Color Constancy and a Changing Illumination

Graham D. Finlayson
School of Computing Science
Simon Fraser University
Vancouver, B.C., Canada V5A 1S6

Abstract

The color constancy problem has proven to be very hard to solve. This is even true in the simple Mondriaan world where a planar patchwork of matte surfaces is viewed under a single illuminant. In this paper we consider the color constancy problem given two images of a Mondriaan viewed under different illuminants.

We show that if surface reflectances are well-modelled by 3 basis functions and illuminants by *up to* 5 basis functions then we can, theoretically, solve for color constancy given 3 surfaces viewed under 2 illuminants. The number of recoverable dimensions in the illuminant depends on the spectral characteristics of the sensors. Specifically if for a given sensor set a von Kries type, diagonal model of color constancy is sufficient then we show that at most 2 illuminant parameters can be retrieved.

Recent work has demonstrated that for the human visual system a diagonal matrix is a good model of color constancy given an appropriate choice of sensor basis. We might predict therefore, that we can recover at most 2 illuminant parameters. We present simulations which indicate that this is in fact the case.

1 Introduction

Under different illuminants the same surface reflects different spectra of light; however, despite this we see the same color. This is the phenomenon of *color constancy*. Despite extensive research there does not yet exist a computational theory sufficient to explain the color constancy performance of a human observer.

A common starting point for color constancy research is the following question:

"Given an image of several surfaces viewed under a single illuminant how can we derive illuminant independent surface descriptors?"

This basic problem is tremendously hard to solve. Edwin Land's famous retinex theory[15,16,17] (and Hurlbert's[13] subsequent extension) are easily shown to be inadequate for the task.[1] Forsyth's recent theory[7] though more powerful requires many chromatically distinct surfaces in each scene.

Consequently the basic color constancy question is relaxed and computational theories often incorporate other factors into the problem formulation. Statistical analyses of reflectance and illuminant spectra are at the heart of many approaches,[18,20,2,11,8] Funt and Ho[10] demonstrated that the chromatic aberration inherent in every lens can provide useful color information, Shafer[23] provides a method of determining the illuminant color given specularities and Funt et. al[9] have shown that mutual reflection occurring at a concave edge ameliorates the color constancy problem.

Tsukada et. al[24] and D'Zmura[4] have considered the color constancy problem where the illumination changes. Specifically they asked:

"Given an image of several surfaces viewed under two illuminants how can we derive illuminant independent surface descriptors?"

This question is particularly relevant since Craven and Foster[3] have recently demonstrated that an illumination change is easily discernible and moreover can be distinguished from reflectance (i.e. false illuminant) changes. D'Zmura[4] has shown that if illuminants and surface reflectances are well described by finite dimensional models each of 3 dimensions (the *3-3* world) then the color constancy problem can be solved given 3 surfaces seen under two illuminants. D'Zmura and Iverson[5] generalized this result and have presented an algorithm which can solve for color constancy given many different assumptions—for example different numbers of reflectances or different model dimensions. However the results of the present paper are not part of D'Zmura and Iverson's general theory.

In this paper we begin by providing an alternate analysis for the *3-3* problem. Unlike D'Zmura's method our approach generalizes to more than 3 patches. Further, in the case where reflectances and illuminants are not precisely 3-dimensional our approach provides a least-squares criterion in solving for color constancy (D'Zmura's analysis is for exact models only). We extend our basic method and demonstrate that it is theoretically possible to solve for 5 parameters in the illuminant and 3 for reflectances given 3 patches seen under two illuminants. Under this generalized model every bijective linear map corresponds to a valid illumination change.

We show that there are classes of sensors for which the 2 illuminant, *3-3* color constancy problem **cannot be solved**. Specifically if for a given sensor set a diagonal matrix is a good model of illuminant change then the color constancy problem cannot be solved.

The world is not precisely *3-3* nor even *5-3* (in illuminant and reflectance dimension); therefore we must ask where our computational model will succeed in solving the color constancy problem. Using the human eye sensitivities we test our algorithm's color constancy performance through simulation experiments. We come to the following surprising conclusion: illumination change is useful in solving for color constancy if we assume the world is *2-3*; that is we assume the **illumination is 2-dimensional**. Moreover the greater the change in illuminant color the greater the likelihood of correctly solving for color constancy.

