

Color constancy: generalized diagonal transforms suffice

Graham D. Finlayson, Mark S. Drew, and Brian V. Funt

Simon Fraser University, Vancouver, British Columbia V5A 1S6, Canada

Received July 22, 1993; revised manuscript received July 13, 1994; accepted July 13, 1994

This study's main result is to show that under the conditions imposed by the Maloney–Wandell color constancy algorithm, whereby illuminants are three dimensional and reflectances two dimensional (the 3–2 world), color constancy can be expressed in terms of a simple independent adjustment of the sensor responses (in other words, as a von Kries adaptation type of coefficient rule algorithm) as long as the sensor space is first transformed to a new basis. A consequence of this result is that any color constancy algorithm that makes 3–2 assumptions, such as the Maloney–Wandell subspace algorithm, Forsyth's MWEXT, and the Funt–Drew lightness algorithm, must effectively calculate a simple von Kries-type scaling of sensor responses, i.e., a diagonal matrix. Our results are strong in the sense that no constraint is placed on the initial spectral sensitivities of the sensors. In addition to purely theoretical arguments, we present results from simulations of von Kries-type color constancy in which the spectra of real illuminants and reflectances along with the human-cone-sensitivity functions are used. The simulations demonstrate that when the cone sensor space is transformed to its new basis in the appropriate manner a diagonal matrix supports nearly optimal color constancy.

Key words: color constancy, von Kries, chromatic adaptation, color balancing.

1. INTRODUCTION

We present a theoretical analysis connecting several color constancy theories (von Kries adaptation,^{1–3} the Land–McCann retinex,⁴ the Maloney–Wandell algorithm,⁵ the Funt–Drew lightness algorithm,⁶ and Forsyth's MWEXT and CRULE algorithms⁷) in which we prove that if illuminants and reflectances are well approximated by finite-dimensional models of low dimension then, under an appropriate change of basis for the sensor space, every one of these methods effectively calculates a simple independent adjustment of coefficients in this new space.

We caution the reader that we use the term von Kries adaptation in a somewhat more general sense than is customary. The term von Kries adaptation in the broad sense is taken to apply to any sensor basis and is not restricted solely to the cone-sensitivity functions. Specifically, von Kries adaptation, with respect to any sensor, is a simple scaling; each scaling component is the reciprocal of the sensor response induced from a reference patch (usually white). For example, if the sensor response for some surface reflectance is x and that of a reference patch is y , then the von Kries-adapted response is x/y .

Generally color cameras, like the human eye, are trichromatic; hence in a color image each pixel is a three-vector, one-component per sensor channel. A color constancy algorithm maps each color vector \mathbf{p} to a descriptor vector \mathbf{d} , which is independent of the illuminant. This mapping is usually considered linear: a matrix transform is applied to color vectors. Indeed, under Forsyth's formulation⁷ of the color constancy problem, the transform must be linear. In this paper we provide a theoretical analysis along with simulation results demonstrating that if the transform is linear then it need only be diagonal. In other words, a diagonal matrix transform suffices as a vehicle for color constancy. Our results are strong in the sense that they place no constraints on the initial spectral sensitivities of the visual system.

The various computational schemes for simulating color constancy apply different structural constraints to the form of the matrix transform. Many authors assume that the transform is a diagonal matrix, and in the Maloney–Wandell model⁵ the transform is a 2×3 projection. Only Forsyth's MWEXT⁷ algorithm places no constraints on the form of the transform. 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 the best) transform?

Our main focus in this paper is on the first of these questions.

From the von Kries adaptation model (see Ref. 8) through the Land–McCann retinex scheme⁴ to Forsyth's recent CRULE theory, the diagonal matrix transform has long been proposed as a feasible mechanism for color constancy. However, West and Brill¹ and D'Zmura and Lennie⁹ cast doubt on the suitability of diagonal matrix theories by demonstrating that for a given set of sensor sensitivities a diagonal matrix supports color constancy only for a restricted set of reflectance and illuminant spectra. With respect to the human-cone sensors, the restricted set of reflectance and illuminant spectra are statistically different from actual measured illuminants and reflectances. Consequently the majority of recent color constancy theories discard the computational simplicity of the diagonal matrix transform for more complex matrix forms that supposedly can model illuminant change better.

In contrast to this trend Finlayson *et al.*^{2,10} recently proved that diagonal matrix transforms can support perfect color constancy under small-dimensional model constraints whereby the illuminant space is linearly spanned

by a two-dimensional (2D) basis and the reflectance space is spanned by a three-dimensional (3D) basis. We term this set of constraints a 2–3 model. The analysis of Finlayson *et al.* employs a generalization, which we use here, of the concept of a diagonal matrix transform in which a fixed sensor transformation \mathcal{T} is permitted prior to the application of a diagonal matrix:

$$\mathbf{d} = \mathcal{D}\mathbf{p} \quad (\text{simple diagonal constancy}), \quad (1)$$

$$\mathcal{T}\mathbf{d} = \mathcal{D}\mathcal{T}\mathbf{p} \quad (\text{generalized diagonal constancy}). \quad (2)$$

In the 2–3 case, given a known reference patch in each scene, the correct diagonal matrix transform can be computed to yield perfect color constancy. The elegant color constancy algorithm of Maloney and Wandell⁵ does not require a reference patch, but it operates under a different set of restrictions. For a trichromatic visual system these restrictions, which we call the 3–2 restrictions, require a 3D illuminant space and a 2D reflectance space.

The main result of this paper is to show that, in a world in which illuminants and reflectances are governed by Maloney's 3–2 restrictions, color constancy can always be formulated as a generalized diagonal matrix transform independent of the spectral characteristics of the sensors. In a world in which these restrictions hold only approximately, a diagonal matrix transform theory of color constancy will still do a good job.

The ramifications of this result for theories of color constancy are widespread. The most immediate implication is that the 3–2 version of Maloney's theory of color constancy is effectively a diagonal-matrix-based theory of color constancy. Specifically, in the 3–2 world, the color vectors of all surfaces viewed under any illuminant will always be a diagonal matrix transform from the color vectors of the same surfaces viewed under a fixed canonical illuminant. Finite-dimensional restrictions are also at the foundation of the Funt–Drew⁶ color constancy algorithm. Their computational method simplifies, through our analysis, to diagonal matrix operations in the 3–2 case and as such reduces to Blake's version of the lightness algorithm.¹¹ Finally, our study plays a unifying role in connecting the theories of Maloney and Forsyth.

Forsyth's research on color constancy consists of two algorithms: MWEXT and the simpler CRULE. In MWEXT color constancy proceeds by parameterization of all the possible matrices that map the gamut of image colors into the gamut of descriptors. The more colorful the image, the smaller the set of possible mappings becomes. Unfortunately this algorithm is extraordinarily complex and, as Forsyth suggests, may not be suitable for machine vision. Restricting color constancy transforms to diagonal matrices results in Forsyth's simpler CRULE algorithm. This algorithm can be implemented efficiently and is a suitable candidate for a machine vision implementation of color constancy. Forsyth⁷ proposes his MWEXT algorithm to solve for color constancy under the 3–2 restrictions. Our results prove that his simpler CRULE algorithm is adequate for this task.

