces. For each number of surfaces, we generated 1000 scenes with the surfaces randomly selected from the reflectance database and a randomly selected illuminant from the test illuminant database. For surface reflectances we used a set of 1995 spectra compiled from several sources. These surfaces included the 24 Macbeth colour checker patches, 1269 Munsell chips, 120 Dupont paint chips, 170 natural objects, the 350 surfaces in Krinov data set [15], and 57 additional surfaces measured by ourselves. The illuminant spectra set was constructed from a number of measured illuminants, augmented where necessary with random linear combinations, in order to have a set which was roughly uniform in (r,g) space. This data set is describe in more detail in [4, chapter 4].

For each algorithm and number of scenes we computed the RMS of the 1000 results. We choose RMS over the average because, on the assumption of roughly normally distributed errors with mean zero, the RMS gives us an estimate of the standard deviation of the algorithm estimates around the target. This is preferable to using the average of the magnitude of the errors, as those values are not normally distributed. Finally, given normal statistics, we can estimate the relative error in the RMS estimate by $1/\sqrt{2N}$ [16, p. 269] For $N=1000$, this is roughly 2%.

For each generated scene we computed the results of the various algorithms. We considered three-dimensional gamut mapping, with and without Finlayson's illumination constraint [9]. We will label the versions without the illumination constraint by CRULE, which is adopted from [1]. When the illumination constraint is added, we use the label ECRULE instead (Extended-CRULE). Solution selection using the maximum volume heuristic is identified by the suffix MV. For averaging in the case of CRULE, we use the suffix AVE, and in the case of ECRULE, we use the suffix ICA, indicating that the average was over the non-convex set (Illumination-Constrained-Average). This gives a total of four algorithms: CRULE-MV, CRULE-AVE, ECRULE-MV, and ECRULE-ICA. Finally, the method described above to reduce diagonal model failure will be indicated by the prefix ND (Non-Diagonal). We test this method in conjunction with each of the four previous algorithms, for a total of eight algorithms. We report the distance in (r,g) chromaticity space between the scene illuminant and the estimate thereof.

In Figure 3 we show the results for the Sony DXC-930 video camera. We see that when solution selection is done by averaging (AVE and ICA), the ND algorithms work distinctly better than their standard counter-parts. On the other hand, when solutions are chosen by the maximum volume heuristic, the ND algorithms performed slightly worse than their standard counterparts, provided that the number of surfaces was not large. Interestingly, as the number of surfaces becomes large, the error in all the ND versions continues to drop to zero, whereas the error in the standard versions levels off well above zero. In [4, page 92] we postulated that this latter behavior was due to the limitations of the diagonal model, and the present results confirm this.

In Figure 4 we show the results for the Kodak DCS-200 digital camera. The sensors of this camera do not support the diagonal model very well [6], and thus it is not surprising that the new extension significantly improves the performance of all four algorithms.




Fig. 3: Algorithm chromaticity performance versus the number of surfaces in generated scenes, showing the main gamut mapping algorithms and their non-diagonal counterparts. These results are for the Sony DXC-930 video camera which has relatively sharp sensors (the diagonal model is a good approximation in general). The error in the plotted values is roughly 2%.




Fig. 4: Algorithm chromaticity performance versus the number of surfaces in generated scenes, showing the main gamut mapping algorithms and their non-diagonal counterparts. These results are for the Kodak DCS-200 digital camera which has relatively dull sensors (the diagonal model is not very accurate). The error in the plotted values is roughly 2%.

We now turn to results with generated data for the solution selection method developed above. For this experiment we included the ND extension to reduce the confound of diagonal model failure. We label the new method with the suffix SCWIA (Surface-Constrained-Weighted-Illuminant-Average) followed by the value of the parameter N in Equation (1). The results are shown in Figure 5. First we point out that solution selection by the original averaging method out-performs the maximum volume heuristic when the number of surfaces is small, but as the number of surfaces increases, the maximum volume heuristic quickly becomes the preferred method.

Turning to the new method, we see that it indeed offers a compromise between these two existing methods, with the new method tending towards the maximum volume method as N increases. More importantly, as long as N is 6 or more, the new method invariably outperforms solution selection by averaging. Furthermore, for N in the range of 9-24, the performance of the new method is better than the maximum volume heuristic, except when the number of surfaces is unusually large. When the number of surfaces becomes large, the maximum volume heuristic eventually wins out.

An important observation is that the results for N in the range of 9-24 are quite close, especially around 8 surfaces. This is fortuitous, as we have previously observed [4] that 8 synthetic surfaces is roughly comparable in difficulty to our image data. Thus we are most interested in improving performance in the range of 4-16 surfaces, and we are encouraged that the results here are not overly sensitive to N , provided that it is roughly correct. Based on our results, $N=12$ appears to be a good compromise value for general purpose use.