Assuming a 2-dimensional illumination has widespread implications. Firstly leading from the work of Maloney-Wandell[20] it is straightforward to show that the color constancy problem can be solved given a single patch under 2 illuminants. Secondly color constancy is still possible even where a diagonal matrix is a good model of illumination change. This is reassuring since Finlayson et. al[6] have shown that under the *2-3* assumptions there exists a sensor basis for all trichromatic visual systems for which a diagonal matrix precisely models all illuminant change. The analysis presented in this paper therefore, serves to strengthen diagonal matrix theories of color constancy.

In section 2 we formulate the color constancy problem under changing illumination. Our solution method for both the *3-3* and *5-3* method is detailed in section 3. In section 4 given the human eye cone sensitivities we consider when the color constancy problem can in practice be solved. Section 5 presents simulations which demonstrate that color constancy can be solved by assuming a *2-3* model and given large color shifts in the illuminant.

2 The Color Constancy Problem

In keeping with D'Zmura we develop our solution method for the Mondriaan world: a Mondriaan is a planar, matte surface with several different, uniformly colored patches. Light striking the Mondriaan is assumed to be of uniform intensity and is spectrally unchanging. Each Mondriaan is assumed to contain at least 3 distinct surfaces.

If $E(\lambda)$ is the illuminant incident to surface reflectance $S(\lambda)$, where λ indexes wavelength, then the reflected color signal is equal to:

$$C(\lambda) = E(\lambda)S(\lambda) \qquad \text{(1: matte reflectance)}$$

The value registered by the kth cone to the color signal $C(\lambda)$ is defined by the integral equation (where ω denotes the visible spectrum):

$$\underline{p}_k = \int_\omega C(\lambda)R_k(\lambda)d\lambda \qquad \text{(2: color response)}$$

The illuminant, reflectance or sensitivity functions are known only for a set of sampled wavelengths (in this paper all spectra are in the range 400nm to 650nm measured at 10nm intervals). Therefore the integral, of equation (2), is approximated as a summation.

2.1 Modelling Reflectance and Illuminant Spectra

Both illuminant spectral power distribution functions and surface spectral reflectance functions are well described by finite-dimensional models of low dimension. A surface reflectance vector $S(\lambda)$ can be approximated as:

$$S(\lambda) \approx \sum_{i=1}^{d_S} S_i(\lambda)\sigma_i \qquad (3)$$

where $S_i(\lambda)$ is a basis function and $\underline{\sigma}$ is a d_S-component column vector of weights. Similarly each illuminant can be written as:

$$E(\lambda) \approx \sum_{j=1}^{d_E} E_j(\lambda)\epsilon_j \qquad (4)$$

where $E_j(\lambda)$ is a basis function and $\underline{\epsilon}$ is a d_E dimensional vector of weights.

Maloney[18] presented a statistical analysis of reflectance spectra and concluded that between 3 and 6 basis vectors are required to model surface reflectance. We will assume a 3-dimensional reflectance model. A similar analysis for daylight illumination was carried out by Judd[14]; daylight illuminants are well represented by 3 basis vectors.

Given finite-dimensional approximations to surface reflectance, the color response eqn. (2) can be rewritten as a matrix equation. A *Lighting Matrix* $\Lambda(\underline{\epsilon})$ maps reflectances, defined by the $\underline{\sigma}$ vector, onto a corresponding color response vector:

$$\underline{p} = \Lambda(\underline{\epsilon})\underline{\sigma} \qquad (5)$$

where $\Lambda(\underline{\epsilon})_{ij} = \int_\omega R_i(\lambda)E(\lambda)S_j(\lambda)d\lambda$. The lighting matrix is dependent on the illuminant weighting vector $\underline{\epsilon}$, with $E(\lambda)$ given by eqn. (4). The lighting matrix corresponding to the ith illuminant basis function is denoted as Λ^i.

2.2 Color Constancy under 2 Illuminants

Let us denote the 2 illuminants by the weight vectors $\underline{\epsilon}^1$ and $\underline{\epsilon}^2$. Reflectances are denoted as $\underline{\sigma}_1, \underline{\sigma}_2, \ldots, \underline{\sigma}_k$ where $k \geq 3$. Arranging the k sigma vectors as the columns of the matrix $\Omega_{1,2,\ldots,k}$ we can write the color responses under the two illuminants as:

$$P^1 = \Lambda(\underline{\epsilon}^1)\Omega_{1,2,\ldots,k} \qquad P^2 = \Lambda(\underline{\epsilon}^2)\Omega_{1,2,\ldots,k} \qquad (6)$$

where the ith column of P^j is the response of the ith surface under the jth illuminant ($j = 1$ or $j = 2$). Given P^1 and P^2 we want to solve for $\underline{\epsilon}^1$, $\underline{\epsilon}^2$ and $\Omega_{1,2,\ldots,k}$.

