Color Constancy: Enhancing von Kries Adaptation via Sensor Transformations

Graham D. Finlayson, Mark S. Drew and Brian V. Funt
School of Computing Science
Simon Fraser University
Vancouver, B.C.
Canada V5A 1S6

ABSTRACT

Von Kries adaptation has long been considered a reasonable vehicle for color constancy. Since the color constancy performance attainable via the von Kries rule strongly depends on the spectral response characteristics of the human cones, we consider the possibility of enhancing von Kries performance by constructing new "sensors" as linear combinations of the fixed cone sensitivity functions.

We show that if surface reflectances are well-modeled by 3 basis functions and illuminants by 2 basis functions then there exists a set of new sensors for which von Kries adaptation can yield perfect color constancy. These new sensors can (like the cones) be described as long-, medium-, and short-wave sensitive; however, both the new long- and medium-wave sensors have sharpened sensitivities—their support is more concentrated. The new short-wave sensor remains relatively unchanged. A similar sharpening of cone sensitivities has previously been observed in test and field spectral sensitivities measured for the human eye.

We present simulation results demonstrating improved von Kries performance using the new sensors even when the restrictions on the illumination and reflectance are relaxed.

1 INTRODUCTION

To a human observer, the world is populated with objects whose color does not change with change in the incident illumination. Moving from indoor tungsten illumination to outdoor blue sky light has only a small effect on our perception of color even though the light spectrum entering our eye differs markedly in the two situations. The ability to discount the effect of the illuminant and thereby perceive only surface properties is called *color constancy.*

Light entering the eye is sampled by long-, medium- and short-wave sensitive cone sensors. Consequently these cone response functions are at the heart of our color perception. Color constancy requires that the values registered by each cone, a 3-vector $\underline{p}$, be transformed into an illuminant-invariant surface descriptor $\underline{d}$. This transform is usually considered linear—a matrix is applied to cone response vectors. Indeed, under Forsyth's formulation[7] of the color constancy problem, the transform *must* be linear. The color constancy problem in equation form is:

$$\underline{d} = \mathcal{Q}\underline{p} \qquad \text{(1 : color constancy equation)}$$

where $\mathcal{Q}$ is the color constancy transform to be determined. Many different computational schemes have been proposed for solving for $\mathcal{Q}$. Each places structural constraints on the form of $\mathcal{Q}$, and further, usually

0-8194-1146-9/93/$6.00

requires that the world satisfy certain other strong constraints (e.g., that each scene must contain a particular distribution of surface reflectances[12]). In studying color constancy algorithms, therefore, we must ask two questions:

1. Independent of the computational scheme for computing the matrix, how well in principle can a particular matrix form discount the effect of the illuminant?

2. How successful is a given color constancy algorithm in solving for the correct (or best) transform?

This paper focuses on the first of these two questions. We show that if illuminants and surface reflectances are well described by finite-dimensional models with 2 and 3 degrees of freedom respectively (henceforth called the *2-3 case*) then a *generalized* diagonal matrix form supports perfect color constancy:

$$\underline{d} = \mathcal{T}^{-1}\mathcal{D}\mathcal{T}\underline{p} \qquad \text{(2 : generalized diagonal matrix color constancy)}$$

The meaning of (2) is that color constancy is achieved by applying a diagonal matrix $\mathcal{D}$ under a basis change $\mathcal{T}$. Of course the operation of the diagonal matrix can be made explicit by premultiplying both sides of equation (2) with $\mathcal{T}$:

$$\mathcal{T}\underline{d} = \mathcal{D}\mathcal{T}\underline{p} \tag{3}$$

The matrix $\mathcal{T}$ effectively creates new sensors by taking a linear combination of the cone response functions. These new sensors are again best described as long-, medium-, and short-wave sensitive; however, both the new long- and medium-wave sensors have *sharpened* sensitivities with their support more concentrated than before. As it happens, the new short-wave sensor remains relatively unchanged. The cone fundamentals measured by Vos and Walraven are contrasted with their sharpened counterparts in Figure 1. A similar sharpening of cone sensitivities has been measured in test and field spectral sensitivity experiments[8,15,20], and more recently in color discrimination experiments[19]. Finlayson et al.[5] have investigated how sharpened sensors better supports diagonal matrix color constancy.

Our work serves to strengthen existing diagonal matrix theories of color constancy; these include von Kries adaptation[22], Land's retinex[17], Horn[13] and Blake's[1] Lightness algorithms and Forsyth's[7] CRULE algorithm. Moreover we will show that the non-diagonal theories of Lennie and D'Zmura[4], Buchsbaum[2] and Gershon[12], when weakly restricted, reduce to diagonal counterparts under a sensor transformation. Our work also challenges the generally accepted belief that color constancy can only be achieved via more complex matrix forms.

