Investigating von Kries-like Adaptation Using Local Linear Models

G. D. Finlayson,[1] S. D. Hordley,[1]* A. Alsam[2]

[1] School of Computing Sciences, University of East Anglia, Norwich, NR4 7TJ, United Kingdom

[2] Norwegian Color Research Laboratory, Gjøvik University College, Gjøvik, Norway

Received 4 February 2004; revised 10 June 2005; accepted 18 August 2005

Abstract: In this article we explore the conditions under which a von Kries-like model of color constancy is exact. This investigation is motivated by the fact that in practice the model has been shown to work well for a wide range of imaging devices despite the fact that previous analysis of the necessary and sufficient conditions for the model to work predicts that it should perform poorly for real lights and surfaces. We present a modified theory that reconciles this apparent contradiction and that is based on the observation that von Kries-like adaptation treats sensor responses independently of one another. Starting from this point we show how to recover, for a single sensor, set of surfaces, and reference illuminant, the set of von Kries illuminants: all lights for which von Kries-like adaptation is a perfect model of illumination change. To help us in this task we use a linear model of surface reflectance, but importantly, we use a local model, that is, a model derived by looking at reflectance only in the region to which the sensor of interest is sensitive. Adopting such a model and treating sensors independently of one another, we show that our new theory accurately predicts the good practical performance of a von Kries-like model. Col Res Appl, 31, 90–101, 2006; DOI 10.1002/col.20186

Key words: color constancy; adaptation

INTRODUCTION

A visual system's response to light reflected from an object depends on the reflectance properties of the object but also on the nature of the light by which the object is illuminated. For example, a white object viewed indoors under reddish tungsten light will appear more red to a visual system than the same object viewed under bluesky daylight. This variation in color is in general an undesirable property of a visual system and much research (e.g., Refs. 1–4) has been devoted to deriving so-called color constancy algorithms that account for the color of the light incident on a scene and deliver stable object colors regardless of the prevailing illumination. Our experience of the world suggests that our own visual system might be able to solve just this problem since we perceive a white object as white regardless of the light under which we view it. However, careful studies[5–8] of this phenomenon are equivocal in their conclusions as to whether and to what extent we are color constant. Furthermore, the mechanism by which this ability might be achieved is also the subject of debate. Leaving our own visual system aside, researchers[1,2,4] in computer vision and digital photography are keen to design algorithms that render their visual systems color constant.

It is common to separate the color constancy problem into two stages: in the first stage the aim is to obtain an estimate of the scene illuminant from an image of the captured scene and in the second stage this illuminant estimate is used to correct the image such that the effect of the illuminant is discounted. This second stage usually amounts to adjusting (or adapting) sensor responses recorded under the scene light to corresponding responses recorded under a reference light. Before we can perform this adjustment we must specify the form that it will take—in mathematical terms we must specify the nature of the mapping between sensor responses under a pair of illuminants. The form this mapping takes is important too in the first stage of color constancy since it governs the type of illuminant estimate that must be obtained from the image data.

In this article we are concerned not with solving the color constancy problem per se but only with the problem of what form the adjustment of sensor responses should take. One of

---

* Correspondence to: S. D. Hordley (e-mail: steve@cmp.uea.ac.uk)

the simplest forms of sensor adjustments that has been proposed is to scale all the responses in a single channel by a fixed scalar. Thus, for a trichromatic device, sensor adjustment (or adaptation) is controlled by three independent scale factors applied separately to the responses of each class of sensor. To understand this model better let us assume that a visual system has three classes of sensor and let us denote the responses of these sensors to some surface viewed under an arbitrary light as (ρ_1, ρ_2, ρ_3). If we then denote the response of the system to the same surface but viewed under a second light by (ρ'_1, ρ'_2, ρ'_3), then the model tells us that the two sets of sensor responses are related by

$$\rho'_1 = \alpha\rho_1, \qquad \rho'_2 = \beta\rho_2, \qquad \rho'_3 = \gamma\rho_3, \tag{1}$$

where α, β, and γ are three scaling factors independent of one another and independent of the surface being viewed. That is, the scale factors depend only on the pair of illuminants in question and are fixed for any surface we might view under them. When the responses ρ_i are cone responses—the responses of light sensitive cells in the retina of the human eye—the model can correctly be called the von Kries model of adaptation. The same term is commonly used to describe the model for any arbitrary sensor response functions but in this article we use the term "von Kries-like model" to denote the model when applied to arbitrary sensor responses and reserve the term "von Kries model" for the case of cone responses. We will also sometimes use the term "diagonal model" to describe the model operating on an arbitrary set of sensors since mathematically the model takes the form of a 3×3 diagonal matrix multiplying sensor responses.

Whether a von Kries-like model is an accurate model of illumination change depends on the nature of the lights, surfaces, and sensors that are being modeled. In this article we set out to answer the following simple question: for a given sensor, what is the set of illuminants for which a von Kries-like model is a perfect model of illumination change? Previously, West and Brill[9] set out necessary and sufficient conditions for such a model to be exact with respect to a given trichromatic sensor set and thus the work in this article is closely related to that prior work. We begin by giving a reformulation of West and Brill's work and we show how to derive, with respect to a given set of sensors and surfaces, the corresponding set of illuminants for which a diagonal model of illumination change is perfect (in the sense that it maps, without error, sensor responses under a pair of illuminants). We show that this theory predicts that for the type of sensor, surface, and illuminant we encounter in the real world, a von Kries-like model will be a poor model of illumination change.

But the poor performance predicted by this analysis is at odds with our experience of using the model in practice. There is evidence[10] to support the theory that von Kries-like adaptation is an accurate model of color constancy both for our own visual system and for a range of other typical imaging devices. In the rest of the article, therefore, we set out to bridge the gap between West and Brill's pessimistic theory and the good performance that can be obtained in practice for a wide class of imaging sensors. To this end we present a new analysis of the conditions under which a von Kries-like model is perfect. Our analysis is based on the observation that scalings applied to the sensor channels in the model are independent of one another. We argue that it is this independence of sensors that is the key to understanding the good practical performance of the diagonal model. To this end we consider each sensor separately and we again set out to derive, for a given sensor, the set of illuminants for which a von Kries-like model is a perfect model of illumination change.

To help us derive this set we employ a linear model representation of surface reflectances as West and Brill did and as many authors[11–14] have done before. However, we introduce an important modification into the representation that is motivated by the fact that we are considering each sensor separately and independently of the others. When modeling reflectance it is usual to consider the region of the visible spectrum for which one or more of the device's sensors has significant response and to derive a set of basis vectors for this whole region. But since we are considering now only a single sensor, we argue that we need only model reflectance in the region of the spectrum that is "seen" by it: that is, in the region of the spectrum across which the individual sensor has significant response. In general, this will be quite a narrow subsection of the whole spectrum. Representing surfaces using a local linear model, localized to the region for which the sensor has significant response, we show how to derive the set of illuminants for which a diagonal model is a perfect model of illumination change.