Next we present some numerical results in the case of the Sony camera which shows the interactions of the two modifications. These are shown in Table 1. The main point illustrated in this table is that the slight disadvantage of the ND method, when used in conjunction with MV, does not carry over to the new solution selection method. To explain further, we note that the positive effect of reducing the diagonal model error can be undermined by the expansion of the canonical gamut, which represents an increase in the size of the feasible sets. The positive effect occurs because these sets are more appropriate, but, all things being equal, their larger size is an increase in ambiguity. Thus when the ND method is used in conjunction with a camera which supports the diagonal model, then, as the results here show, the method can lead to a decrease in performance. In our experiments on generated data, the negative effect is present in the case of MV, but in the case of averaging, the effect is always slightly positive. When ND is used in conjunction with the new solution method, the results are also minimally compromised by this negative effect. This is very promising, because, in general, the diagonal model will be less appropriate, and the method will go from having little negative impact, to having a substantial positive effect. This has already been shown in the case of the Kodak DCS-200 camera, as well as when the number of surfaces is large. Increasing the number of surfaces does not, of course, reduce the efficacy of the diagonal model, but under these conditions, the diagonal model becomes a limiting factor.




Fig. 5: Algorithm chromaticity performance versus the number of surfaces in generated scenes, showing the selected gamut mapping algorithms, including ones with the new solution selection method. These results are for the Sony DXC-930 video camera. The error in the plotted values is roughly 2%.

Number of Surfaces	4	8	16
ECRULE-MV	0.064	0.044	0.032
ECRULE-ICA	0.058	0.050	0.044
ECRULE-SCWIA-3	0.057	0.051	0.045
ECRULE-SCWIA-6	0.054	0.043	0.036
ECRULE-SCWIA-9	0.054	0.041	0.032
ECRULE-SCWIA-12	0.055	0.040	0.031
ECRULE-SCWIA-18	0.057	0.040	0.030
ECRULE-SCWIA-24	0.058	0.041	0.029
ND-ECRULE-MV	0.065	0.047	0.033
ND-ECRULE-ICA	0.057	0.049	0.043
ND-ECRULE-SCWIA-3	0.060	0.054	0.048
ND-ECRULE-SCWIA-6	0.054	0.044	0.036
ND-ECRULE-SCWIA-9	0.054	0.041	0.031
ND-ECRULE-SCWIA-12	0.055	0.041	0.030
ND-ECRULE-SCWIA-18	0.057	0.041	0.029
ND-ECRULE-SCWIA-24	0.059	0.042	0.029

Table 1: Algorithm chromaticity performance for some of the algorithms developed here, together with the original methods, for generated scenes with 4, 8, and 16 surfaces. The numbers are the RMS value of 1000 measurements. The error in the values is roughly 2%.

Finally we turn to results with 321 carefully calibrated images. These images were of 30 scenes under 11 different illuminants (9 were culled due to problems). The images are describe more fully in [4]. Figure 6 shows the 30 scenes used. We provide the results of some of the algorithms discussed above, as well as several comparison methods. We use NOTHING to indicated the result of no colour constancy processing, and AVE-ILLUM for guessing that the illuminant is the average of a normalized illuminant database. The method labeled MAX estimates the illuminant RGB by the maximum found in each channel. GW estimates the illuminant based on the image average on the assumption that the average is the response to a perfect grey. DB-GW is similar, except that the average is now assumed to be the response to grey as defined by the average of a reflectance database. CIP-ICA is essentially a chromaticity version of ECRULE-ICA described in [11]. The method labeled NEURAL-NET is another chromaticity oriented algorithm which uses a neural net to estimate the illuminant chromaticity [17-19]. C-by-C-MAP is the Colour by Correlation method using the maximum posterior estimator [20]. Finally, C-by-C-MSE is Colour by Correlation using the minimum mean square error estimate. All these comparison methods are described in detail in [4]

Table 2 shows the results over the 321 test images. The results from the image data generally confirm those from the generated data in the case of the new selection method. On the other hand, the ND method improves matters significantly in only one, case, has essentially no effect in several others, and when used in

conjunction with the new selection method, it has a small negative effect. Since the camera used already supports the diagonal model well, these varied results are understandable.

Algorithm	Solution Selection Method (If Applicable)					
	MV	AVE/ICA	SCWIA-6	SCWIA-9	SCWIA-12	SCWIA-15
CRULE	0.045	0.046				
ECRULE	0.041	0.047	0.043	0.039	0.037	0.037
ND-CRULE	0.047	0.039				
ND-ECRULE	0.042	0.048	0.045	0.041	0.040	0.040
NOTHING	0.125					
AVE-ILLUM	0.094					
GW	0.106					
DB-GW	0.088					
MAX	0.062					
CIP-ICA	0.081					
NEURAL-NET	0.069					
C-by-C-MAP	0.072					
C-by-C-MMSE	0.070					

Table 2: The image data results of the new algorithms compared to related algorithms. The numbers presented here are the RMS value of the results for 321 images. Assuming normal statistics, the error in these numbers is roughly 4%.