In Section 2 we provide the necessary definitions required to develop a mathematical model to describe the color response of the human eye and go on to state formally the color constancy problem. In Section 3 we develop techniques for finding the sensor transform $\mathcal{T}$ which affords perfect diagonal matrix color constancy under the *2-3* constraints. It should be noted that this analysis is quite general and does not place restrictions on the possible form of the initial set of sensors. Obviously, the human cone sensitivities are of interest in the study of human color constancy; nevertheless, the analysis presented applies in a straightforward manner to the type of RGB filters found in color cameras used in machine vision.

In Section 4 we discuss our results in the context of existing computational theories of color constancy. In Section 5 we present simulation results which evaluate the performance of generalized von Kries adaptation for the case of real surface reflectances and illuminant spectra. Finally in Section 6 we present a brief review of the psychophysical evidence for sharpened cone sensitivities.

2 THE MODEL

The light reflected from a surface depends not only on the spectral properties of illumination and surface reflectance but also on other confounding factors such as specularities and mutual illumination. To simplify our analysis we will, in line with many other authors, develop our method for the simplified Mondrian world; a Mondrian is a planar, matte surface with several different, uniformly colored patches. We assume that the light striking the Mondrian is of uniform intensity and is spectrally unchanging. In this world the only factor confounding the retrieval of surface descriptors is illumination.

Light reflected from a Mondrian falls onto the retina where, for simplicity, we assume that at each location X all 3 classes of cone sensors are present. The value registered by the kth cone, p_k^X (a scalar), is equal to the integral of its response function multiplied by the incoming color signal. The color signal is the product of the illuminant spectral power distribution and the surface spectral reflectance function. For convenience, we arrange the index X such that each p_k^X corresponds to a unique surface reflectance:

$$\underline{p}_k^X = \int_\omega C^X(\lambda)\, R_k(\lambda)\, d\lambda \qquad \text{(4: color response)}$$

where λ is wavelength, $R_k(\lambda)$ is the response function of the kth sensor, $C^X(\lambda)$ is the color signal at X and the integral is taken over the visible spectrum ω. The color signal is the product of a single surface reflectance $S(\lambda)$ multiplied by the ambient illumination $E(\lambda)$: $C(\lambda) = E(\lambda)S(\lambda)$.

2.1 Finite-Dimensional Models

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:

$$S(\lambda) \approx \sum_{i=1}^{d_S} S_i(\lambda)\sigma_i \qquad (5)$$

where $S_i(\lambda)$ is a basis function and $\underline{\sigma}$ is a d_S-component column vector of weights. Maloney[18] presents evidence which suggests surface reflectances can be well modeled by a set of between 3 and 6 basis vectors. Similarly we can model illuminants with a low–dimension basis set:

$$E(\lambda) \approx \sum_{j=1}^{d_E} E_j(\lambda)\epsilon_j \qquad (6)$$

where $E_j(\lambda)$ is a basis function and $\underline{\epsilon}$ is a d_E dimensional vector of weights. Judd[14] measured 605 daylight illuminants and showed that overall they are well modeled by a set of 3 basis functions.

2.2 Lighting Matrices

Given finite-dimensional approximations to surface reflectance, the color response eqn. (4) can be rewritten as a matrix transform. A *lighting matrix* $\Lambda(\underline{\epsilon})$ maps reflectances, defined by the $\underline{\sigma}$ vector, onto a corresponding color response vector:

$$\underline{p} = \Lambda(\underline{\epsilon})\underline{\sigma} \tag{7}$$

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. (6).

2.3 The Color Constancy Problem

Within the framework of eqn. (1), any color constancy algorithm must aim to transform the color response vector $\underline{p}$, due to the surface $\underline{\sigma}$, to its corresponding illuminant independent descriptor $\underline{d}$; however, there is no consistent definition for a descriptor. Under von Kries adaptation[22] for example, the color response adapted relative to the response of a white reflectance is used as a descriptor. Specifically the adapted response in the long-wave cone channel is equal to the long-wave response of the surface $\underline{\sigma}$ divided by the long-wave response of the white reflectance $\underline{\sigma}_w$ (the medium- and short-wave adapted responses are similarly defined). The *simple* von Kries descriptor is formally stated in eqn. (8); note that the function *diag* converts a vector into a diagonal matrix, the ith row of the vector is mapped to the ith diagonal component of the matrix. Forsyth, in contrast, defines a descriptor to be the color response of a surface seen under a canonical illuminant, defined by the weight vector $\underline{c}$ (eqn. (9)).