An empirical evaluation of the new theory reveals two facts. First, if we allow different linear bases for each sensor then a lower dimensional basis suffices to model surfaces accurately than when we derive a model across the support of all three sensors. Second, the set of illuminants we derive based on these local linear models is much closer to actual real-world illuminants than the illuminants recovered using the analysis of West and Brill. Furthermore, we investigate the correlation between the actual errors that occur in mapping sensor responses using a diagonal model and the error between the theoretical illuminants we derive and the actual illuminant. We find better correlation of the errors when the illuminants are derived using our local analysis than when we use the analysis of West and Brill. Thus, we achieve our aim by deriving a theory that predicts the good practical performance of von Kries-like models of adaptation.

The rest of the article is organized as follows. In the next section we introduce a model of image formation and discuss some possible models for illumination change. The third section contains a reformulation of West and Brill's original work on the necessary and sufficient conditions for von Kries-like color constancy to be exact. We reanalyze this theory in the next section on a per-sensor basis and in the fifth section we evaluate the practical performance of both theories. Finally, we present a short conclusion to the article.

IMAGE FORMATION AND ILLUMINATION CHANGE

We adopt a simple model of image formation in which color is formed by the interaction of three different factors: light, surface, and sensor. A light is characterized by $e(\lambda)$, its spectral power distribution (SPD), which defines how much energy it emits at each wavelength. A surface is represented by $s(\lambda)$, its surface reflectance function defining what proportion of light incident upon it is reflected on a per-wavelength basis. Finally, the ith sensor of a device or visual system is denoted by its spectral sensitivity function $q_i(\lambda)$, which characterizes what fraction of the light incident upon it the sensor absorbs, again on a per-wavelength basis. The response of a sensor, which we denote ρ_i, can thus be expressed as

$$\rho_i = \int_{\omega} e(\lambda)s(\lambda)q_i(\lambda)d\lambda, \tag{2}$$

where ω represents the range of wavelengths for which the device has significant sensitivity. Our own visual system is trichromatic and most visual systems we encounter in computer vision or digital photography are trichromatic and thus we make the assumption in this article that at a primary level "color" is represented as a triplet of sensor responses, which we denote $\underline{\rho}$ (here, and throughout the article, an underline denotes a vector quantity).

In the context of this article we are interested in understanding what happens to these responses when the illuminant e changes and surface and sensor remain fixed. It is not obvious from Eq. (2) that a change in e will give rise to a simple change in ρ_i since the interaction among light, surface, and sensor is complex. In practice, however, quite simple mechanisms are able to model illumination change well and it is the main aim of this article to explain better why this is the case. To help in this task it is useful to first introduce a number of simplifying assumptions.

The first such simplification is to represent light, surface, and sensor, not by continuous functions but rather by their values at a set of n discrete sample points across the range of the visible spectrum. Such a representation is well justified by the Sampling Theorem[15] and a common and well-justified[16] sampling for most imaging systems is 31 samples at 10-nm intervals in the range of 400–700 nm. With such a representation the continuous integral in Eq. (2) is now approximated by a summation,

$$\rho_i = \sum_{k=1}^{n} e(\lambda_k)s(\lambda_k)q_i(\lambda_k)\Delta\lambda, \tag{3}$$

where the scalar $\Delta\lambda$ accounts for the size of the sampling interval. Henceforth, we represent the discrete functions of light, surface, and sensor by vectors E, S, and Q_i, whose elements represent the function values at each discrete sample point so that we can rewrite Eq. (3) in matrix form,

$$\rho_i = S^{\mathrm{T}}\mathrm{diag}(Q_i)E, \tag{4}$$

where T denotes the transpose operator, $\Delta\lambda$ has been subsumed into Q_i, and diag is the diagonal operator that transforms its $n \times 1$ vector argument into an $n \times n$ diagonal matrix, whose nonzero entries are the elements of the vector.

A further simplification that is often made is to represent lights, surfaces, or both by low-dimensional linear models. Such a simplification is justified by statistical analyses of large sets of typical daylights[17] and surfaces.[12,14] In this article we will make use of such a representation for surfaces so that any surface s is represented:

$$s(\lambda) = \sum_{k=1}^{M} \sigma_k S_k(\lambda). \tag{5}$$

That is, an arbitrary surface can be represented as a linear combination of a set of fixed basis functions S_k, $k = 1 \cdots M$ so that a surface is completely characterized by the M weights σ_k, where M is typically much smaller than n, the number of sample points used to represent lights and surfaces. The image formation equation now becomes

$$\rho_i = ([S_1 \cdots S_M]\underline{\sigma})^{\mathrm{T}}\mathrm{diag}(Q_i)E. \tag{6}$$

Such a representation is helpful since it can, for example, lead to a simplified model of illumination change. For example, it can be shown[4] that if surface reflectance is three-dimensional, then sensor responses to the same surface under two different lights are a linear transform apart,

$$\underline{\rho}^{E_1} = A\underline{\rho}^{E_2}, \tag{7}$$

where A is a 3 × 3 matrix.

Perhaps the most widely adopted model of illumination change simplifies this linear relationship even further and restricts the linear transform to be diagonal so that corresponding sensor responses under two different illuminants are related by

$$\begin{pmatrix} \rho_1^{E_1} \\ \rho_2^{E_1} \\ \rho_3^{E_1} \end{pmatrix} = \begin{pmatrix} \alpha & 0 & 0 \\ 0 & \beta & 0 \\ 0 & 0 & \gamma \end{pmatrix} \begin{pmatrix} \rho_1^{E_2} \\ \rho_2^{E_2} \\ \rho_3^{E_2} \end{pmatrix}. \tag{8}$$

This so-called diagonal or von Kries-like model of illumination change is exactly the model introduced in the introduction of this article. Von Kries[18] first proposed the model as a way of explaining the results of certain asymmetric matching experiments performed by him. It has since often been suggested that such an adaptation mechanism is (partly) responsible for the color constancy that our own visual system exhibits. We note, however, that there is evidence[19] to suggest that human color constancy is not achieved using only von Kries adaptation since cone signals are only available within the retina while color constancy has been shown to depend on binocular factors and therefore some processing must occur beyond the point at which von Kries adaptation is possible. Independently of our own visual system von Kries-like models are widely used in the machine vision literature.

In this article we are interested in understanding the conditions under which the model will be perfect: that is, for which class of sensors, surfaces, and illuminants is the model exact in the sense that mapping between sensor responses under a pair of illuminants by a diagonal transform is error free? It has been shown by Finlayson *et al.*[20,21] that the adequacy of the model depends on the sensors to which it is applied. Indeed, they have shown that a slightly more general form of the diagonal model provides an accurate model of illumination change for most sensor sets. In their model that they called the generalized coefficient model, illumination change is again modeled by a diagonal transform, but sensor responses are first transformed by a fixed 3 × 3 linear transform prior to the application of this transform. In this case responses under different illuminants are related,

$$T_s\underline{\rho}^{E_1} = DT_s\underline{\rho}^{E_2} \Rightarrow \underline{\rho}^{E_1} = T_s^{-1}DT_s\underline{\rho}^{E_2}, \tag{9}$$

where $D = \mathrm{diag}(\alpha, \beta, \gamma)$ and the fixed matrix T_s can be interpreted as a transformation of the original sensors to some new set of sensors. Finlayson *et al.* set out to derive transforms such that the resulting new set of sensors better supported von Kries-like adaptation. That is, they derived new sensor sets, linear transformations of the original sensors, for which Eq. (8) is a good model of illumination change. The original basis for this work is the fact that a sensor that is responsive to only a single wavelength of light renders the diagonal model a perfect model of illumination change. Given this observation, Finlayson *et al.* set out to derive sensors that were closer to this ideal case, i.e., more narrow-band, a technique they referred to as spectral sharpening.