Generalized diagonal matrix transforms also relate to the problems of color correction and color balancing. White-point mapping, a common scheme for color balancing, is based on a von Kries-style adjustment of the three sensor channels aimed at making a white patch

in a scene appear white in an image. Simply adjusting white to look white does not guarantee, however, that the other colors will be correctly reproduced. By using a generalized diagonal matrix transformation instead of a simple diagonal matrix transformation one should obtain much better results.

In Section 2 we provide the definitions required for developing a mathematical model for color image formation and color constancy. In Section 3 we develop techniques for finding the sensor transform \mathcal{T} that affords perfect diagonal matrix color constancy under 3–2 restrictions. Note that this analysis does not place restrictions on the possible form of the initial set of sensors. In Section 4 we formally connect our results with other computational theories of color constancy. Finally, in Section 5 we present simulation results that evaluate the performance of generalized diagonal matrix color constancy. Appendix A discusses the role of complex numbers in theories of color constancy.

2. 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, in line with many other authors, we develop our theory for the simplified Mondrian world. A Mondrian is a planar surface composed of several overlapping matte (Lambertian) patches. We assume that the light that strikes the Mondrian is of uniform intensity and is spectrally unchanging. In the Mondrian world the only factor confounding the retrieval of surface descriptors is illumination.

Light reflected from a Mondrian falls onto a planar array of sensors, and at each location X in the sensor array there are three different classes of sensors. The value registered by the k th sensor p_k^X (a scalar) is equal to the integral of its response function multiplied by the incoming color signal. For convenience we arrange the index X such that each p_k^X corresponds to a unique surface reflectance:

$$p_k^X = \int_{\omega} C^X(\lambda)R_k(\lambda)d\lambda \quad (\text{color observation}), \quad (3)$$

where λ is the wavelength, $R_k(\lambda)$ is the response function of the k th 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 $E(\lambda)$: $C(\lambda) = E(\lambda)S(\lambda)$. Henceforth we drop the index X .

A. Finite-Dimensional Models

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

$$S(\lambda) \approx \sum_{i=1}^{d_S} S_i(\lambda)\sigma_i, \quad (4)$$

where $S_i(\lambda)$ is a basis function and σ is a d_S -component column vector of weights. Maloney¹² presents evidence that suggests that surface reflectances can be well

modeled by a set of between three and six basis vectors. Similarly, we can model illuminants with a low-dimension basis set:

$$E(\lambda) \approx \sum_{j=1}^{d_E} E_j(\lambda) \epsilon_j, \quad (5)$$

where $E_j(\lambda)$ is a basis function and ϵ is a d_E -dimensional vector of weights. Judd *et al.*¹³ measured 605 daylight illuminants and showed that they are well modeled by a set of three basis functions.

Basis functions are generally chosen by performance of a principal-component analysis of each data set (reflectances and illuminants) in isolation.^{14–16} This type of analysis is weak in the sense that it does not take into account how illuminant, reflectance, and sensor interact in forming a color vector [Eq. (3)]. Recently Marimont *et al.*¹⁷ developed a method for deriving reflectance and illuminant basis functions that best model color observations; Eq. (3) is the foundation of their method. They conclude that a 3D basis set for surface reflectance and a 3D basis set for illumination are sufficient to model the color observations of the 462 Munsell chips¹⁵ under a wide range of blackbody illuminants.

B. Lighting and Surface Matrices

Given finite-dimensional approximations to surface reflectance, a color observation equation (3) can be rewritten as a matrix transform. A lighting matrix $\Lambda(\epsilon)$ maps reflectances, defined by the σ vector, onto a corresponding color vector:

$$\mathbf{p} = \Lambda(\epsilon)\sigma, \quad (6)$$

where $\Lambda(\epsilon)_{ij} = \int_{\omega} R_i(\lambda)E_j(\lambda)S_j(\lambda)d\lambda$. The lighting matrix is dependent on the illuminant weighting vector ϵ , with $E(\lambda)$ being given by relation (5). The roles of illumination and reflectance are symmetric; we can write a color observation as a surface matrix transforming an ϵ vector:

$$\mathbf{p} = \Omega(\sigma)\epsilon, \quad (7)$$

where $\Omega(\sigma)_{ij} = \int_{\omega} R_i(\lambda)E_j(\lambda)S(\lambda)d\lambda$, with $S(\lambda)$ being defined in relation (4). This symmetry is a key part of the analysis presented in Section 3.

C. Color Constancy Problem

The aim of any color constancy algorithm is to transform the color observation vector \mathbf{p} to its corresponding illuminant-independent descriptor \mathbf{d} :

$$\mathbf{d} = \mathcal{Q}\mathbf{p}, \quad (8)$$

where \mathcal{Q} is a linear transform. However, there is no consistent definition for a descriptor. For example, Maloney and Wandell⁵ use the surface weight vector σ for the descriptor [Eq. (9)]; in contrast, Forsyth defines a descriptor to be the observation of a surface seen under a canonical illuminant, defined by the weight vector \mathbf{c} [Eq. (10)]:

$$\mathbf{d}^M = [\Lambda(\epsilon)]^{-1}\Lambda(\epsilon)\sigma \quad (\text{Maloney's descriptor}), \quad (9)$$

$$\mathbf{d}^F = \Lambda(\mathbf{c})[\Lambda(\epsilon)]^{-1}\Lambda(\epsilon)\sigma \quad (\text{Forsyth's descriptor}). \quad (10)$$

Because each color constancy algorithm applies a linear transform to color vectors, different descriptor definitions differ only by a fixed linear transform; for example, $\mathbf{d}^F = \Lambda(\mathbf{c})\mathbf{d}^M$. Therefore demonstrating the adequacy of a diagonal matrix for one descriptor form demonstrates its adequacy for color constancy in general. If a diagonal matrix is the vehicle for color constancy for any given descriptor, then we can say that color constancy is in general a diagonal matrix problem. As an example we show below that under the Maloney–Wandell 3–2 conditions a diagonal matrix is the vehicle for color constancy given by Forsyth descriptors. Thus, although the Maloney–Wandell algorithm does not explicitly calculate a diagonal matrix, we can say that by the equivalence of descriptor forms it is effectively a diagonal matrix theory of color constancy. In the analysis of Section 3 we use Forsyth's descriptor form.