$$\underline{d}^V = diag(\Lambda(\underline{\epsilon})\underline{\sigma}_w)^{-1}\Lambda(\underline{\epsilon})\underline{\sigma} \quad \text{(8 von Kries descriptor)}$$

$$\underline{d}^F = \Lambda(\underline{c})[\Lambda(\underline{\epsilon})]^{-1}\Lambda(\underline{\epsilon})\underline{\sigma} \quad \text{(9 Forsyth's descriptor)}$$

Since both descriptors are a linear transform from the cone response vector, they are themselves related by a linear transform. If, for all illuminants $\Lambda(\underline{c})[\Lambda(\underline{\epsilon})]^{-1}$ is a diagonal matrix then this implies that von Kries adaptation will afford perfect color constancy. In Section 3 we show, assuming *2-3* world conditions, that there exists a transformation of the cone basis such that $\Lambda(\underline{c})[\Lambda(\underline{\epsilon})]^{-1}$ is always diagonal.

3 DIAGONAL MATRIX TRANSFORM AND THE *2-3* CASE

In this section we show that if illumination and surface reflectance are described perfectly by finite dimensional models of 2 and 3 degrees of freedom respectively then there exists a transformed sensor basis in which a diagonal matrix supports perfect color constancy.

Under the *2-3* restrictions, every illuminant is described by a 2–vector with components $(\epsilon_1, \epsilon_2)^T$. Consequently, the color response of a surface characterized by $\underline{\sigma}$ can be written with respect to the two Λ matrices associated with the two basis directions in ϵ–space,

$$\Lambda(1) \leftrightarrow (1,0)^T\,, \quad \Lambda(2) \leftrightarrow (0,1)^T\,, \tag{10}$$

$$\Lambda(\underline{\epsilon}) = \epsilon_1\Lambda(1) + \epsilon_2\Lambda(2) \tag{11}$$

Therefore the color response of $\underline{\sigma}$ (eqn. (7)) becomes

$$\underline{p} = \epsilon_1 \Lambda(1)\underline{\sigma} + \epsilon_2 \Lambda(2)\underline{\sigma} \quad (12)$$

Let us define our canonical illuminant, $\underline{c}$, to be the first illuminant basis function. Now consider the relationship between the color response of a surface seen under illuminant $\underline{c}$ with that for the same surface illuminated by the second illuminant basis function. It is immediately clear that the color response under the second illuminant basis function is a linear transform, $\mathcal{M}$, away from its response vector viewed with respect to the canonical illuminant:

$$\Lambda(2)\underline{\sigma} = \mathcal{M}\Lambda(1)\underline{\sigma} \quad (13)$$

$$\mathcal{M} = \Lambda(2)[\Lambda(1)]^{-1} \quad (14)$$

Now we can rewriting eqn. (12):

$$\underline{p} = (\epsilon_1 \mathcal{I} + \epsilon_2 \mathcal{M})\Lambda(1)\underline{\sigma} \quad (15)$$

where $\mathcal{I}$ is the identity matrix.

Therefore we have shown that the color response of any surface viewed under an arbitrary illuminant is a fixed linear transform from its response with respect to the canonical illuminant. Furthermore, this linear transform is necessarily a linear combination of the identity matrix $\mathcal{I}$ and the matrix $\mathcal{M}$.

We defined a *generalized diagonal transform* (eqn. (2)) as a diagonal matrix applied with respect to a transformed sensor basis. That there exists a generalized diagonal transform mapping surface color responses between illuminants now follows from the eigenvector decomposition of $\mathcal{M}$. Suppose $\mathcal{M}$ has the (unique) decomposition

$$\mathcal{M} = \mathcal{T}^{-1}\mathcal{D}\mathcal{T} \ . \quad (16)$$

We can also express the identity matrix $\mathcal{I}$ in terms of the eigenvectors of $\mathcal{M}$:

$$\mathcal{I} = \mathcal{T}^{-1}\mathcal{I}\mathcal{T} \quad (17)$$

Consequently we can rewrite eqn. (12) as a generalized diagonal matrix transform:

$$\mathcal{T}\underline{p} = (\epsilon_1 \mathcal{I} + \epsilon_2 \mathcal{D})\mathcal{T}\Lambda(1)\underline{\sigma} \quad (18)$$

and hence

$$(\epsilon_1 \mathcal{I} + \epsilon_2 \mathcal{D})^{-1}\mathcal{T}\underline{p} = \mathcal{T}\Lambda(1)\underline{\sigma} \quad (19)$$

Using Forsyth's definition of a descriptor (eqn. (9)), the color response of a surface under a canonical illuminant, equation (19) is a statement of diagonal matrix color constancy. All surfaces described by the 3-dimensional reflectance set seen under an arbitrary illuminant (falling in the span of the 2-dimensional illuminant space) can be mapped to the canonical illuminant through the application of a diagonal matrix. An implication of equation (19) is that in the *2-3* world all color constancy transforms of the form $\Lambda(c)[\Lambda(e)]^{-1}$ have the same eigenvectors.