While Finlayson *et al.*'s work suggests that a diagonal model suffices for a large class of sensors, it does not tell us for which lights and surfaces the model is exact. In previous work West and Brill set out the necessary and sufficient conditions for the model to give perfect color constancy and it is their work that motivates this article. In the next section we reformulate West and Brill's analysis and show how to derive the set of illuminants for which a von Kries-like model is exact.

WEST AND BRILL'S ANALYSIS

West and Brill[9] set out the mathematical conditions for a von Kries-like model (Eq. 8) to be exactly invariant to a change in illumination. That is, they asked the question: For what lights, surfaces, and sensors can sensor responses under a pair of lights be exactly related by a diagonal transform? Their analysis consists of two main cases: first, illumination is represented as a finite-dimensional linear model and the authors derive the conditions for a reflectance to be illuminant-invariant with respect to a von Kries-like model. In a second case the authors adopt a linear model of reflectance and derive the illuminants that are von Kries invariant with respect to this model of reflectance and some reference illuminant. That is, they recover the set of illuminants such that when surfaces are viewed under them the resulting sensor responses are exactly related by a diagonal transform. Here we present our own solution to this second case.

We begin by adopting a linear basis of surface reflectance functions as defined by Eq. (5). Now, let $\underline{\rho}_k^{\mathrm{c}}$ represent the triplet of responses of a device to the kth basis function $S_k(\lambda)$ when viewed under a reference illuminant c and let $\underline{\rho}_k^{\mathrm{o}}$ be the corresponding response under an arbitrary illuminant o. It can be shown that if responses $\underline{\rho}_k^{\mathrm{o}}$, $(k = 1 \cdots M)$ are related to responses $\underline{\rho}_k^{\mathrm{c}}$, $(k = 1 \cdots M)$ by a von Kries-like transform then so too are the responses to any other surface representable in the basis. This follows from two facts: first, that if a surface is a linear combination of basis functions then, by the linear nature of image formation, the response it induces in a device is that same linear combination of the responses induced by the basis functions, and second, the fact that if responses are related by a diagonal transform then so too are linear combinations of these responses.

So, considering just the basis functions, let $\underline{\mathrm{q}}_i^{\mathrm{c}}$ be an $M \times 1$ vector representing the response of the ith sensor to the basis functions when viewed under illuminant c. That is, for a trichromatic device,

$$\begin{aligned}
\underline{\mathrm{q}}_1^{\mathrm{c}} &= [S_1 \cdots S_M]^{\mathrm{T}}\mathrm{diag}(Q_1)E^{\mathrm{c}}\\
\underline{\mathrm{q}}_2^{\mathrm{c}} &= [S_1 \cdots S_M]^{\mathrm{T}}\mathrm{diag}(Q_2)E^{\mathrm{c}}\\
\underline{\mathrm{q}}_3^{\mathrm{c}} &= [S_1 \cdots S_M]^{\mathrm{T}}\mathrm{diag}(Q_3)E^{\mathrm{c}}.
\end{aligned} \tag{10}$$

Alternatively, we can write

$$\underline{\mathrm{r}}^{\mathrm{c}} = B^{\mathrm{T}}E^{\mathrm{c}}, \tag{11}$$

where $\underline{\mathrm{r}}^{\mathrm{c}}$ is a $3M \times 1$ vector representing the three vectors $\underline{\mathrm{q}}_i^{\mathrm{c}}$, $i = 1 \cdots 3$, stacked one on top of the other and $\boldsymbol{B}^{\mathrm{T}}$ a $3M \times 31$ matrix representing the three matrices $[S_1 \cdots S_M]^{\mathrm{T}}\mathrm{diag}(Q_i)$, $i = 1 \cdots 3$, stacked in the same way.

We can construct a similar equation for any arbitrary illuminant E^{o} by replacing $\underline{\mathrm{r}}^{\mathrm{c}}$ and E^{c} in Eq. (11) with the relevant terms. Now, if E^{o} is to be von Kries invariant with respect to the reference light E^{c} then there must exist a diagonal matrix $\boldsymbol{D}$ such that

$$\underline{\rho}^{\mathrm{o}} = \boldsymbol{D}\underline{\rho}^{\mathrm{c}} \tag{12}$$

for all responses corresponding to a surface within the span of the surface basis. That is, responses under the two lights must be related by a 3 × 3 diagonal transform. Now, any 3 × 3 diagonal matrix has three degrees of freedom and so can be expressed as a linear combination of any three linearly independent basis matrices,

$$\boldsymbol{D} = a_1\boldsymbol{D}_1 + a_2\boldsymbol{D}_2 + a_3\boldsymbol{D}_3, \tag{13}$$

and thus the von Kries invariant illuminant can in turn be expressed as the same linear combination of three illuminant spectra,

$$E^{o} = a_1E^1 + a_2E^2 + a_3E^3, \tag{14}$$

where E^i, $(i = 1 \cdots 3)$ are solutions to the equations

$$\underline{r}^i = \boldsymbol{B}^{\mathrm{T}}E^i, \qquad i = 1 \cdots 3, \tag{15}$$

where the elements of $\underline{r}^i$ are related to $\underline{r}^{\mathrm{c}}$ in the following way. Recall that $\underline{r}^{\mathrm{c}}$ is an $3M \times 1$ vector such that its first M elements represent the response of the first sensor to the M basis functions, the next M elements represent the response of the second sensor to the same functions, and the final M elements represent the response of the third sensor to the same functions. Then $\underline{r}^i$ is also a $3M \times 1$ vector whose first M elements are the elements of $\underline{r}^{\mathrm{c}}$ each multiplied by the first diagonal element of D_i. The next M elements of $\underline{r}^i$ are the elements of $\underline{r}^{\mathrm{c}}$ each multiplied by the second diagonal element of D_i and the final M elements are the elements of $\underline{r}^{\mathrm{c}}$ each multiplied by the third diagonal element of D_i.