3 The Color Constancy Solution

We propose solving for color constancy in 2 steps. First we calculate the linear transform mapping the color responses between illuminants. This transform is independent of $\Omega_{1,2,\ldots,k}$ and, as we shall show, provides an elegant means of determining $\underline{\epsilon}^1$ and $\underline{\epsilon}^2$. By calculating $[\Lambda(\underline{\epsilon}^1)]^{-1}$, we can easily recover $\Omega_{1,2,\ldots,k}$.

In the *3-3* world lighting matrices have 3 rows and 3 columns. An implication of this is that color responses under one illuminant can be mapped to corresponding responses under a second illuminant by the application of a 3×3 linear transform; we call this an *illuminant map*.

$$\mathcal{M}^{1,2}\Lambda(\underline{\epsilon}^1) = \Lambda(\underline{\epsilon}^2), \qquad \mathcal{M}^{1,2} = \Lambda(\underline{\epsilon}^2)[\Lambda(\underline{\epsilon}^1)]^{-1}, \qquad \text{(7: illuminant map)}$$

Theorem 1 *The color constancy problem can only be solved if for each pair of illuminants* $s_1\underline{\epsilon}^1$ *and* $s_2\underline{\epsilon}^2$ (s_1, s_2 *are scalars) the corresponding illuminant map* $\mathcal{M}^{1,2}$ *as unique, up to a scaling, over all other illuminant pairs.*

Proof of Theorem 1: Assume we have 4 illuminants $\underline{\epsilon}^a\underline{\epsilon}^b\underline{\epsilon}^c$ and $\underline{\epsilon}^d$ such that $\mathcal{M}^{a,b} = \mathcal{M}^{c,d}$ where $\underline{\epsilon}^b$ and $\underline{\epsilon}^d$ are linearly independent. If $\Lambda(\underline{\epsilon}^a) = s\Lambda(\underline{\epsilon}^c)$ then Theorem 1 follows since b and d are, by assumption, linearly independent and consequently $\mathcal{M}^{a,b} \neq \mathcal{M}^{c,d}$. Otherwise let $\Omega_{1,2,3}$ denote a matrix of 3 reflectances. Let us define $\Omega^+_{1,2,3}$ such that:

$$P^b = \Lambda(\underline{\in}^b)\Omega_{1,2,3} \equiv P^d = \Lambda(\underline{\in}^d)\Omega^+_{1,2,3} \qquad (8)$$

Thus the reflectances $\Omega_{1,2,3}$ viewed under illuminant $\underline{\in}^b$ cannot be distinguished from the reflectances $\Omega^+_{1,2,3}$ viewed under illuminant $\underline{\in}^d$. That is, the color constancy problem cannot be solved if $\mathcal{M}^{1,2}$ is not unique over all other illuminant pairs. If $\mathcal{M}^{1,2}$ is not unique we cannot hope to separately recover both the $\underline{\in}$ vectors (or at least the $\mathcal{M}^{1,2}$ matrix) as well as the $\underline{\sigma}$ vectors.

Let us assume that all pairs of illuminants, in our 3-dimensional span, have a corresponding unique illuminant map. Given the mapping $\mathcal{M}^{1,2}$ then:

$$\mathcal{M}^{1,2}\left[\Lambda^1 \in_1^1 + \Lambda^2 \in_2^1 + \Lambda^3 \in_3^1\right] = \left[\Lambda^1 \in_1^2 + \Lambda^2 \in_2^2 + \Lambda^3 \in_3^2\right] \qquad (9)$$