Of course, depending on the spectral characteristics of our illuminant and reflectance basis functions, the matrix $\mathcal{M}$ may have complex eigenvalues and eigenvectors. However as described in Finlayson et. al[6], generalized diagonal matrix color constancy holds equally well even when the sensor transformation is complex.

	CIE A	D48	D55	D65	D75	D100
3-parameter	0.04	0.3	0.11	0.23	0.16	0.22
2-parameter	4.8	5.4	4.1	1.4	1.1	5.7

Table 1: Percentage error of best fitting 2- and 3-parameter spectra to 6 test illuminants

4 IMPLICATIONS FOR COMPUTATIONAL THEORIES OF COLOR CONSTANCY

Many computational theories of color constancy assume that illumination and surface reflectance are well described by finite-dimensional models with few parameters. Commonly both reflectance and illumination are restricted to be 3-dimensional. This is a pragmatic choice since data analyses have demonstrated that 3-parameter models provide reasonable approximations to real illuminant and real reflectance spectra. Moreover restricting reflectances to a 3-parameter model is necessary if metamerism is to be avoided and color constancy rendered soluble.

One of the supposed advantages of color constancy algorithms founded on finite-dimensional models is that the color constancy transform is not constrained to be diagonal. The reasoning is that an independent adjustment of each sensor channel (i.e., the application of a diagonal matrix) is too simple an operation to discount the illuminant effectively. For example, Lennie and D'Zmura[4], assuming 3-parameter models of illumination and reflectance, state that "Because any reasonable basis functions for illuminants and reflectances are spectrally broadband, the integrals that describe the link between reflectance and quantum catches are typically non-zero [and hence] the three numbers used in scaling cone signals cannot undo what it takes nine numbers to describe".

Our analysis of Section 3 directly contradicts this viewpoint, however. Restricting illumination to be 2-dimensional implies that, under an appropriate transformation of the sensor basis, *perfect color constancy is afforded by a diagonal matrix.*. As long as illumination is restricted to a 2-dimensional space, then Lennie and Dzmura's algorithm is simply a diagonal matrix theory of color constancy. Similarly the computational methods discussed by Buchsbaum[2] and Gershon[12] both assume 3-parameter models of illumination and reflectance. Again, restricting illumination to a 2-dimensional space implies that these approaches also reduce to diagonal matrix color constancy.

Assuming that good color constancy is theoretically possible given a 3-parameter model of illumination, we must ask how good it still will be if we drop to a 2-parameter model. To answer this question, we performed a simulation experiment with the following 6 illuminants: 5 daylight phases (correlated color temperatures ranging from 4800 to 10000K) measured by Judd[14] and CIE standard illuminant A[23] (color temperature 2856K). A singular value decomposition of these spectra was used to derive optimal 3- and 2-parameter basis sets (optimal in the sense of minimizing the sum of squares of all the spectra to their best fitting 3- and 2- parameter approximations). In Table 1 we tabulate the percentage error between the real illuminant spectra and their closest fitting 2- and 3- parameter fits.

Clearly a 3-parameter model describes all 6 of our test illuminants well—all fitted errors are less than 0.25%. Also, reasonable fits are still obtained when we move to a 2-parameter model; errors ranging between 1.1 and 5.7%. In both cases the approximate spectra are close enough that we may venture to guess that a diagonal matrix may in fact support fairly good color constancy even when the 2-3 restrictions are relaxed.

This guess is tested in Section 5.

It is perhaps unsurprising that we can model the illuminants so well—after all we derived our 2- and 3-parameter models from a small set of only 6 illuminants. To more stringently test the adequacy of our models, we generated 13 Planckian black-body radiator spectra with color temperatures in the range 3200K to 10000K. The mean fitted error to the (previously derived) 3- and 2-parameter models are 5.2 and 9.3% respectively.