For each value of i, Eq. (15) represents a linear system of $3M$ equations in n unknowns. In general, $3M < n$ (the number of basis functions, M is less than the number of sample points n divided by 3), which implies that the system is underconstrained and thus has infinitely many solutions.[11] In the case that the system has solutions, any solution can be expressed as a particular solution to Eq. (15) plus a solution to the corresponding homogeneous system:[11]

$$\underline{0} = B^{\mathrm{T}}E^{\mathrm{h}}. \tag{16}$$

This homogeneous system is independent of $\underline{r}^i$ and thus independent of the basis matrices D_i. So to find all von Kries illuminants we must find three particular solutions corresponding to the three sets of equations in (15) together with the solutions to Eq. (16). That is, any von Kries invariant illuminant can be expressed as

$$E^{\mathrm{o}} = E^{\mathrm{p}} + E^{\mathrm{h}} = a_1E^1 + a_2E^2 + a_3E^3 + E^{\mathrm{h}}. \tag{17}$$

A particular solution to Eq. (15) can be found,

$$\boldsymbol{E}^{\mathrm{p}} = (\boldsymbol{B}\boldsymbol{B}^{\mathrm{T}})^{-1}\boldsymbol{B}\underline{r}^i, \tag{18}$$

where $(\boldsymbol{B}\boldsymbol{B}^{\mathrm{T}})^{-1}\boldsymbol{B}$ is called the pseudo-inverse of $\boldsymbol{B}$ and is sometimes written $\boldsymbol{B}^{+}$ so that $E^{\mathrm{p}} = B^{+}\underline{r}^i$. The illuminant E^{h}, a solution to the homogeneous system, is any illuminant in the space orthogonal to the columns of $\boldsymbol{B}$, that is any vector in the null space of $\boldsymbol{B}$. This orthogonal space can be represented by the columns of a matrix $\boldsymbol{B}^{\perp}$, which is calculated

$$\boldsymbol{B}^{\perp} = I - \boldsymbol{B}^{\mathrm{T}}\boldsymbol{B}^{+}, \tag{19}$$

where I denotes the identity matrix. Thus, the orthogonal space is defined by all vectors in the span of the columns of $\boldsymbol{B}^{\perp}$. Physically we can say that any surface in the basis $[S_1 \cdots S_M]^{\mathrm{T}}$ when viewed under an illuminant E^{h} in the space $\boldsymbol{B}^{\perp}$ will elicit zero response from the sensors.

In summary, any von Kries invariant illuminant can be expressed as a linear combination of a set of basis functions,

$$E^{\mathrm{o}} = B_{v\mathrm{K}}\epsilon, \tag{20}$$

where $B_{v\mathrm{K}} = [E^{\mathrm{p}} \quad B^{\perp}]$ and ϵ is the vector of weights corresponding to the particular von Kries invariant illuminant.

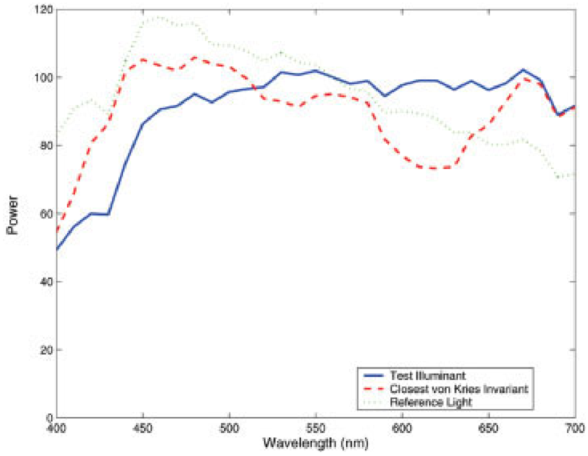

FIG. 1. An example of the predictions of West and Brill's theory. The dotted line and solid lines show the reference and test illuminant, respectively. The dashed line shows the closest von Kries invariant illuminant to the test light. [Color figure can be viewed in the online issue, which is available at www.interscience.wiley.com.]

An Example

West and Brill's analysis allows us to predict for a given reference illuminant and reflectance basis whether an arbitrary test light will be von Kries invariant for a given set of sensors. To test the theory we derived von Kries invariant illuminants with respect to the standard CIE color matching functions.[22] For our surface basis we derived a three-dimensional linear model based on a principal component[23] analysis of a composite of three published sets of reflectance data: 462 Munsell chips,[24] 170 object reflectances,[25] and 219 natural reflectances.[14] We took as a reference light a D65 daylight illuminant illustrated by the dotted line in Fig. 1. Then we determined the set of von Kries invariant illuminants using the analysis above. Given an arbitrary test illuminant we can then determine the von Kries invariant illuminant that is closest to it. That is, suppose we have a test illuminant E^{t} we can find the closest illuminant $\hat{E}^{\mathrm{t}}$, which is a linear combination of the columns of $\boldsymbol{B}_{v\mathrm{K}}$:

$$E^{\mathrm{t}} \approx \hat{E}^{\mathrm{t}} = B_{v\mathrm{K}}\underline{\hat{\epsilon}}. \tag{21}$$

The weights $\hat{\epsilon}$ that deliver the closest (in a least-squares sense) von Kries invariant illuminant can be calculated:

$$\underline{\hat{\epsilon}} = (\boldsymbol{B}^{\mathrm{T}}_{v\mathrm{K}}\boldsymbol{B}_{v\mathrm{K}})^{-1}\boldsymbol{B}^{\mathrm{T}}_{v\mathrm{K}}. \tag{22}$$

As an example, consider the test illuminant D50 shown by the solid line in Fig. 1. The closest illuminant to this that is von Kries invariant (calculated according to Eqs. (21) and (22)) with respect to the reference light is shown in Fig. 1 by the dashed line.

What conclusions can we draw from Fig. 1? Clearly, if the test light was von Kries invariant (that is, afforded perfect color constancy by the mechanism of a diagonal model) then it would be in the von Kries invariant set and the solid and dashed lines in Fig. 1 would be identical. That they are different implies that the test light is not von Kries invariant but it does not give us a quantitative measure of the degree of invariance. However, if we assume for the moment that the closer an illuminant is to the von Kries invariant set, the better the level of diagonal model color constancy we will achieve, then the implication of Fig. 1 is that the test illuminant will give quite poor diagonal model color constancy. If we repeat the analysis for a range of different illuminants we find similar results that might suggest that the diagonal model is in practice a poor model of illumination change. This conclusion, however, is at odds with the application of the model in practice, since were we to use the model to map between sensor responses, we would find that it performs reasonably well. There are two possible explanations for this apparent contradiction: first, the assumption that the closeness of an illuminant to the von Kries invariant set is a good measure as to what degree an illuminant supports the diagonal model might be wrong. Alternatively, a different theory is required to better explain the performance of the model in practice. Our aim in this article is to present such a theory. We set out this new theory in the next section and in the fifth section we present quantitative analysis that supports our contention that the theory better explains the practical performance we observe.

REASSESSING VON KRIES-LIKE ADAPTATION

The starting point for our analysis is the recognition that a von Kries-like model of illumination change treats sensors independently of one another. The model says that when illumination changes the sensor responses of each channel change by scale factors that are independent of one another. Thus, we argue, rather than considering all three sensors together and deriving the set of von Kries invariant illuminants as we did previously, that we should instead consider each sensor in isolation and derive for each an individual set of von Kries illuminants. In this way we hope to develop a theory that explains the good practical performance of von Kries-like models.