Let $\mathcal{P}$ denote a 9×3 matrix where the ith column contains the basic lighting matrix Λ^i stretched out columnwise. Similarly let $\mathcal{Q}$ denote the 9×3 matrix where the ith column contains $\mathcal{M}^{1,2}\Lambda^i$. Rewriting equation (9):

$$\mathcal{Q}\underline{\in}^1 = \mathcal{P}\underline{\in}^2 \qquad (10)$$

The columns of the matrix $\mathcal{P}$ are a basis for a 3-dimensional subspace of 9-space. Similarly the columns of $\mathcal{Q}$ are a basis for a 3-dimensional subspace of 9-space. The solution of equation (10) is the intersection of these two spaces. The intersection is easily found by the method of principal angles.[12] This method finds the vectors $\underline{\in}^1$ and $\underline{\in}^2$ which maximizes:

$$\cos^2(\theta) = \left(\frac{(\mathcal{Q}\underline{\in}^1)\cdot(\mathcal{P}\underline{\in}^2)}{\left|\mathcal{Q}\underline{\in}^1\right| \left|\mathcal{P}\underline{\in}^2\right|} \right)^2 \qquad (11)$$

That is $\underline{\in}^1$ and $\underline{\in}^2$ are chosen such that the angle, θ, between $\mathcal{Q}\underline{\in}^1$ and $\mathcal{P}\underline{\in}^2$ is minimized. Thus even when there does not exist an exact solution to equation 10, the method of principal angles provides a least-squares criterion for returning the *best* answer.

3.1 Robust Color Constancy

While reflectances may in general be well described by a 3-parameter reflectance model a particular set of three reflectances may be poorly modelled. Consequently $\mathcal{M}^{1,2}$ will be incorrectly estimated. In this case the color constancy algorithm may return incorrect estimates for the reflectance and illuminant parameters. However if the parameters of the illuminant map are derived from the observations of many (greater than three) distinct reflectances then we would expect improved color constancy.

Let P^1 denote a $3 \times n$ matrix of n reflectances observed under an arbitrary illuminant. Similarly P^2 is the $3 \times n$ matrix of observations of the same reflectances viewed under a second illuminant. The best illuminant map, in the least-squares sense, taking P^1 onto P^2 is defined by the Moore-Penrose inverse:

$$\mathcal{M}^{1,2}P^1 \approx P^2, \qquad \mathcal{M}^{1,2} = P^2(P^1)^t[P^1(P^1)^t]^{-1} \qquad (12)$$

In section 5 we present simulations where the illuminant map is derived first from 3 and then from 6 reflectances. The greater the number of reflectances the better the color constancy performance.

3.2 Solving for More Illuminant Parameters

The columns of $\mathcal{P}$ and $\mathcal{Q}$ of equation (9) are bases for 3-dimensional subspaces of 9-space. It is possible, therefore, that $\mathcal{P}$ and $\mathcal{Q}$ have a null-intersection—the combined 9×6 matrix $[\mathcal{P}\,\mathcal{Q}]$ has full rank; all columns are linearly independent. A null intersection is indicative of the fact that the world is not *3-3*. We might ask therefore, if it is possible to extend our model assumptions such that any illuminant map falls within our model.

The columns of $\mathcal{P}$ correspond to the three basis lighting matrices; the columns of $\mathcal{Q}$ correspond to this basis transformed by an illuminant mapping. Clearly if we increase the dimension of the illuminant model from 3 to 5 then $\mathcal{P}$ and $\mathcal{Q}$ become 9×5 matrices and are bases for 5-dimensional subspaces. In the 5-dimensional case the combined matrix $[\mathcal{P}\,\mathcal{Q}]$ has 9 rows and 10 columns. If the first 9 columns of $[\mathcal{P}\,\mathcal{Q}]$ are linearly independent then they form a basis for 9-space. Consequently the 10th column is guaranteed to be linearly dependent on the first 9. Thus in the *5-3* case intersection is assured; indeed all bijective linear maps correspond to a valid illuminant mapping.

It is interesting to note that the *5-3* world does not belong to the general color constancy formulation of D'Zmura and Iverson.[5] Our work, therefore, supplements this general theory.

4 When Color Constancy Can Be Solved

So far we have assumed that the illuminant map is unique, and consequently from Theorem 1, there exists a solution to the color constancy problem. If the world is *3-3* and the illuminant mapping is not unique then we show here that there exists a sensor basis such that all illuminant mappings are diagonal matrices.