5 SIMULATIONS OF VON KRIES ADAPTATION

The *2-3* world is only an approximation of reality—reflectances are not precisely 3-dimensional nor are illuminants 2-dimensional—and as such a diagonal matrix affords only approximate color constancy. In this section we carry out simulations, using measured reflectances and measured illuminants, to evaluate the performance of *simple* and *generalized* von Kries adaptation—equations (8) and (20).

$$\underline{d}^V = \mathcal{T}^{-1} diag(\mathcal{T}\Lambda(\underline{\epsilon})\underline{\sigma}_w)^{-1}\mathcal{T}\Lambda(\underline{\epsilon})\underline{\sigma} \tag{20}$$

The color responses of surfaces viewed under different illuminants are generated using equation (4). The human cone fundamentals measured by Vos and Walraven[21] are used as our sensors, the 462 Munsell Spectra[3] for surfaces, and the 5 Judd Daylight phases[14] (D48, D55, D65, D75 and D100) and CIE A[23] for illuminants. All spectra are sampled at 10nm (nanometer) intervals from 400 to 650nm. Consequently the integral of eqn.(4) is approximated as a summation. The sensor transformation $\mathcal{T}$ was calculated via the technique outlined in Section 3. Illumination was modeled using the 2-dimensional basis described in Section 4. A 3-dimensional reflectance basis was derived by performing a singular value decomposition of the Munsell spectra. Figure 1 contrasts the Vos Walraven fundamentals with the transformed sensitivities.

We chose the Munsell spectra closest to uniform white as the white reference patch and the adapted responses calculated for D55 as *canonical* descriptors—these provide a reference for determining constancy performance. Under each of the remaining 5 *test* illuminants we calculated the adapted responses corresponding to each of the 462 reflectances. The Euclidean distance between these descriptors and their canonical counterparts, normalized with respect to each descriptor's length, provides a measurement of constancy performance. The percent normalized fitted distance (NFD) metric is defined as:

$$\mathrm{NFD} = 100 * \frac{\| \underline{d}^e - \underline{d}^c \|}{\| \underline{d}^c \|} \tag{20}$$

where c denotes the canonical (D55) illuminant and e a second arbitrary illuminant. For each of our 5 test illuminants we calculated the following 3 cumulative NFD histograms:

1. the NFD error for simple von Kries adaptation, eqn.(8).
2. the NFD error for generalized von Kries adaptation, eqn. (20).
3. the optimal color constancy performance for a general linear transform.