To begin, let us consider a single sensor Q_i from some arbitrary visual system. In theory this sensor can have sensitivity across the whole range of the visible spectrum but in practice many sensors have significant sensitivity across only a limited range of this spectrum. For example, consider the typical camera sensors in Fig. 2. Taken together the sensors have significant sensitivity in the range 400–700 nm but considered alone each sensor has significant sensitivity on only a restricted interval of this range. What is more, the work of Finlayson *et al.*[21] has shown that a diagonal matrix transformation of sensor responses is a better model of illumination change for sensors that have been “sharpened.” That is, if we apply a fixed 3×3 linear transform to a set of sensors such that a new set of sensors are derived that have significant sensitivity in only a restricted region of the spectrum, then the responses of these modified sensors under a pair of different illuminants can more accurately be related by a diagonal matrix transform. Thus, without loss of generality we can assume that an arbitrary sensor is effectively sensitive[†] only in a subregion of the visible spectrum.

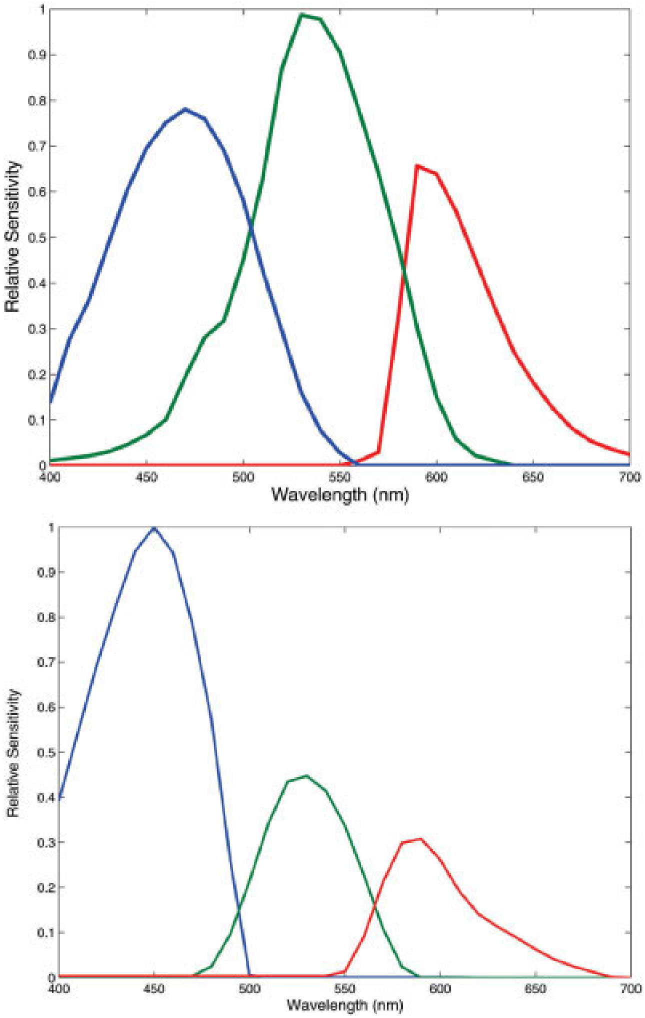

FIG. 2. Two sets of typical device sensors. [Color figure can be viewed in the online issue, which is available at www.interscience.wiley.com.]

Local Linear Models of Reflectance

For this arbitrary sensor we would like to derive the set of von Kries invariant illuminants and, like West and Brill, we do so by first adopting a linear model of surface reflectance. However, we introduce here an important modification to the usual definition of a linear model, which reflects the fact that we are treating sensors independently of one another. Suppose that our sensor is active (has significant nonzero responsivity) in a wavelength range $\lambda_s \leq \lambda \leq \lambda_e$. It follows that when modeling surface reflectances we are interested only in their behavior within this range of wavelengths since variations outside this range do not affect the sensor.

[†] By effectively sensitive we mean that the sensor has significant sensitivity.

Thus, we define a local linear model of surface reflectance with respect to an arbitrary sensor by

$$s^{\text{local}}(\lambda) = \sum_{k=1}^{N_i} \sigma_k S_k^{\text{local}}(\lambda), \qquad \lambda_s \leq \lambda \leq \lambda_e. \tag{23}$$

The basis functions themselves can be derived using techniques of statistical analysis set out by a number of previous authors.[12,14] But, importantly, the derived basis functions capture the variation in reflectances across the region of the spectrum pertinent to the sensor of interest. The number and nature of these basis functions will vary depending on the active range of the sensor Q_i but, importantly, we have found that local models typically capture a greater degree of variation in the reflectance data than does a global model with the same dimensionality. Section 5 contains an empirical analysis of global versus local linear models of reflectance.

von Kries Invariant Illuminants wrt a Single Sensor

Given this sensor specific model of surface reflectance let us now consider how we can derive the set of von Kries invariant illuminants for the sensor Q_i. As in the West and Brill analysis, to ensure that an illuminant is a von Kries invariant illuminant it is sufficient to ensure that the diagonal model holds for the surface basis functions since the result for an arbitrary surface within this basis follows from the linear nature of image formation. Thus, we begin by defining a reference illuminant E^{c} and a linear basis $[S_1^{\text{local}} \cdots S_{M_i}^{\text{local}}]$ (defined with respect to the active region of Q_i). Given these definitions we can express the response of the sensor to the basis vectors by an $M \times 1$ vector $\underline{r}^{\text{c}}$; thus,

$$\underline{r}^{\text{c}} = [S_1^{\text{local}} \cdots S_{M_i}^{\text{local}}]^{\text{T}} \text{diag}(Q_i) E^{\text{c}}. \tag{24}$$

For the case of a single sensor a von Kries illuminant is any illuminant o such that the responses of the sensor under it, $\underline{r}^{\text{o}}$, are related to the responses under the reference light by a simple scale factor:

$$\underline{r}^{\text{o}} = \alpha \underline{r}^{\text{c}}. \tag{25}$$

Thus, a von Kries illuminant is any illuminant E^{o} that satisfies the equation

$$[S_1^{\text{local}} \cdots S_{M_i}^{\text{local}}]^{\text{T}} \text{diag}(Q_i) E^{\text{o}} = \underline{r}^{\text{o}} = \alpha \underline{r}^{\text{c}}. \tag{26}$$

Equation (26) is the single sensor version of Eq. (15) and represents a system of linear equations. There are M_i equations, corresponding to the dimension of the surface reflectance basis, and the unknowns are the n elements of the illuminant E^{o}. In general M_i will be strictly less than n and so the system in Eq. (26) is underconstrained. Once again we can use basic results of linear algebra[26] to characterize the solutions to this system as the sum of a particular solution to Eq. (26) plus a solution to the corresponding homogeneous system

$$[S_1^{\text{local}} \cdots S_{M_i}^{\text{local}}]^{\text{T}} \text{diag}(Q_i) E^{\text{h}} = \underline{0}. \tag{27}$$

By inspection, αE^{c} is a solution to Eq. (26); thus, it remains to determine solutions to Eq. (27). Algebraically Eq. (27) can be interpreted as meaning that E^{h} is orthogonal to the space defined by the rows of the matrix $\boldsymbol{P} = [S_1^{\text{local}} \cdots S_{M_i}^{\text{local}}]^{\text{T}} \text{diag}(Q_i)$ or, in physical terms, that E^{h} is an illuminant that elicits a zero response in the sensor (for any surface in the basis). Solutions to Eq. (27) can be determined by finding the null space of the matrix $\boldsymbol{P}$, which we denote $\boldsymbol{P}^{\perp}$ and calculate thus:

$$\boldsymbol{P}^{\perp} = I - \boldsymbol{P}^{\text{T}} \boldsymbol{P}^{+}. \tag{28}$$