Theorem 2 *If the illuminant mapping is non-unique then there exists a sensor transformation $\mathcal{T}$, such that for all illuminants* 1 *and* 2, *$\mathcal{T}\mathcal{M}^{1,2}\mathcal{T}^{-1}$ is a diagonal matrix.*

Proof of Theorem 2: Assume we have 4 illuminants $\underline{\in}^a$, $\underline{\in}^b$, $\underline{\in}^c$ and $\underline{\in}^d$ such that $\mathcal{M}^{a,b} = \mathcal{M}^{c,d}$ ($\mathcal{M}^{a,b} \neq \mathcal{T}$) and $\Lambda(\underline{\in}^a) \neq s_1\Lambda(\underline{\in}^c)$ and $\Lambda(\underline{\in}^b) \neq s_2\Lambda(\underline{\in}^d)$ (where s_1 and s_2 are scalars). Because we are assuming a *3-3* world, there are only 3 linearly independent lighting matrices and therefore we can choose scalars α, β and γ such that:

$$\Lambda(\underline{\in}^a) + \alpha\Lambda(\underline{\in}^b) = \beta\Lambda(\underline{\in}^c) + \gamma\Lambda(\underline{\in}^d) \qquad (13)$$

Denoting the identity matrix as I, we can write the mapping of illuminant a to illuminant c as:

$$\mathcal{M}^{a,c} = [\beta I + \gamma \mathrm{M}^{c,d}]^{-1}\,[I + \alpha\mathcal{M}^{a,b}] \qquad (14)$$

Rewriting both $\mathcal{M}^{a,b}$ and $\mathcal{M}^{c,d}$, by assumption they are equal, as $\mathcal{T}^{-1}\mathcal{D}\mathcal{T}$ and I as $\mathcal{T}^{-1}I\mathcal{T}$ equation (14) becomes:

$$\mathcal{M}^{a,c} = \mathcal{T}^{-1}[\beta I + \gamma\mathcal{D}]^{-1}\mathcal{T}\mathcal{T}^{-1}[I + \alpha\mathcal{D}]\mathcal{T} \qquad (15)$$

$$\mathcal{M}^{a,c} = \mathcal{T}^{-1}[\beta I + \gamma\mathcal{D}]^{-1}[I + \alpha\mathcal{D}]\mathcal{T} \qquad (16)$$

The lighting matrices $\Lambda(\underline{\in}^a)$, $\Lambda(\underline{\in}^b)$ and $\Lambda(\underline{\in}^c)$ are linearly independent and span the space of lighting matrices; any lighting matrix can be written as a linear combination of these three. Since $\mathcal{M}^{a,b}$ and $\mathcal{M}^{a,c}$ have the same eigenvectors all illuminant mappings from a must have the same eigenvectors:

$$\mathcal{M}^{a,s_1a+s_2b+s_3c} = \mathcal{T}^{-1}\left[s_1I + s_2\mathcal{D}^{a,b} + s_3\mathcal{D}^{a,c}\right]\mathcal{T} \qquad (17)$$

s_1, s_2 and s_3 are scalars defining an arbitrary illuminant and $\mathcal{D}^{a,b}$ and $\mathcal{D}^{a,c}$ are the diagonal matrices of eigenvalues for $\mathcal{M}^{a,b}$ and $\mathcal{M}^{a,c}$ respectively. It is a simple step to show that all illuminant mappings share the same eigenvectors. Consider the illuminant mapping $\mathcal{M}^{s_1a+s_2b+s_3c,t_1a+t_2b+t_3c}$ where s_i and t_j are arbitrary scalars. Employing equation (17) we can write this as:

$$\mathcal{M}^{s_1a+s_2b+s_3c,t_1a+t_2b+t_3c} = \qquad (18)$$

$$\mathcal{T}^{-1}\left[t_1I + t_2\mathcal{D}^{a,b} + t_3\mathcal{D}^{a,c}\right]\mathcal{T}\mathcal{T}^{-1}\left[s_1I + s_2\mathcal{D}^{a,b} + s_3\mathcal{D}^{a,c}\right]^{-1}\mathcal{T}$$

which is equal to

$$\mathcal{M}^{s_1a+s_2b+s_3c,t_1a+t_2b+t_3c} = \qquad (19)$$

$$\mathcal{T}^{-1}\left[t_1I + t_2\mathcal{D}^{a,b} + t_3\mathcal{D}^{a,c}\right]\left[s_1I + s_2\mathcal{D}^{a,b} + s_3\mathcal{D}^{a,c}\right]^{-1}\mathcal{T}$$

This completes the proof of Theorem 2.