D. Illuminant Invariance

Color constancy seeks illuminant-invariant color descriptors. A closely related problem is to find illuminant-invariant relationships between color vectors instead. One candidate relationship is the diagonal matrix mapping between the color vectors of the two surfaces:

$$\mathcal{D}^{ij}\mathbf{p}^{i,x} = \mathbf{p}^{j,x}. \quad (11)$$

Here i and j index two different surface reflectances, x indexes an illuminant, and \mathcal{D}^{ij} is a diagonal matrix. It is important to note that \mathcal{D}^{ij} denotes the entire 3×3 diagonal matrix relating $\mathbf{p}^{i,x}$ and $\mathbf{p}^{j,x}$, not the ij component of a matrix \mathcal{D} . We refer to Eq. (11) as diagonal invariance. Diagonal invariance is sometimes referred to as ratio invariance because the diagonal elements of \mathcal{D}^{ij} equal the ratios of the components of $\mathbf{p}^{j,x}$ over $\mathbf{p}^{i,x}$.

Diagonal invariance will be said to hold if for all illuminants x a fixed diagonal matrix \mathcal{D}^{ij} maps the color vector for surface i to the color vector for surface j . Diagonal invariance plays a key role in the lightness computations of Horn¹⁸ and Blake,¹¹ in the image-segmentation research of Hurlbert,¹⁹ and in the object-recognition research of Funt and Finlayson.²⁰ Brill²¹ develops a more general theory of illuminant invariance whereby the relationship between surfaces can be a general linear transform.

3. DIAGONAL TRANSFORMS AND THE 3–2 CASE

Finlayson *et al.*² proved that, assuming that illumination is 2D and reflectance 3D (the 2–3 case), there exists a transformed sensor basis in which a diagonal matrix supports perfect color constancy. In this section we prove the equivalent result for the 3–2 case.

Theorem 1. If illumination is 3D and surface reflectance 2D, then there exists a sensor transform \mathcal{T} for which a diagonal matrix supports perfect color constancy.

We prove theorem 1 in two stages. First we demonstrate a symmetry between diagonal invariance and diagonal matrix color constancy. Then we prove the existence of a sensor transform that supports diagonal invariance.

Lemma 1. A diagonal matrix supports perfect color constancy if and only if there is diagonal invariance.

Proof. When a diagonal matrix supports perfect color constancy, illumination change is exactly modeled by a diagonal matrix:

$$\mathbf{p}^{i,c} = \mathcal{D}^{ec} \mathbf{p}^{i,e}, \quad \mathbf{p}^{j,c} = \mathcal{D}^{ec} \mathbf{p}^{j,e} \quad (\text{same surfaces}), \quad (12)$$

where i, j index surface reflectance and the diagonal matrix \mathcal{D}^{ec} maps the observation of surfaces under an arbitrary illuminant e to their observation with respect to the canonical illuminant c . Clearly we can map $\mathbf{p}^{i,e}$ to $\mathbf{p}^{j,e}$ by applying a diagonal matrix:

$$\mathbf{p}^{i,e} = \mathcal{D}^{ij} \mathbf{p}^{j,e}. \quad (13)$$

Applying the color constancy transform \mathcal{D}^{ec} to both sides of Eq. (13), we can see that

$$\mathcal{D}^{ec} \mathbf{p}^{i,e} = \mathcal{D}^{ec} \mathcal{D}^{ij} \mathbf{p}^{j,e}. \quad (14)$$

Because transformation by diagonal matrices is commutative we can rewrite Eq. (14) as

$$\mathcal{D}^{ec} \mathbf{p}^{i,e} = \mathcal{D}^{ij} \mathcal{D}^{ec} \mathbf{p}^{j,e}. \quad (15)$$

Substituting Eqs. (12) into Eq. (15), we can see that

$$\mathbf{p}^{i,c} = \mathcal{D}^{ij} \mathbf{p}^{j,c}. \quad (16)$$

Equation (16) is a statement of diagonal invariance. The above argument is clearly symmetric: given diagonal invariance, diagonal matrix color constancy must follow. For the proof, we need only change the meaning of the superscripts in Eqs. (12)–(16) so that the first index denotes the illuminant and the second index the reflectance (\mathcal{D}^{ec} becomes a diagonal invariant and \mathcal{D}^{ij} a color constancy transform).

Lemma 2. Given 3–2 restrictions, there exists a transformation of the sensor response functions for which, independent of the illuminant, color vectors are diagonally invariant.

Proof. Under the 3–2 restrictions the color observation of a reflectance σ under an illuminant ϵ can be written in terms of two surface matrices. To see this, first note that matrix $\Omega(\sigma)$ in Eq. (7) can be decomposed into two parts. If the two-vector σ has components $(\sigma_1, \sigma_2)^T$ then, defining two special Ω matrices associated with the two basis directions in σ space,

$$\Omega(1) \leftrightarrow (1, 0)^T, \quad \Omega(2) \leftrightarrow (0, 1)^T,$$

we have

$$\Omega(\sigma) = \sigma_1 \Omega(1) + \sigma_2 \Omega(2).$$

Therefore Eq. (7) becomes

$$\mathbf{p} = \sigma_1 \Omega(1) \epsilon + \sigma_2 \Omega(2) \epsilon. \quad (17)$$

Let us define a canonical surface reflectance \mathbf{s} and examine its relationship to the color observation of other

surfaces. Without loss of generality we choose the first surface basis function as the canonical surface. The observation of the second surface basis function is an illuminant-independent linear transform from the color observation of the canonical surface:

$$\Omega(2) \epsilon = \mathcal{M} \Omega(1) \epsilon, \quad (18)$$

$$\mathcal{M} = \Omega(2) [\Omega(1)]^{-1}. \quad (19)$$

Now we can rewrite Eq. (17), the general observation of arbitrary surfaces, as a fixed transform from the observation of the canonical surface:

$$\mathbf{p} = [\sigma_1 I + \sigma_2 \mathcal{M}] \Omega(1) \epsilon, \quad (20)$$

where I is the identity matrix. Therefore we have shown that we can map the observation of the canonical surface to the observation of any other surface reflectance by applying a linear combination of the identity matrix I and the matrix \mathcal{M} . We define a generalized diagonal transform as a basis transformation followed by a diagonal matrix transform. The existence of a generalized diagonal transform that maps the observation of the canonical surface follows from the eigenvector decomposition of \mathcal{M} :

$$\mathcal{M} = \mathcal{T}^{-1} \mathcal{D} \mathcal{T}. \quad (21)$$

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

$$I = \mathcal{T}^{-1} I \mathcal{T}. \quad (22)$$

Consequently we can rewrite Eq. (17) as a generalized diagonal matrix transform:

$$\mathcal{T} \mathbf{p} = [\sigma_1 I + \sigma_2 \mathcal{D}] \mathcal{T} \Omega(1) \epsilon. \quad (23)$$