The dimensionality of this space is $n - M_i$ and thus any solution to Eq. (27) can be written as a linear combination of the columns of $\boldsymbol{P}^{\perp}$:

$$E^{\text{h}} = b_1 \underline{p}_1^{\perp} + b_2 \underline{p}_2^{\perp} + \cdots + b_{n-M_i} \underline{p}_{n-M_i}^{\perp}, \tag{29}$$

from which it follows that any von Kries illuminant can be expressed as

$$E^{\text{o}} = \alpha E^{\text{c}} + E^{\text{h}} = \alpha E^{\text{c}} + b_1 \underline{p}_1^{\perp} + \cdots + b_{n-M_i} \underline{p}_{n-M_i}^{\perp}. \tag{30}$$

Both α and the elements $b_1, b_2, \cdots b_{n-M_i}$ can be arbitrary scalars so that Eq. (30) tells us that there is an infinite set of von Kries illuminants with respect to any given reference illuminant, surface reflectance basis, and sensor. Defining a matrix $\boldsymbol{B}_{v\text{K}}^{\text{local}} = [\text{E}^{\text{c}} \quad \underline{p}_1 \cdots \underline{p}_{n-M_i}]$ we can write any such von Kries invariant illuminant as a linear combination of the columns of $\boldsymbol{B}_{v\text{K}}^{\text{local}}$,

$$E^{\text{o}} = \boldsymbol{B}_{v\text{K}}^{\text{local}} \boldsymbol{\epsilon}, \tag{31}$$

where the elements of $\boldsymbol{\epsilon}$ are the relevant scalars that define the illuminant.

Such an analysis can be performed for any arbitrary sensor; thus, for a trichromatic device we can derive three sets of von Kries illuminants, one set for each sensor. In the next section we investigate the practical implications of this theory.

EMPIRICAL ANALYSIS

In testing the theory we are interested in determining, for common sensor sets and surfaces, the degree to which common illuminants are von Kries invariant. We report here results for four different sensor sets: two sets of camera sensors (illustrated in Fig. 2) and two sets of sharpened color matching functions illustrated in Fig. 3. The top sensors in Fig. 3 were sharpened using data-based sharpening[9] while the bottom part shows sensors using the method of perfect sharpening.[20] As described in Section 2, sharpening a set of sensors amounts to transforming them by a fixed 3 × 3 linear transform. The relevant transforms $\boldsymbol{T}_{\text{PS}}$ and $\boldsymbol{T}_{\text{DS}}$ for perfectly sharpened and data-based sharpened color matching functions are

$$\boldsymbol{T}_{\text{PS}} = \begin{pmatrix} 0.8800 & -0.4748 & 0.0142 \\ -0.3917 & 0.9172 & 0.0731 \\ -0.1536 & 0.0922 & 0.9832 \end{pmatrix} \tag{32}$$

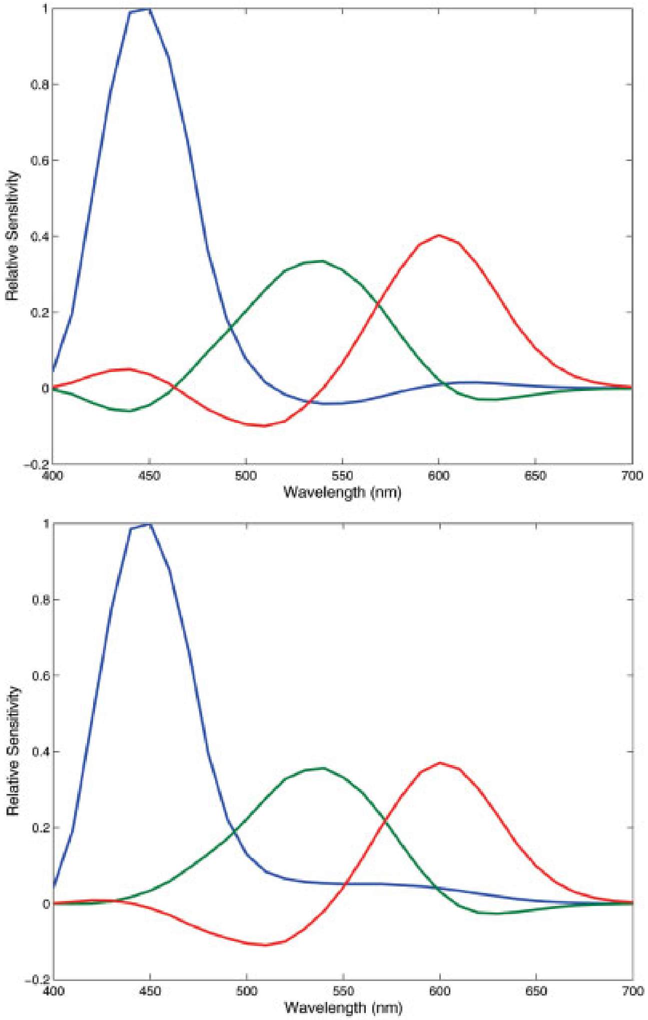

FIG. 3. Sharpened color matching functions. Top and bottom sensors are sensors resulting from data-based and perfect sharpening, respectively. [Color figure can be viewed in the online issue, which is available at www.interscience.wiley.com.]

$$T_{\mathrm{DS}} = \begin{pmatrix} 0.9214 & 0.4769 & 0.0560 \\ -0.3752 & 0.8786 & -0.0859 \\ -0.1010 & -0.0228 & 0.9947 \end{pmatrix}. \qquad (33)$$

Note that by sharpening the color matching functions we expect to improve the efficacy of a diagonal model of illumination change and we note that any advantage that this sharpening provides will be shared both by West and Brill's method and by our own theory.

Global vs Local Linear Models

Analysis is based on a composite of three published reflectance sets: 462 Munsell chips,[24] 170 object reflectances,[25] and 219 natural reflectances.[26] All linear models are derived from these three sets using a principal component analysis of the data. Let S be the 31 × 851 matrix of surface reflectance functions. We use singular value decomposition of the 31 × 31 covariance matrix (Σ) of S to derive the matrices U, Λ, and V such that $\Sigma = U\Lambda V$ where the columns of U represent the principal components of the data set and the nonzero elements of Λ are the associated singular values.

West and Brill's analysis requires a single, global linear model of reflectance, whereas our analysis uses three separate local models derived per sensor, that is, we perform the analysis above on the wavelength range to which a sensor has significant sensitivity. Figure 4 compares the first three basis functions derived locally (solid line) for short-, medium-, and long-wavelength sensitive (data-based sharpened) sensors with a globally (dashed line) derived model of surface reflectance. Note that the wavelength scale varies in each graph according to the effective sensitivity of the sensor. In the case of the data-based sharpened sensors these ranges are 540–700, 460–610, and 400–520 nm. It should be noted that outside these ranges the sensors do have some small residual sensitivity.