6 Conclusion

We have described two improvements to gamut mapping colour constancy. These improvements are important because earlier work has shown that this approach is already one of the most promising. For the first improvement we modified the canonical gamuts used by these algorithms to account for expected failures of the diagonal model. When used with a camera which does not support the diagonal model very well, the new method was clearly superior. When used with a camera with sharp sensors, the resulting method improved gamut mapping algorithms when the solution was chosen by averaging. When the maximum volume heuristic was used, there was a slight decrease in performance. This decrease was erased when the method was combined with the second improvement. Furthermore, we posit that any decreases in performance must be balanced against the increased stability of the new method as the number of surfaces becomes large.

We are also encouraged by the results of the new method for choosing the solution. Our findings contribute to the understanding of the relative behavior of the two existing methods. Furthermore, the flexibility of the new method allows us to select a variant which works better than either of the two existing methods for the kind input we are most interested in.

7 References

1. D. Forsyth, A novel algorithm for color constancy, *International Journal of Computer Vision*, 5, pp. 5-36 (1990).
2. K. Barnard, Computational colour constancy: taking theory into practice, M.Sc. Thesis, Simon Fraser University, School of Computing (1995).
3. B. Funt, K. Barnard, and L. Martin, Is Colour Constancy Good Enough?, *Proc. 5th European Conference on Computer Vision*, pp. I:445-459 (1998).
4. K. Barnard, Practical colour constancy, Ph.D. Thesis, Simon Fraser University, School of Computing Science (1999), <ftp://fas.sfu.ca/pub/cs/theses/1999/KobusBarnardPhD.ps.gz>
5. G. D. Finlayson, M. S. Drew, and B. V. Funt, Spectral Sharpening: Sensor Transformations for Improved Color Constancy, *Journal of the Optical Society of America A*, **11**, pp. 1553-1563 (1994).
6. K. Barnard and B. Funt, Experiments in Sensor Sharpening for Color Constancy, *Proc. IS&T/SID Sixth Color Imaging Conference: Color Science, Systems and Applications*, pp. 43-46 (1998).
7. K. Barnard, Color constancy with fluorescent surfaces, *Proc. IS&T/SID Seventh Color Imaging Conference: Color Science, Systems and Applications* (1999).
8. G. D. Finlayson, Coefficient Color Constancy, : Simon Fraser University, School of Computing (1995).
9. G. D. Finlayson, Color in perspective, *IEEE Transactions on Pattern Analysis and Machine Intelligence*, **18**, pp. 1034-1038 (1996).
10. G. Finlayson and S. Hordley, A theory of selection for gamut mapping colour constancy, *Proc. IEEE Conference on Computer Vision and Pattern Recognition* (1998).
11. G. Finlayson and S. Hordley, Selection for gamut mapping colour constancy, *Proc. British Machine Vision Conference* (1997).
12. J. A. Worthey, Limitations of color constancy, *Journal of the Optical Society of America [Suppl.]*, **2**, pp. 1014-1026 (1985).
13. J. A. Worthey and M. H. Brill, Heuristic analysis of von Kries color constancy, *Journal of the Optical Society of America A*, **3**, pp. 1708-1712 (1986).
14. P. L. Vora, J. E. Farrell, J. D. Tietz, and D. H. Brainard, Digital color cameras--Spectral response, (1997), available from <http://color.psych.ucsb.edu/hyperspectral/>.
15. E. L. Krinov, *Spectral Reflectance Properties of Natural Formations*: National Research Council of Canada, 1947.
16. J. L. Devore, *Probability and Statistics for Engineering and the Sciences*. Monterey, CA: Brooks/Cole Publishing Company, 1982.
17. B. Funt, V. Cardei, and K. Barnard, Learning Color Constancy, *Proc. IS&T/SID Fourth Color Imaging Conference: Color Science, Systems and Applications*, pp. 58-60 (1996).
18. V. Cardei, B. Funt, and K. Barnard, Modeling color constancy with neural networks, *Proc. International Conference on Vision Recognition, Action: Neural Models of Mind and Machine* (1997).
19. V. Cardei, B. Funt, and K. Barnard, Adaptive Illuminant Estimation Using Neural Networks, *Proc. International Conference on Artificial Neural Networks* (1998).
20. G. D. Finlayson, P. H. Hubel, and S. Hordley, Color by Correlation, *Proc. IS&T/SID Fifth Color Imaging Conference: Color Science, Systems and Applications*, pp. 6-11 (1997).