Corollary 1 *If the illuminant mapping is non-unique, then from each Lighting matriz* $\Lambda(\in)$ *the set of all matrices of the form* $\mathcal{T}^{-1}\mathcal{D}\mathcal{T}$ *exactly characterize the set of valid illumanant maps. The spaces spanned by* $\mathcal{P}$ *and* $\mathcal{Q}$ *intersect in all 3 dimensions.*

Theorem 2 provides a useful test, see Figure 2, for determining whether a change in illumination adds extra information to the color constancy process. If the eigenvectors of the mapping taking the first lighting matrix to the second is *different* from the eigenvectors of the mapping taking the first to the third then a change of lighting adds information to the color constancy process.

$$\mathcal{TM}^{1,2}\mathcal{T}^{-1} = \mathcal{D},\ \ \mathcal{TM}^{1,3}\mathcal{T}^{-1} \neq \mathcal{D}',$$

Figure 1. Color Constancy Check for 3-3 world: If check fails then color constancy is as hard under 2 illuminants as under a single illuminant.

The question of when a change of illumination is useful in solving for color constancy raises a paradox. Entrenched in color constancy research is the notion that if illumination change is well modelled by a diagonal matrix then it is *easier* to solve for color constancy and this is in direct conflict with what we have shown.

In the next section we present simulations which go some way to resolving this paradox. We examine our algorithm's performance given reflectances viewed under pairs of illuminants. The best color constancy is attained assuming *2-3* conditions, as oppose to *3-3* or *3-5*. Under *2-3* conditions illumination change is always perfectly modelled by a diagonal matrix (with respect to an appropriate basis),[6] and more importantly by losing one degree of freedom in the illuminant model, every diagonal matrix corresponds uniquely to a pair of illuminants; that is, color constancy is soluble.

5 Simulations

For our illuminants we chose a set of 7 Planckian black body radiators with correlated color temperatures 2000K, 2856K (CIE A), 4000K, 6000K, 8000K, 10000K and 20000K. Reflectances are drawn from the set of 462 Munsell[22] spectra. Illuminant bases of dimension 2, 3 and 5 are derived from an ensemble set of 14 illuminants containing the Planckian radiators, 5 daylight phases and CIE B and CIE

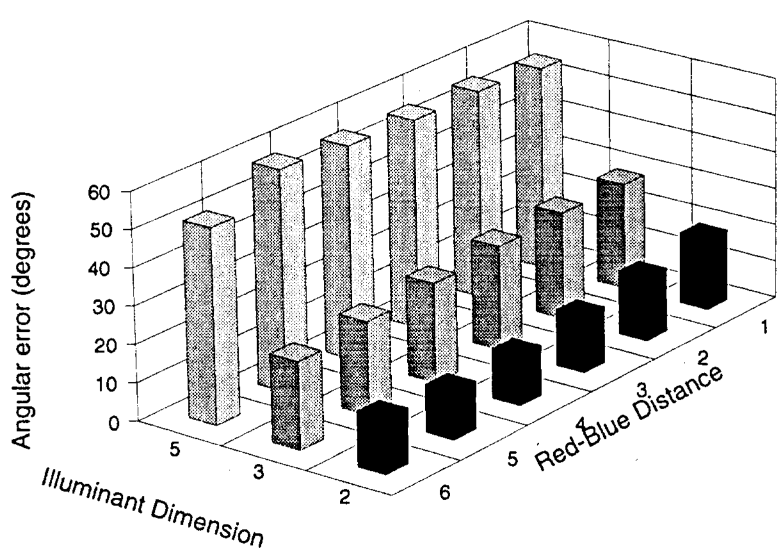

Figure 2. Randomly selected sets of 3 reflectances are imaged under 2 illuminants. The average angular error in the recovery of the second illuminant was calculated for illuminant models of 2-, 3- and 5-dimensions and red-blue distances of 1 through 6.

C. A 3-dimensional reflectance basis is derived from the complete Munsell set.

We proceed in the following manner. First, 3 reflectances are randomly drawn from the set of 462 Munsell Spectra. For each pair of black-body radiators, $B_i(\lambda)$, $B_j(\lambda)$, we simulate the color response of the eye (equation 2) using the cone fundamentals measured by Vos and Walraven.[25] We run our color constancy algorithm three separate times; using the 2-, 3- and 5-dimensional illuminant models. We record the recovery error as the angle between $B_j(\lambda)$ and that returned by our algorithm ($\underline{\in}^2$ of equation (10)):

$$err = angle(B_j(\lambda), \sum_{k=1}^{d_E} E_k(\lambda) \in_k^2) \tag{20}$$

As the recovered spectra better approximates the actual spectra so the error decreases towards zero. This experiment was repeated 5-times and the average angular error calculated for each illuminant pair.