Table I compares the total proportion of variance accounted for by the various models for between 2 and 6 basis functions. Significantly, using a local rather than a global linear model means that more of the variation in the reflectance data is accounted for using the same number of basis functions. In the example here, 3 local basis functions account for more than 99% of the variance in the data, whereas only 98% is accounted for using a three-dimensional global model. For each sensor set tested we used the same global basis for use in West and Brill's analysis but derived local linear models individually for all sensors tested. The local linear models for different sensor sets are similar in the sense that a lower dimensional basis is able to capture a greater degree of variance in the data than is the global model. We note here that if we use three-dimensional local linear models to recover von Kries invariant sensors using the theory set out in Section 4, then in total we use 9-basis functions, whereas using a three-dimensional model in West and Brill's analysis uses a total of 3 basis functions; thus, in one sense we are using a higher dimensional model in the local analysis. However, the important point is that for any one sensor we only use a three-dimensional basis and this basis captures more of the relevant variation in the surface data than does the global model. In addition, we point out that using a higher dimensional global model (e.g., a nine-dimensional model) does not change the results of the empirical analysis that follows.

Comparative Analysis

Given a set of sensors and corresponding surface bases we chose a reference illuminant and used the analysis set out previously to derive the corresponding set of von Kries invariant illuminants. We report results here using D65 illumination as the reference light and note that similar trends were found in the results when a different reference light was chosen. In conducting the experiment we chose a test light and used both West and Brill's analysis and our own analysis to determine the closest von Kries invariant illuminant to this light.

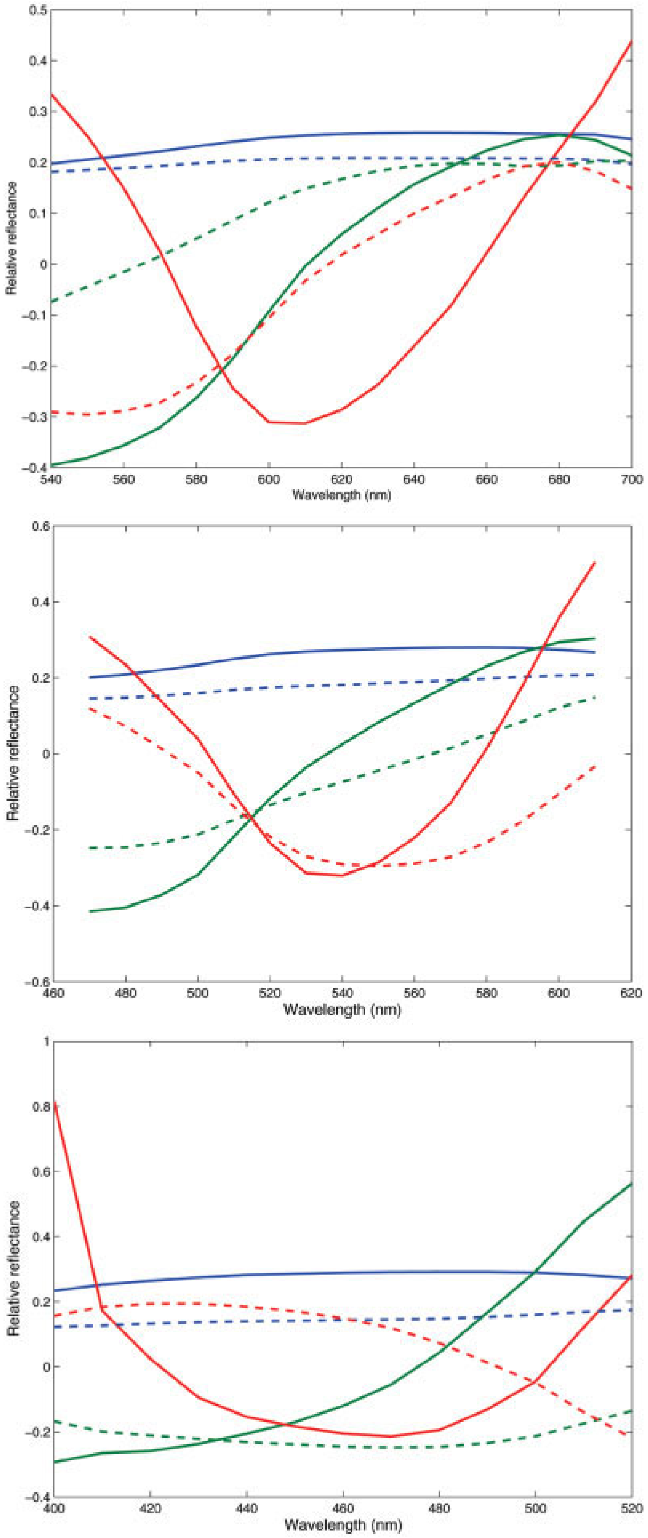

FIG. 4. Principal components, global (dashed lines) vs local (solid lines) linear model (data-based sharp). [Color figure can be viewed in the online issue, which is available at www.interscience.wiley.com.]

Figure 1 shows a sample result for the case of a D50 test illuminant (solid line) and its corresponding von Kries illuminant derived using West and Brill's analysis (dashed line). The corresponding result obtained using our own analysis is shown in Fig. 5. The upper plot in Fig. 5 shows the test illuminant SPD (solid line) and the three von Kries invariant illuminants (dashed lines) recovered using the local analysis. Visually, it is clear that in each case the von Kries invariant illuminant derived with a local analysis is closer to the test illuminant than that recovered using West

TABLE I. Total proportion of variance. Local linear models for data-based sharp sensors.

	2	3	4	5	6
Global	0.9134	0.9759	0.9855	0.9914	0.9950
Local red	0.9844	0.9935	0.9973	0.9992	0.9996
Local green	0.9659	0.9951	0.9981	0.9994	0.9997
Local blue	0.9681	0.9903	0.9965	0.9987	0.9994