Equation (23) states that diagonal invariance holds between the canonical surface and all other surfaces given the fixed sensor transformation \mathcal{T} . In fact, Eq. (23) implies that diagonal invariance holds between any two surfaces. Let i and j index two arbitrary surfaces described by 2D vectors σ^i and σ^j . From Eq. (23), under any illuminant, we can write $\mathcal{T} \mathbf{p}^i$ and $\mathcal{T} \mathbf{p}^j$ as fixed diagonal transforms of $\mathcal{T} \mathbf{p}^s$ (the observation of the canonical surface):

$$\mathcal{T} \mathbf{p}^i = [\sigma_1^i I + \sigma_2^i \mathcal{D}] \mathcal{T} \mathbf{p}^s, \quad (24)$$

$$\mathcal{T} \mathbf{p}^j = [\sigma_1^j I + \sigma_2^j \mathcal{D}] \mathcal{T} \mathbf{p}^s. \quad (25)$$

Clearly we can write $\mathcal{T} \mathbf{p}^i$ as a diagonal matrix premultiplying $\mathcal{T} \mathbf{p}^j$:

$$\mathcal{T} \mathbf{p}^i = \mathcal{D}^{ij} \mathcal{T} \mathbf{p}^j, \quad (26)$$

where

$$\mathcal{D}^{ij} = [\sigma_1^i I + \sigma_2^i \mathcal{D}] [\sigma_1^j I + \sigma_2^j \mathcal{D}]^{-1}. \quad (27)$$

This step completes the proof of lemma 2. In the 3–2 case there exists a sensor transformation \mathcal{T} with respect to which there is diagonal invariance, and this invariance implies that a diagonal matrix is sufficient to support perfect color constancy (lemma 1). Therefore the proof of theorem 1 is also complete. ■

The crucial step in the above derivation is the eigenvector decomposition of the transform matrix \mathcal{M} . To re-

late this analysis to traditional theories of diagonal matrix color constancy we would like the eigenvalues of \mathcal{M} to be real valued. However, whether they are real valued depends on the form of the surface matrices (and hence the initial sensor spectral sensitivities).

On first consideration complex eigenvalues appear problematic; e.g., transforming the sensors by a complex matrix of eigenvectors does not have a plausible physical interpretation. The problem lies in the fact that the new sensors would be partly imaginary; however, we show in Appendix A that complex eigenvalues fit seamlessly into our generalized theory of diagonal matrix color constancy.

4. IMPLICATIONS FOR OTHER THEORIES OF COLOR CONSTANCY

Under the 3–2 conditions the lighting matrices $\Lambda(\epsilon)$ are 3×2 injective maps [i.e., color vectors are linear combinations of the two column vectors of $\Lambda(\epsilon)$], and surfaces seen under a single illuminant span a plane in the 3D receptor space. Maloney and Wandell⁵ exploit this plane constraint in their algorithm for color constancy. Maloney¹² proves that each illuminant corresponds to a unique plane of response vectors. Given uniqueness, Maloney and Wandell⁵ present an algorithm that can determine the illuminant weight vector ϵ and hence the pseudo-inverse $[\Lambda(\epsilon)]^{-1}$. Consequently the surface weight vector (or the Maloney descriptor) can be recovered with Eq. (9).

We present an alternative diagonal color constancy algorithm for the 3–2 world. We solve for color constancy in terms of Forsyth descriptors [Eq. (10)] and therefore explicitly solve for the diagonal matrix that maps the gamut of observed responses into the gamut of canonical responses.

A. Diagonal Color Constancy

In the 3–2 world the response vectors for surface reflectances under the canonical illuminant lie on the canonical plane \mathcal{P}^c (the canonical gamut). The span of the canonical plane is defined by the column vectors \mathbf{v}_1^c and \mathbf{v}_2^c of a 3×2 spanning matrix V^c and is calculated prior to the color constancy computation.

Under each of the other illuminants, response vectors for surfaces lie on the observed plane \mathcal{P}^o (the image gamut). Our goal is to solve for the diagonal matrix that maps \mathcal{P}^o to \mathcal{P}^c . If this mapping is unique, then there is a single solution to the color constancy problem. The diagonal mapping is unique if we make the following assumptions:

Assumption 1. There are two linearly independent surfaces in our image.

Assumption 2. None of the components of the surface normals \mathbf{n}^c , \mathbf{n}^o (of planes \mathcal{P}^c and \mathcal{P}^o) is equal to zero: $n_i^c \neq 0$ and $n_j^o \neq 0$ ($i, j = 1, 2, 3$).

If assumption 1 holds, then we can solve for the spanning matrix \mathcal{V}^o : the columns of \mathcal{V}^o , \mathbf{v}_1^o and \mathbf{v}_2^o , are simply observation vectors of any two distinct surfaces. Assumption 2 states that the normal of the canonical or the observed planes cannot lie on the x – y , the x – z , or the y – z plane. If all planes are equally likely, the probability that a plane normal has a zero component is vanishingly small. Maloney¹² sets out metameric black and unique light conditions to characterize when the

Maloney–Wandell⁵ subspace algorithm can solve the color constancy problem. These conditions are captured by assumption 2.

Theorem 2 (uniqueness theorem). Given assumptions 1 and 2, the diagonal transform that maps \mathcal{P}^o onto \mathcal{P}^c is unique.

Lemma 3. The only diagonal matrices that map the planes \mathcal{P}^c and \mathcal{P}^o onto themselves are the identity matrix I and the scalar multiples of the identity matrix γI .

Proof of Lemma 3. Let \mathcal{V} be a 3×2 matrix that defines the span of a plane (either \mathcal{P}^c or \mathcal{P}^o), with \mathbf{n} denoting its plane normal and \mathcal{D} being a diagonal matrix. Writing lemma 3 in mathematical notation, we would like to determine the conditions on \mathcal{D} for which

$$[\mathcal{D}\mathbf{n}]^t \mathcal{V} = [0 \ 0]. \tag{28}$$

Without loss of generality let us write the columns \mathbf{v}_1 and \mathbf{v}_2 of the spanning matrix \mathcal{V} in terms of the surface normal \mathbf{n} :

$$\mathbf{v}_1 = [n_2 \quad -n_1 \quad 0]^t, \tag{29}$$

$$\mathbf{v}_2 = [0 \quad n_3 \quad -n_2]^t. \tag{30}$$

It is easy to verify that \mathbf{v}_1 and \mathbf{v}_2 are linearly independent and orthogonal to \mathbf{n} . If \mathcal{D} has a zero component, where $\mathcal{D}_{ii} = 0$ for some i , then all vectors in the plane with spanning matrix $\mathcal{D}\mathcal{V}$ must have a zero as their i th component. In this case a vector that has a nonzero i th component and is zero elsewhere is normal to all vectors in the span of $\mathcal{D}\mathcal{V}$. If \mathcal{D} satisfies Eq. (28), then the matrices \mathcal{V} and $\mathcal{D}\mathcal{V}$ span the same plane, and the normal of \mathcal{P}^c or \mathcal{P}^o has a zero component. Under assumption 2 this is not the case. Consequently all components of the diagonal matrix \mathcal{D} are nonzero. Writing Eq. (28) in full, we have

$$\begin{aligned} \mathcal{D}_{11}n_1n_2 - \mathcal{D}_{22}n_2n_1 &= 0, \\ \mathcal{D}_{22}n_2n_3 - \mathcal{D}_{33}n_3n_2 &= 0. \end{aligned} \tag{31}$$