The 2000K black-body radiator is a red biased spectrum, the 2856K radiator is still red but has a greater blue component. This trend continues with each of the 4000K, 6000K, 8000K, 10000K and 20000K radiators becoming progressively bluer. If the first illuminant is the 2000K spectra and the second 2856K then these spectra differ by 1 red-blue position. If the second illuminant is 4000K then this distance is 2. Subject to this red to blue ordering in our illuminant set we further average the error values. We calculate the average angular error given a red-blue difference of 1, 2, 3, 4, 5 and 6. We graph the color constancy performance, given these red-blue distances, for 2-, 3- and 5-dimensional illumination models in Figure 2.

In all cases the 2-dimensional assumption returns better color constancy. Moreover as the color constancy performance generally improves as the red-blue distance increases—a distance greater than 2 and the average angular error is less than or equal to 14 degrees. Under the 3-dimensional assumption the angular error is much higher, on the order of 18 degrees throughout. However there is a discernible performance improvement given red-blue distances of greater than 4. The 5-dimensional assumption returns extremely poor color constancy with angular error always larger than 50 degrees.

We repeated this experiment for random selections of 6 patches. The results are graphed in Figure 3. Both the 2- and 3-dimensional assumptions show marked improvement; with the 2-dimensional assumption still supporting substantially better color constancy. As before the 5-dimensional assumption pays very poor dividends with the minimum angular error of 48 degrees. The error distribution is graphed in Figure 3.

That the 2-dimensional illuminant assumptions returns the best color constancy, at first glance, appears surprising. However previous simulations[21] have demonstrated that *2-3* assumptions provide a reasonable model for approximating cone responses. Moving to higher dimensional illumination models improves, slightly, on this approximation but at the cost of introducing very many more valid illuminant mappings. As the number of mappings increases so does the likelihood of a false match. A more theoretical treatment of this observation lies out with the scope of this paper.

6 Conclusion

We have developed a computational framework for solving for color constancy under a change of illuminant. The framework is general in the sense that the computation remains the same under different illuminant model assumptions.

We derived a test to determine whether a change in illumination adds new information to the color constancy problem given a 3-dimensional illuminant. If a diagonal matrix, with respect to a sensor basis, is a good model of illuminant change then a change of illumination does not add new information. This is paradoxical in that it contradicts the established view that diagonal matrix color constancy is *easier* than non-diagonal methods.

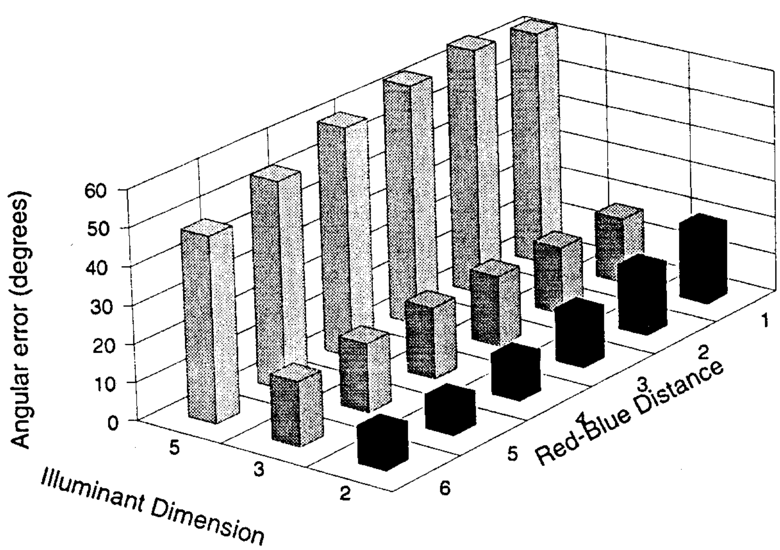

Figure 3. Randomly selected groups of 6 reflectances are imaged under 2 illuminants. The average angular error in the recovery of the second illuminant was calculated for illuminant models of 2-, 3- and 5-dimensions and red-blue distances of 1 through 6.