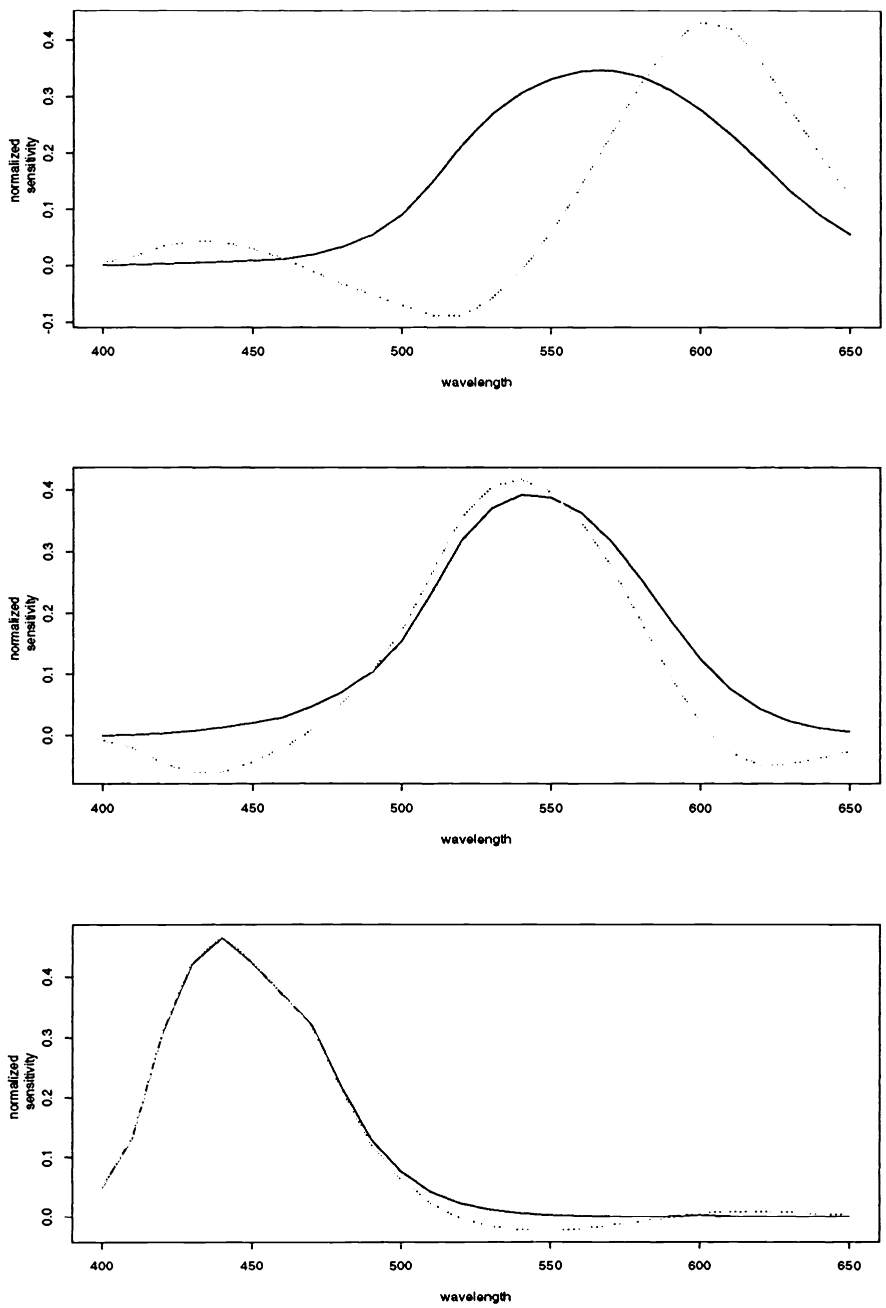

Figure 1: Result of sensor transformation $\mathcal{T}$. Solid lines: Vos–Walraven cone fundamentals; dashed lines: transformed sensors.

Optimal color constancy performance is defined to be a least-squares fit relating the responses of all surfaces viewed under an illuminant e to the corresponding responses with respect to the canonical illuminant c. This optimal case, serves as a control for evaluating the color constancy performance afforded by a diagonal matrix.

Figure 2 displays these 3 cumulative histograms for the test illuminants CIE A, D48, D65, D75 and D100. In all cases generalized von Kries adaptation outperforms, by a large margin, simple von Kries adaptation. Generalized von Kries adaptation also compares favorably with optimal color constancy. Only for the extremes in test illuminants, CIE A and D100, is there a significant performance difference.

6 PSYCHOPHYSICAL EVIDENCE

Sharpened sensitivities, similar to those shown in Figure 1, have been detected both in field- and test-sensitivity experiments (for a review of these terms see Foster[9]). Sperling and Harwerth[20] measured the test spectral sensitivities of human subjects conditioned to a large white background and, in correspondence with our transformed sensors, found sharpened peaks at 530nm and 610nm with no sharpening of the blue mechanism. Furthermore, consistent with our computational analysis, these authors propose that a linear combination of the cone responses accounts for the sharpening.

More recently, Foster[8] observed that field- and test-sensitivity spectra show sharpened peaks when derived in the presence of a small monochromatic auxiliary field coincident with the test field. Foster[11] extended this work by performing a hybrid experiment with a white, spatially-coincident auxiliary field; and sharpened sensitivities again were found. In both cases, these experimentally determined, sharpened sensitivities agree with our theoretical results. Like Sperling, Foster[10] verified that the sharpened sensitivities were a linear combination of the cone sensitivities.

Krastel[16] has measured spectral field sensitivities under changing illumination where, like Sperling, a white conditioning field is employed. Illumination color was changed by placing colored filters in front of the eye. The same test spectral sensitivity curve is measured under both a reddish and bluish illuminant. This suggests that the eye's sharpened mechanisms are unaffected by illumination. Kalloniatis[15] has measured cone spectral sensitivities under white adapting fields of different intensity and found the sharpened sensitivities to be independent of the intensity of the adapting field.

Poirson and Wandell[19] have developed techniques for measuring the spectral sensitivity of the eye with respect to the task of color discrimination. For color discrimination among briefly presented targets, the spectral sensitivity curve has relatively sharp peaks at 530nm and 610nm.

Although the general correspondence between our sharpened sensors and the above psychophysical results does not imply that sharpening in humans exists for the purpose of color constancy, at least the evidence that a linear combination of the cone responses is employed somewhere in the visual system lends plausibility to the idea that sharpening might be used in human color constancy processing.