Because none of the variables (\mathcal{D}_{ii} , n_j , $j = 1, 2, 3$) is equal to zero, these equations are satisfied if and only if $\mathcal{D} = \gamma I$. ■

Proof of Theorem 2. Let us assume that there are two diagonal matrices, \mathcal{D}_1 and \mathcal{D}_2 , which differ by more than a simple scaling, mapping \mathcal{P}^o onto \mathcal{P}^c :

$$\mathcal{D}_1 \mathcal{V}^o = \mathcal{V}^c A_1, \tag{32}$$

$$\mathcal{D}_2 \mathcal{V}^o = \mathcal{V}^c A_2, \tag{33}$$

where A_1 and A_2 are 2×2 matrices (of full rank) that transform the span \mathcal{V}^c . Both the matrices \mathcal{D}_1 and \mathcal{D}_2 must have full rank; otherwise the normal to \mathcal{V}^c contains a zero component. Solving for \mathcal{V}^o in Eq. (32) and substituting the result into Eq. (33), we can see that

$$\mathcal{D}_2[\mathcal{D}_1]^{-1} \mathcal{V}^c A_1 = \mathcal{V}^c A_2. \tag{34}$$

According to lemma 3, only the identity matrix maps the canonical plane onto itself. Hence $\mathcal{D}_1 = \gamma \mathcal{D}_2$ (where

γ is a scalar), which contradicts our initial assumption; thus theorem 2 follows. ■

The truth of theorem 2 depends on whether assumption 2 holds. Let us assume, without loss of generality, that the first component of the surface normal is zero (i.e., $\mathbf{n} = [0 \ n_2 \ n_3]^t$, violating assumption 2). If the vector \mathbf{v} lies on the plane orthogonal to \mathbf{n} , then its dot product with \mathbf{v} equals zero: $v_2n_2 + v_3n_3 = 0$. The vector \mathbf{v} , transformed by a diagonal matrix with components $\mathcal{D}_{11} = a$ and $\mathcal{D}_{22} = \mathcal{D}_{33} = k$, is also orthogonal to \mathbf{n} : $kv_2n_2 + kv_3n_3 = 0$. In this case it follows that lemma 2 and theorem 2 are no longer true.

If the diagonal matrix that maps \mathcal{P}^o to \mathcal{P}^c is unique, it is easily determined. We can map the first spanning vector of \mathcal{V}^o , \mathbf{v}_1^o , onto \mathcal{P}^c by applying a linear combination of two diagonal matrices:

$$[\alpha \mathcal{D}^{11} + \beta \mathcal{D}^{12}] \mathbf{v}_1^o = \alpha \mathbf{v}_1^c + \beta \mathbf{v}_2^c. \quad (35)$$

Similarly, \mathbf{v}_2^o can be mapped onto \mathcal{P}^c by application of linear combinations of the diagonal matrices \mathcal{D}^{21} and \mathcal{D}^{22} . Because the diagonal matrix that maps \mathcal{P}^o to \mathcal{P}^c is unique, the set of diagonal matrices spanned by \mathcal{D}^{11} and \mathcal{D}^{12} must intersect those spanned by \mathcal{D}^{21} and \mathcal{D}^{22} in a unique diagonal matrix.

We can use this property to develop a simple algorithm, called S-CRULE (simplified CRULE), for color constancy. The algorithm requires two distinct colors in the image. It proceeds in three stages:

1. Find the set D_1 of diagonal matrices that map the first image color to the set of canonical colors.
2. Find the set D_2 of diagonal matrices that map the second image color to the set of canonical colors.
3. The unique diagonal matrix that maps all image colors to their observation under the canonical illuminant is equal to $D_1 \cap D_2$.

This algorithm is closely related to Forsyth's CRULE.⁷ One difference, however, is that through our analysis we can solve for the unique diagonal matrix by examining the color observations of two surfaces. In contrast, CRULE would examine all observed response vectors. If we relax the 3–2 model restrictions, both Maloney's algorithm and S-CRULE cannot solve the color constancy problem.

However, Forsyth's general CRULE can achieve color constancy even when the 3–2 conditions are relaxed. For example, if a vision system has ideal narrow-band sensors, then a diagonal matrix is always, without any restriction on illuminant or reflectance spectra, a perfect vehicle for color constancy.⁷ Therefore, for trichromatic color constancy, Maloney's theory is a subtheory of Forsyth's CRULE. Previously Forsyth had proposed that his more complex MWEXT algorithm would be required for solving for color constancy under 3–2 conditions.

B. Other Theories

The color constancy problem is made more difficult if the illuminant intensity varies across the image. Horn¹⁸ presented an algorithm for removing intensity gradients from images of a Mondrian world. Unfortunately his approach imposed strong constraints on the form of the Mondrian boundary. Later Blake¹¹ extended this algorithm to permit less restrictive boundary constraints.

Key to their algorithms is diagonal invariance and hence diagonal matrix color constancy. Therefore lightness recovery can be improved by the addition of a sensor transformation.

Funt and Drew⁶ presented a nondiagonal lightness algorithm for illuminants and reflectances that are well approximated by finite-dimensional models. Their method is independent of the sensor spectral sensitivities; however, we showed in Section 3 that diagonal invariance holds for arbitrary spectral sensitivity functions under an appropriate sensor transformation. Our analysis therefore circumvents the need for a nondiagonal lightness theory: the Funt–Drew algorithm reduces to Blake's algorithm under a sensor transformation in the case of 3–2 world conditions.

Land's retinex theory²² and its precursor, von Kries adaptation,¹ assume that color constancy is achieved if each image contains a known reference patch. Both require that a good approximation of diagonal invariance hold between observations of arbitrary surfaces. We have shown that perfect diagonal invariance holds for all sensor sets given 3–2 [and 2–3 (Refs. 2 and 10)] constraints.

Video cameras cannot account for changing illumination. Consequently images obtained under different illuminants must be balanced before being displayed to a human observer. This balancing usually takes the form of a simple scaling in each color channel (the color video image is transformed by a diagonal matrix). To ensure that illumination change is successfully corrected, video cameras are normally equipped with narrow-band sensors. The results given above indicate that a diagonal matrix transform is a suitable balancing technique independent of the sensor sensitivities used: broadband sensors are as suitable a choice as narrow-band sensors.

5. EXPERIMENTAL RESULTS

Real illuminants are not 3D, and real surfaces are not 2D (the 3–2 conditions provide only an approximation of actual color observations), and hence a diagonal matrix can achieve only approximate color constancy. Here we perform simulations, using measured surface reflectances and measured illuminants, that compare the performance of diagonal matrix and generalized diagonal matrix color constancy. We employ the analysis of Section 3 to derive the generalized diagonal matrix transform but test its efficacy with a von Kries-type algorithm. The von Kries algorithm has the advantage of not restricting response vectors to lying on a plane.