Simulation experiments go some way to resolving this paradox. We show that a 2-dimensional illuminant assumption supports better color constancy than 3- or 5-dimensional assumptions. A 2-dimensional assumption is completely consistent with diagonal theories of color constancy.

References

1. D.A. Brainard and B.A. Wandell. Analysis of the retinex theory of color vision. *J. Opt. Soc. Am. A*, **36**:1651-1661, 1986.
2. G. Buchsbaum. A spatial processor model for object colour perception. *Journal of the Franklin Institute*, **310**:1-26, 1980.
3. B.J. Craven and D.H. Foster. An operational approach to colour constancy. *Vision Research*, **32**:1359-1366, 1992.
4. M. D'Zmura. Color constancy: surface color from changing illumination. *J. Opt. Soc. Am. A*, **9**:491-493, 1992.
5. M. D'Zmura and G. Iverson. Color constancy. i. basic theory of two-stage linear recovery of spectral descriptions for lights and surfaces. *J. Opt. Soc. Am. A*, **10**:2148-2165, 1993.
6. G.D. Finlayson, M.S. Drew, and B.V. Funt. Enhancing von Kries adaptation via sensor transformations. In *Human Vision, Visual Processing, and Digital Display IV. SPIE*, 1-4 Feb. 1993.
7. D. Forsyth. A novel algorithm for color constancy. *Int. J. Comput. Vision*, **5**:5-36, 1990.
8. B.V. Funt and M.S. Drew. Color constancy computation in near-Mondrian scenes using a finite dimensional linear model. In *Computer Vision ant Pattern Recognition Proceedings*, pages 544-549. *IEEE Computer Society*, June 1988.
9. B.V. Funt, M.S. Drew, and J. Ho. Color constancy from mutual reflection. *Int. J. Computer Vision*, **6**:5-24, 1991.
10. B.V. Funt and J. Ho. Color from black and white. In *Proceedings of the Second International Conference on Computer Vision*, Tarpon Springs Dec 5-8, pages 2-8. *IEEE*, 1988; and *Int. J. Computer Vision*, **3**:109-117, 1989.
11. R. Gershon, A.D. Jepson, and J.K. Tsotsos. From [r,g,b] to surface reflectance: Computing color constant descriptors in images. *Perception*, pages 755-758, 1988.
12. G.H. Golub and C.F. van Loan. Matriz Computations. *John Hopkins U. Press*, 1983.
13. A.C. Hurlbert. The Computation of Color (*PhD Thesis*). MIT Artificial Intelligence Laboratory, 1989.
14. D.B. Judd, D.L. MacAdam, and G. Wyszecki. Spectral distribution of typical daylight as a function of correlated color temperature. *J. Opt. Soc. Am.*, **54**:1031-1040, August 1964.
15. E.H. Land. The retinex theory of color vision. *Scientific American*, pages 108-129, 1977.
16. E.H. Land. Recent advances in retinex theory. *Vision Res.*, **26**:7-21, 1986.
17. E.H. Land and J.J. McCann. Lightness and retinex theory. *J. Opt. Soc. Amer.*, **61**:1-11, 1971.
18. L.T. Maloney. Computational Approaches to Color Constancy. *PhD Thesis*, Stanford University, 1985. Applied Psychology Lab.
19. L.T. Maloney. Evaluation of linear models of surface spectral reflectance with small numbers of parameters. *J. Opt. Soc. Am. A*, **3**:1673-1683, 1986.
20. L.T. Maloney and B.A. Wandell. Color constancy: a method for recovering surface spectral reflectance. *J. Opt. Soc. Am. A*, **3**:29-33, 1986.
21. D.H. Marimont and B.A. Wandell. "Linear models of surface and illuminant spectra", *J. Opt. Soc. Am. A*, **9**(11):1905-1913, 92.
22. Newhall, Nickerson, and D.B. Judd. Final report of the OSA subcommittee on the spacing of the munsell colors. *J. Opt. Soc. Am. A*, **33**:385-418, 1943.
23. S.A. Shafer. Using color to separate reflection components, *Color Res. Appl.*, **10**:210-218, 1985.
24. M. Tsukada and Y. Ohta. "An approach to color constancy using multiple images". *In Proc. Third International Conference on Computer Vision. IEEE Computer Society*, **December 1990**.
25. J. J Vos and P. L. Walraven. On the derevation of the foveal receptor primaries, *Vision Research*, **11**: 799-818, 1971.

published previously in SPIE, Vol. 2179, page 352