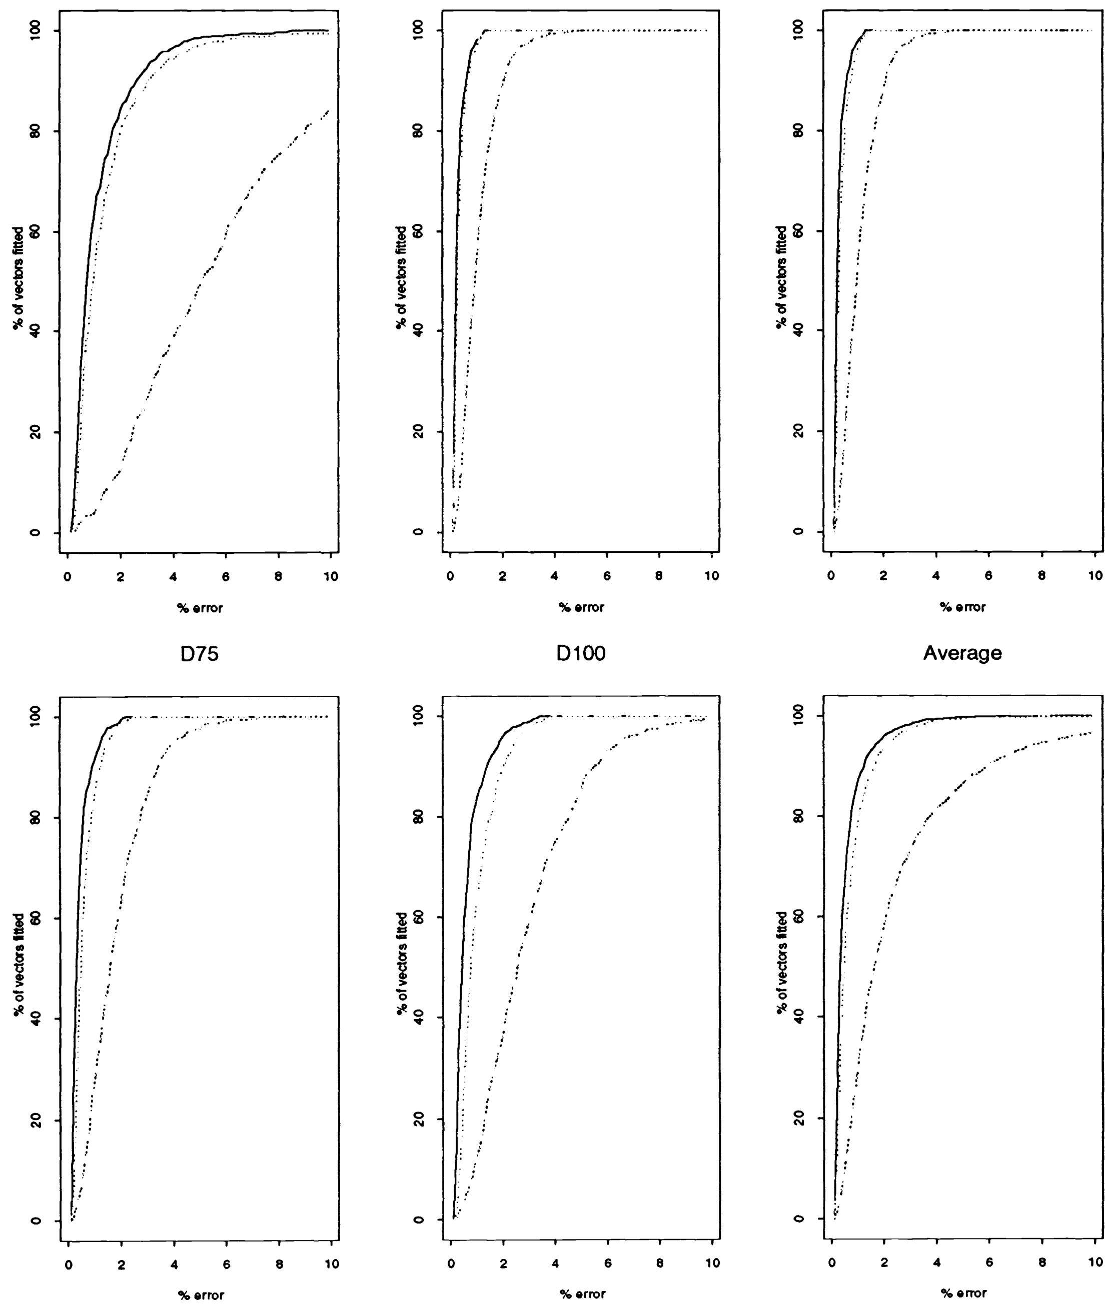

Figure 2: Cumulative histograms showing improved performance of generalized diagonal color constancy. Short dashed lines: simple von Kries adaptation; Long dashed lines: generalized von Kries adaptation; solid lines: optimal (non-diagonal) color constancy.