The color observations of surfaces viewed under different illuminants are generated with Eq. (3). We use the human-cone responses measured by Vos and Walraven²³ as our sensors, the 462 Munsell spectra¹⁵ for surfaces, and the 5 Judd daylight phases¹³ (D48, D55, D65, D75, and D100) and CIE A⁸ for illuminants. All spectra are sampled at 10-nm intervals from 400 to 650 nm. Consequently the integral of Eq. (3) is approximated as a summation.

The sensor transformation \mathcal{T} was calculated by the technique outlined in Section 3. Singular-value decompositions of the Munsell and the illuminant spectra were performed to derive the required surface and illuminant basis functions. Figure 1 displays the cone functions

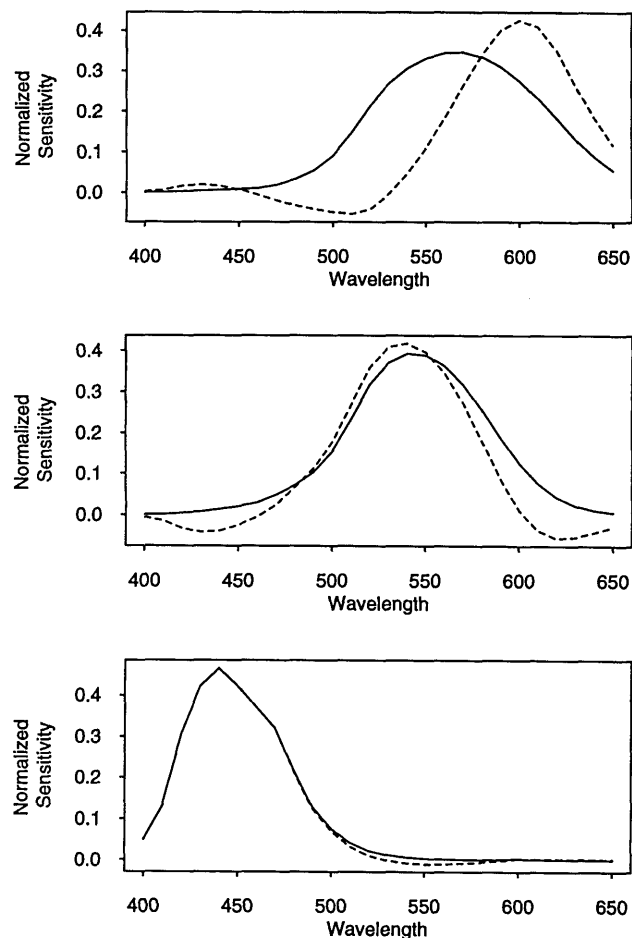


Fig. 1. Result of sensor transformation \mathcal{T} . Solid curves, Vos-Walraven cone fundamentals; dashed curves, transformed sensors.

before and after the fixed sensor transformation \mathcal{T} . Notice that the transformed sensors appear more narrow-band; this appearance is consistent with the pragmatic observation that narrow-band sensors afford better diagonal matrix color constancy. A similar narrowing was observed in various psychophysical experiments^{10,24-28} involving the human visual system.

In Fig. 2 we contrast the transformed sensors derived assuming 3-2 conditions with those assuming 2-3 conditions.² Both these sensor sets are remarkably similar to each other. This similarity is not altogether surprising: both the illuminant and the reflectance basis sets are statistically quite similar.

There are many algorithms for diagonal matrix color constancy; each differs in its strategy for determining the diagonal matrix. Here we present simulation results for von Kries adaptation,³ or white-patch normalization. The starting point for that algorithm is diagonal invariance. A color vector \mathbf{p}_i is assumed to be diagonally invariant to the observation of a white patch \mathbf{p}_w .

$$\mathbf{p}^i = \mathcal{D}^{iw} \mathbf{p}^w. \quad (36)$$

Hence it is the diagonal matrix \mathcal{D}^{iw} that is independent of the illuminant, and consequently this matrix can be used as a descriptor. Usually \mathcal{D}^{iw} is written in vector (or descriptor) form \mathbf{d}^{iw} , where $d_k^{iw} = p_k^i/p_k^w$. By the

symmetry between diagonal matrix color constancy and diagonal invariance we can rewrite Eq. (36) as a color constancy transform:

$$\mathbf{d}^{iw} = [\text{diag}(\mathbf{p}^w)]^{-1} \mathbf{p}^i, \quad (37)$$

where the function diag converts the vector \mathbf{p}^w to a diagonal matrix (diagonal elements correspond to the rows of \mathbf{p}^w). Arbitrarily we chose the white-patch descriptor vectors calculated for illuminant D55 as the canonical descriptor vectors; these provide a reference for determining color constancy performance. Under each of the other five illuminants we calculate white-patch descriptors. The Euclidean distance between these descriptors and their canonical counterparts, normalized with respect to the canonical descriptor's length, provides a measurement of constancy performance. The percent normalized distance (ND) metric is defined as

$$\text{ND} = 100 * \frac{\|\mathbf{d}^{iw,e} - \mathbf{d}^{iw,c}\|}{\|\mathbf{d}^{iw,c}\|}, \quad (38)$$

where $\mathbf{d}^{iw,c}$ denotes a canonical descriptor and $\mathbf{d}^{iw,e}$ denotes a descriptor for some other illuminant e . For each illuminant we calculated the three listed cumulative ND histograms:

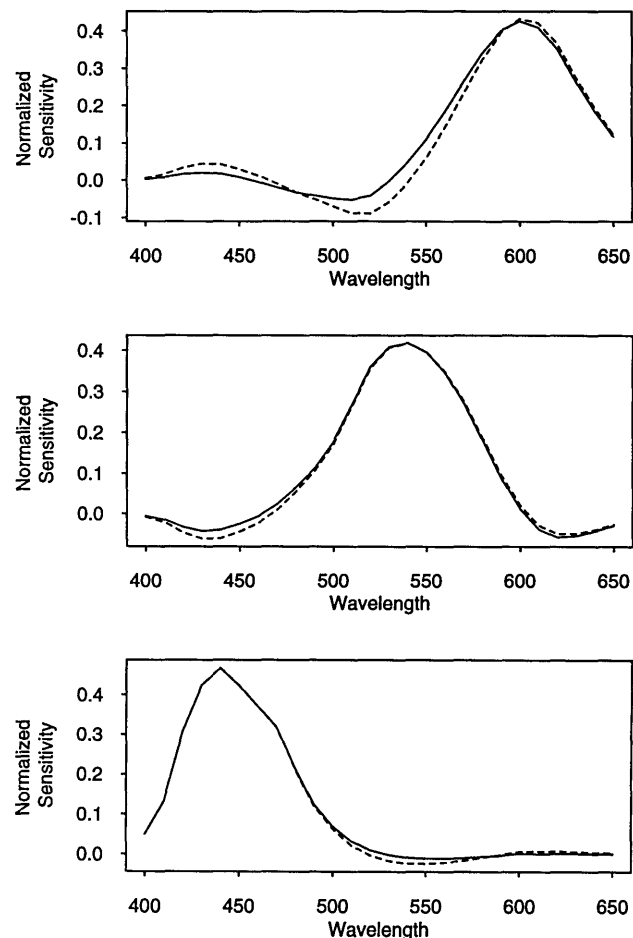


Fig. 2. Comparison of transformed sensors derived under 3-2 and 2-3 model assumptions. Solid curves, sensors derived assuming a 3-2 world; dashed curves, sensors derived assuming a 2-3 world.

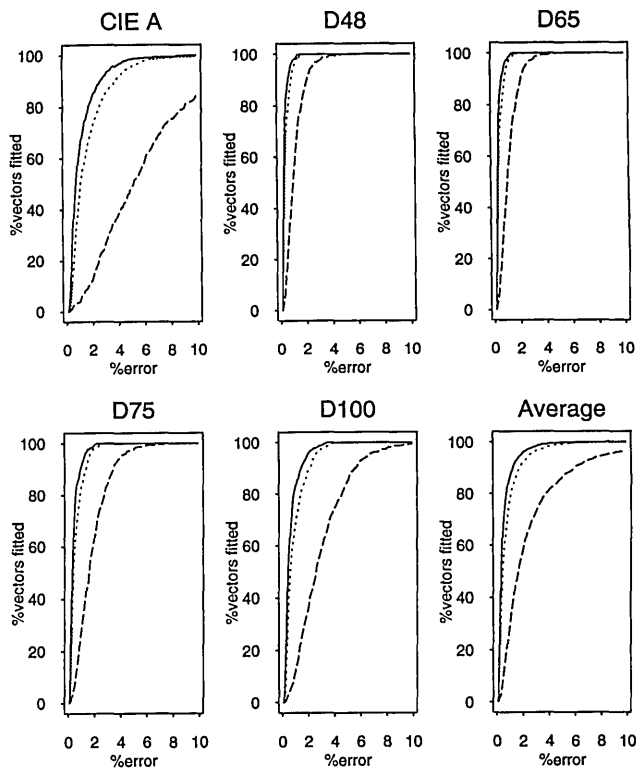


Fig. 3. Cumulative histograms showing improved performance of generalized diagonal color constancy. Dashed curves, simple diagonal color constancy; dotted curves, generalized diagonal color constancy; solid curves, optimal (nondiagonal) color constancy.

- (1) The ND error of white-patch-normalized responses for the cone functions;
- (2) The ND error of generalized white-patch-normalized responses, generalized in the sense of $\mathbf{d}^{i,w,e} = \mathcal{T}^{-1}[\text{diag}(\mathcal{T}\mathbf{p}^w)]^{-1}\mathcal{T}\mathbf{p}^{i,e}$ (we apply \mathcal{T}^{-1} after applying the diagonal matrix to ensure that all our comparisons are made with respect to the same sensor basis);
- (3) The optimal color constancy performance for a general linear transform.

We define optimal color constancy performance to be a least-squares fit that relates the observations of all surfaces under an illuminant e to their corresponding observations under the canonical illuminant c . This optimal case serves as a control for evaluating the color constancy performance afforded by a diagonal matrix.

Figure 3 displays these three cumulative histograms for the test illuminants CIE A, D48, D65, D75, and D100 (the dashed curves denote simple white-patch normalization, the dotted curves denote generalized white-patch normalization, and the solid curves denote the optimal constancy performance). In all cases generalized diagonal matrix color constancy outperforms, by a large margin, simple diagonal matrix constancy. Generalized diagonal matrix constancy also compares favorably with optimal color constancy. Only for the extremes in test illuminants (CIE A and to a lesser extent D100) is there a significant performance difference.

Vrhel and Trussell²⁹ considered the suitability of diagonal and nondiagonal matrices operating on cone responses as vehicles for color balancing. They concluded that, whereas a nondiagonal matrix performed well, a di-

agonal matrix (or a white-point mapping) was inadequate for color balancing. Our results refute this conclusion.

6. CONCLUSION

A diagonal matrix is the simplest possible vehicle for color constancy. Indeed, its inherent simplicity has motivated research into more complex matrix forms: if a diagonal matrix can give good color constancy, a nondiagonal matrix, which has nine instead of three parameters, must be able to support better color constancy, or so the reasoning goes. The analysis presented in this paper concludes that this is in fact not the case. Under the Maloney–Wandell world constraints a diagonal matrix, in conjunction with an appropriate fixed transformation of the sensor basis, has been shown to suffice for the support of perfect color constancy. This result is strong in the sense that no constraints are placed on the spectral sensitivities of the sensors.

Our simulation studies investigated whether the optimal sensors as expressed in the new sensor basis that was derived for the 3–2 world would continue to support good color constancy when the 3–2 restrictions were relaxed. For many real reflectances imaged under real illuminants a diagonal matrix continued to give nearly optimal color constancy.

Our analysis establishes a relationship among several theories of color constancy. For a world in which illumination is three dimensional and surface reflectance two dimensional the Maloney–Wandell⁵ algorithm, Forsyth's MWEXT,⁷ and the Funt–Drew lightness theory⁶ are all effectively diagonal matrix theories of color constancy because a diagonal matrix is always the vehicle for color constancy with respect to the Forsyth descriptor. Moreover, diagonal transforms are already at the heart of Forsyth's CRULE and von Kries adaptation. We contend that these nondiagonal algorithms are more complex than necessary and can be simplified by a fixed transformation of the sensor basis.

APPENDIX A: COMPLEX EIGENVALUES

Complex eigenvalues may arise in the eigenvector decomposition of the transform matrix \mathcal{M} , but, as we show below, they do not present a serious problem.

In traditional theories of diagonal matrix color constancy it is clear that each diagonal constancy transform can be expressed as the sum of three basis transforms. Indeed, it is this condition that makes diagonal matrix color constancy so appealing. For example, suppose that we observe the color vector \mathbf{p} and that this corresponds to the descriptor \mathbf{d} . This information is sufficient to solve for the constancy transform:

$$\mathbf{d} = \mathcal{D}\mathbf{p}, \quad \mathcal{D}_{kk} = \frac{d_k}{p_k}. \quad (\text{A1})$$

This same uniqueness condition is clearly true in generalized diagonal matrix color constancy if the sensor transformation \mathcal{T} is real valued. In fact, the uniqueness condition also holds in the general case in which the elements of \mathcal{T} can have complex terms.

Theorem 3. Under any sensor transformation \mathcal{T} (where \mathcal{T} can have complex elements) there are exactly

three linearly independent diagonal matrices consistent with generalized diagonal matrix color constancy. Consequently the mapping between a color vector and its descriptor is unique.