7 CONCLUSION

A diagonal matrix transform is at the foundation of many color constancy theories; these include von Kries adaptation, Land's retinex and Forsyth's CRULE. If illuminants are 2-dimensional and reflectances 3-dimensional, then we have shown that there exists a transformed cone basis in which a diagonal matrix will support perfect color constancy. Moreover, our analysis is quite general in the sense that subject to the dimensional restrictions on illuminants and reflectances, a diagonal matrix will suffice for color constancy for all trichromatic visual systems.

Relaxing these restrictions, we performed simulations of von Kries adaptation using actual measured illuminant and reflectance spectra. While the initial cone sensitivities afforded poor color constancy relative to the optimal possible. The transformed cone basis supported close to optimal color constancy. Indeed given the transformed basis, the performance is sufficiently good that there seems little reason to employ a more complex matrix form to support color constancy.

8 ACKNOWLEGEMENT

This work has been supported by the Natural Sciences and Engineering Research Council of Canada under grant #4322.

References

[1] A. Blake. Boundary conditions for lightness computation in Mondrian world. *Computer Vision, Graphics, and Image Processing*, 32:314–327, 1985.

[2] G. Buchsbaum. A spatial processor model for object colour perception. *Journal of the Franklin Institute*, 310:1–26, 1980.

[3] J. Cohen. Dependency of the spectral reflectance curves of the Munsell color chips. *Psychon. Sci.*, 1:369–370, 1964.

[4] M. D'Zmura and P. Lennie. Mechanisms of color constancy. *J. Opt. Soc. Am. A*, 3:1662–1672, 1986.

[5] G.D. Finlayson, M.S. Drew, and B.V. Funt. Spectral sharpening: an optimal sensor transformation for colour constancy. Technical Report CSS/LCCR TR 91-14, Simon Fraser University School of Computing Science, 1991.

[6] G.D. Finlayson, M.S. Drew, and B.V. Funt. Diagonal transforms suffice for color constancy. In *Proc. Fourth International Conference on Computer Vision*. IEEE Computer Society & European Vision Society, May 1992. To appear.

[7] D. Forsyth. A novel algorithm for color constancy. *Int. J. Comput. Vision*, 5:5–36, 1990.

[8] D.H. Foster. Changes in field spectral sensitivities of red-, green- and blue-sensitive colour mechanisms obtained on small backround fields. *Vision Research*, 21:1433–1455, 1981.

[9] D.H. Foster. Colour vision. *Contemporary Physics*, 25:477–497, 1984.

[10] D.H. Foster and R.S. Snelgar. Initial analysis of opponent-colour interactions revealed in sharpened field sensitivities. In J.D. Mollon and L.T. Sharpe, editors, *Colour Vision : Physiology and Psychophysics*, pages 303–312. Academic Press, 1983.

[11] D.H. Foster and R.S. Snelgar. Test and field spectral sensitivities of colour mechanisms obtained on small white backgrounds: action of unitary opponent-colour processes? *Vision Research*, 23:787–797, 1983.

[12] R. Gershon, A.D. Jepson, and J.K. Tsotsos. Ambient illumination and the determination of material changes. *J. Opt. Soc. Am. A*, 3:1700–1707, 1986.

[13] B. K. P. Horn. Determining lightness from an image. *Computer Vision, Graphics, and Image Processing*, 3:277–299, 1974.

[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] M. Kalloniatis and R.S. Harwerth. Spectral sensitivity and adaptation characteristics of cone mechanisms under white-light adaptation. *J. Opt. Soc. Am. A*, 7:1912–1928, 1990.

[16] H. Krastel, W. Jaeger, and S. Braun. An increment-threshold evaluation of mechanisms underlying colour constancy. In J.D. Mollon and L.T. Sharpe, editors, *Colour Vision : Physiology and Psychophysics*, pages 545–552. Academic Press, 1983.

[17] E.H. Land and J.J. McCann. Lightness and retinex theory. *J. Opt. Soc. Amer.*, 61:1–11, 1971.

[18] 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.

[19] A.B. Poirson and B.A. Wandell. Task-dependent color discrimination. *J. Opt. Soc. Am. A*, 7:776–782, 1990.

[20] H.G. Sperling and R.S. Harwerth. Red-green cone interactions in the increment-threshold spectral sensitivity of primates. *Science*, 172:180–184, 1971.

[21] J.J Vos and P.L. Walraven. On the derevation of the foveal receptor primaries. *Vision Research*, 11:799–818, 1971.

[22] G. West and M.H. Brill. Necessary and sufficient conditions for von kries chromatic adaption to give colour constancy. *J. Math. Biol.*, 15:249–258, 1982.

[23] G. Wyszecki and W.S. Stiles. *Color Science: Concepts and Methods, Quantitative Data and Formulas.* Wiley, New York, 2nd edition, 1982.