Proof. Our original statement of diagonal matrix color constancy, Eq. (2), can be written in the following mathematically equivalent form:

$$\mathbf{d} = \mathcal{T}^{-1} \mathcal{D} \mathcal{T} \mathbf{p}. \quad (\text{A2})$$

Both \mathbf{d} and \mathbf{p} are real-valued vectors, and hence $\mathcal{T}^{-1} \mathcal{D} \mathcal{T}$ must be a real-valued matrix. Theorem 2 follows if we can demonstrate that there exist only three linearly independent, real-valued matrices with the same eigenvectors, the columns of \mathcal{T}^{-1} .

A diagonal matrix \mathcal{D} has six variable components, three real and three imaginary numbers. Consequently there are in general six linearly independent matrices that share the same eigenvectors. The matrices $\mathcal{T}^{-1} I \mathcal{T}$, $\mathcal{T}^{-1} \mathcal{D} \mathcal{T}$, and $\mathcal{T}^{-1} \mathcal{D}^{-1} \mathcal{T}$ are linearly independent, real-valued matrices. Similarly, $\mathcal{T}^{-1} I j \mathcal{T}$, $\mathcal{T}^{-1} \mathcal{D} j \mathcal{T}$, and $\mathcal{T}^{-1} \mathcal{D}^{-1} j \mathcal{T}$ are all linearly independent, purely imaginary matrices (j is the square root of -1). The sum of imaginary numbers is always imaginary, and conversely the sum of real numbers is always real; hence these six matrices span the set of all matrices with eigenvectors \mathcal{T}^{-1} . Including complex numbers in the field over which we form a span, this means that only three matrices form a basis for the span of all real-valued matrices with eigenvectors \mathcal{T}^{-1} . This statement completes the proof for theorem 3. ■

Theorem 3 states that generalized diagonal matrix constancy holds equally well even when the sensor transformation is complex. For any sensor transformation the diagonal color constancy transform can be expressed as the sum of three diagonal basis matrices \mathcal{D} , \mathcal{D}^{-1} , and I . The mapping \mathcal{D}^{ij} , in Eq. (13), that takes $\mathbf{p}^{i,e}$ to $\mathbf{p}^{j,e}$ is still unique and is independent of the illuminant.

ACKNOWLEDGMENTS

The authors appreciate the helpful comments of Michael Brill. This research was supported by the Simon Fraser University Centre for Systems Science and the Natural Sciences and Engineering Research Council of Canada.

REFERENCES AND NOTES

- G. West and M. H. Brill, "Necessary and sufficient conditions for von Kries chromatic adaptation to give colour constancy," *J. Math. Biol.* **15**, 249–258 (1982).
- 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*, J. P. Allebach and B. E. Rogowitz, eds., Proc. Soc. Photo-Opt. Instrum. Eng. **1913**, 473–484 (1993).
- Adaptation made with linear combinations of the adapted cone functions is sometimes referred to as second-site adaptation (see, e.g., Ref. 9). This can be confusing, however, because second-site adaptation implies a second adaptation stage, whereas we use only a single adaptation stage, with the difference being that the adaptation is applied to sensors derived as linear transformations of the cone-sensitivity functions.
- E. H. Land and J. J. McCann, "Lightness and retinex theory," *J. Opt. Soc. Am.* **61**, 1–11 (1971).
- 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).
- B. V. Funt and M. S. Drew, "Color constancy computation in near-Mondrian scenes using a finite dimensional linear model," in *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition* (Institute of Electrical and Electronics Engineers, New York, 1988), pp. 544–549.
- D. Forsyth, "A novel algorithm for color constancy," *Int. J. Comput. Vision* **5**, 5–36 (1990).
- G. Wyszecki and W. S. Stiles, *Color Science: Concepts and Methods, Quantitative Data and Formulas*, 2nd ed. (Wiley, New York, 1982).
- M. D'Zmura and P. Lennie, "Mechanisms of color constancy," *J. Opt. Soc. Am. A* **3**, 1662–1672 (1986).
- G. D. Finlayson, M. S. Drew, and B. V. Funt, "Spectral sharpening: sensor transformations for improved color constancy," *J. Opt. Soc. Am. A* **11**, 1553–1563 (1994).
- A. Blake, "Boundary conditions for lightness computation in Mondrian world," *Comput. Vision Graphics Image Process.* **32**, 314–327 (1985).
- L. T. Maloney, "Computational approaches to color constancy," Ph.D. dissertation (Stanford University, Stanford, Calif., 1985).
- 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 (1964).
- J. Cohen, "Dependency of the spectral reflectance curves of the Munsell color chips," *Psychon. Sci.* **1**, 369–370 (1964).
- S. M. Newhall, D. Nickerson, and D. B. Judd, "Final report of the OSA subcommittee on the spacing of the Munsell colors," *J. Opt. Soc. Am.* **33**, 385–418 (1943).
- 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).
- D. H. Marimont, B. A. Wandell, and A. B. Poirson, "Linear models of surface and illuminant spectra," *J. Opt. Soc. Am. A* **9**, 1905–1913 (1992).
- B. K. P. Horn, "Determining lightness from an image," *Comput. Vision Graphics Image Process.* **3**, 277–299 (1974).
- A. Hurlbert, "Formal connections between lightness algorithms," *J. Opt. Soc. Am. A* **3**, 1684–1692 (1986).
- B. V. Funt and G. D. Finlayson, "Color constant color indexing," *IEEE Pattern Anal. Mach. Intell.* (to be published).
- M. H. Brill, "A device performing illuminant-invariant assessment of chromatic relations," *J. Theor. Biol.* **71**, 473–478 (1978).
- E. H. Land, "The retinex theory of color vision," *Sci. Am.* **237**(6), 108–129 (1977).
- J. J. Vos and P. L. Walraven, "On the derivation of the foveal receptor primaries," *Vision Res.* **11**, 799–818 (1971).
- D. H. Foster and R. S. Snelgar, "Initial analysis of opponent-colour interactions revealed in sharpened field sensitivities," in *Colour Vision: Physiology and Psychophysics*, J. D. Mollon and L. T. Sharpe, eds. (Academic, New York, 1983), pp. 303–312.
- W. Jaeger, H. Krastel, and S. Braun, "An increment-threshold evaluation of mechanisms underlying colour constancy," in *Colour Vision: Physiology and Psychophysics*, J. D. Mollon and L. T. Sharpe, eds. (Academic, New York, 1983), pp. 545–552.
- H. G. Sperling and R. S. Harwerth, "Red-green cone interactions in the increment-threshold spectral sensitivity of primates," *Science* **172**, 180–184 (1971).
- 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).
- A. B. Poirson and B. A. Wandell, "Task-dependent color discrimination," *J. Opt. Soc. Am. A* **7**, 776–782 (1990).
- M. J. Vrhel and H. J. Trussell, "Physical device illumination correction," in *Device-Independent Color Imaging and Imaging Systems Integration*, R. J. Motta and H. A. Berberian, eds., Proc. Soc. Photo-Opt. Instrum. Eng. **1909**, 84–91 (1993).