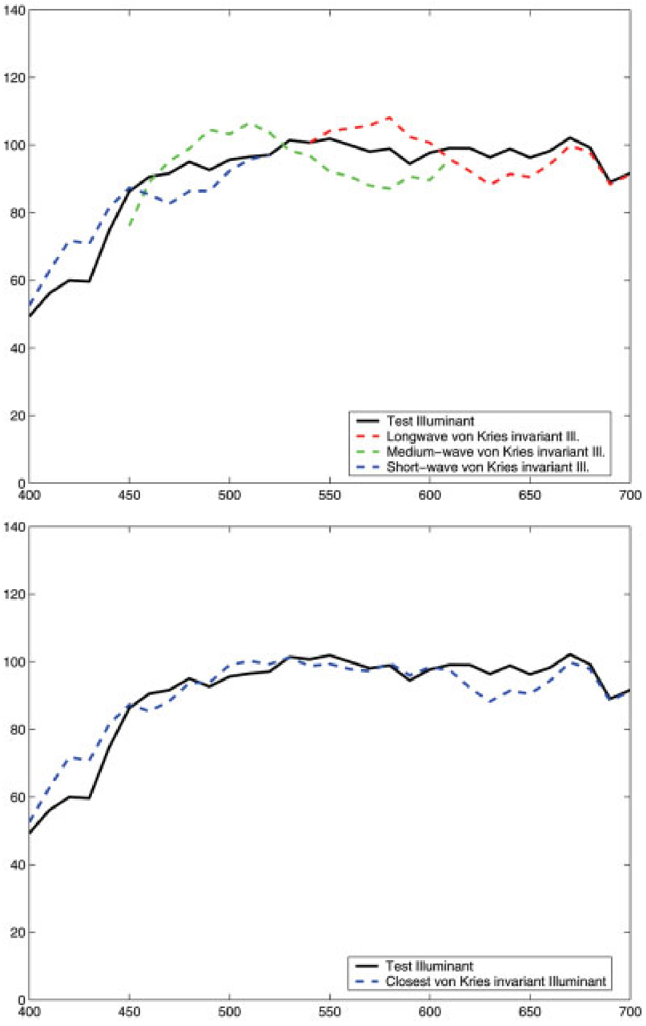

FIG. 5. An example of the von Kries invariant illuminants predicted with our local analysis. The top plot shows the von Kries invariant illuminants (dashed lines) recovered on a per-sensor basis corresponding to the test light illustrated by the solid line. The bottom graph shows the same test light (solid line) and an approximate von Kries invariant illuminant obtained by averaging the three separate recovered lights. [Color figure can be viewed in the online issue, which is available at www.interscience.wiley.com.]

TABLE II. Errors between illuminants predicted by the two theories and the actual test illuminant, averaged over 14 different lights.

	Median error	
Sensors	West and Brill	Local analysis
DB Shp, CMFs	0.109	0.050
Perf Shp CMFs	0.127	0.051
Sony camera sensors	0.071	0.053
Average, local analysis	0.131	0.045

and Brill's original analysis (see Fig. 1). The lower plot in Fig. 5 compares the test illuminant SPD (solid line) to an "average" illuminant, derived from the three illuminant spectra recovered using our analysis and each of the three individual sensors. In the regions of the spectrum where only one sensor has nonzero sensitivity, this average illuminant is exactly the same as the illuminant recovered for the sensor. However, in the regions where two or more sensors have nonzero sensitivity, the spectral power of the average illuminant is found by averaging the power of all illuminants recovered from the sensors with nonzero sensitivity. This implies that unless the three sensors are completely disjoint (in the sense that there is no region of the spectrum for which more than one sensor has significant sensitivity) this average illuminant will be only approximately von Kries invariant. However, forming an average illuminant in this way provides us with a single illuminant that can be directly compared to the actual test illuminant and the illuminant recovered using West and Brill's original analysis.

As noted earlier, if the test illuminant were perfectly von Kries invariant with respect to the reference light, then the test light and recovered lights should be identical. That they are not implies at least some degree of noninvariance, both for the light predicted using West and Brill's original analysis and for those predicted using the theory set out in this article. It is not clear, though, how exactly differences between the actual illuminant and its closest von Kries invariant illuminant will manifest themselves as errors in a diagonal mapping between lights. We might make the assumption that the closer an illuminant is to the von Kries invariant set, the better will be its diagonal model performance. But we should really test this assumption by comparing the errors we get when using a diagonal model to map between illuminants with the error between a test light and its closest von Kries illuminant. In the context of such a test we can assess whether one of the two theories discussed in this work is better than the other by looking at the correlation between the errors predicted by the theory and the actual diagonal model errors. That is, we would expect the errors predicted by the theory to be highly correlated to the actual diagonal model errors, if that theory is a good prediction of diagonal model performance.

For a given set of surfaces and sensors we can test the accuracy of a diagonal model of illumination change in the following way. Let us denote by e^c and e^o a reference and test illuminant, respectively. For these two lights we can calculate sensor responses for all surfaces viewed under each light. Let us denote by $\underline{q}_i^c$ and $\underline{q}_i^o$ the vectors of responses of the ith sensor under the reference and test light, respectively. Assuming a diagonal model of illumination change these two vectors are related:

$$\underline{q}_i^c = d_i \underline{q}_i^o, \tag{34}$$

where d_i is a scalar mapping the sensor responses. For any pair of illuminants we can calculate the best (in a least-squares sense) scalar that maps between their sensor responses. Let us denote this mapping $\hat{d}_i$. Then the error in the diagonal model for the ith sensor can be written as

$$\text{Error}_{\text{Diag}} = (\underline{q}_i^c - \hat{d}_i \underline{q}_i^o)^T (\underline{q}_i^c - \hat{d}_i \underline{q}_i^o). \tag{35}$$

Now, for the reference light e^c we can use our theory to determine the set of von Kries invariant illuminants and we can determine the closest (in a least-squares sense) von Kries illuminant to the test light E_o. Let e_i^o represent the portion of the test illuminant spectrum in the region in which the ith sensor has significant sensitivity and $\hat{e}_i^o$ represent the corresponding portion of the closest von Kries invariant light. We can measure the error between these two spectra in a least-squares sense as

$$\text{Error}_{\text{ill}} = \frac{(e_i^o - \hat{e}_i^o)^T (e_i^o - \hat{e}_i^o)}{(e_i^o)^T (e_i^o)}. \tag{36}$$

The smaller the error between these two spectra, the closer is the test light to the von Kries invariant set. Table II shows the median value of the error measure defined in Eq. (36) calculated over 14 different illuminants (shown in Fig. 6) for four different sensor sets: the color matching functions (CMFs) sharpened using the method of data-based sharpening, CMFs sharpened using the method of perfect sharpening, and the sensors from a Sony and a Nikon digital still camera. Results are shown for the illuminant predicted by the West and Brill analysis and also for the average illuminant calculated using the analysis presented in this article. In all cases the local analysis presented in this article recovers an illuminant that is closer to the actual test illuminant than

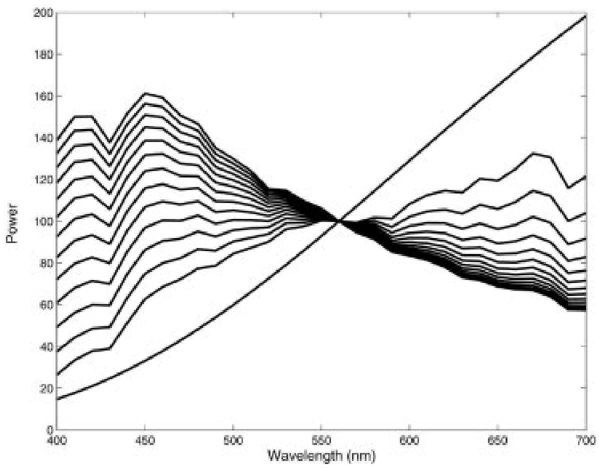

FIG. 6. Illuminant SPDs.

TABLE III. Correlation of diagonal model errors and the errors between illuminant spectra for a range of different sensors, averaged over 14 different lights.

	Wav. Int.			
Sensors and method	Sensor 1	Sensor 2	Sensor 3	All sensors
DB Shp, CMFs, West and Brill	0.89	0.97	0.95	0.94
DB Shp, CMFs, local analysis	0.99	0.97	0.97	0.98
Perf Shp, CMFs, West and Brill	0.87	0.98	0.99	0.95
Perf. Shp, CMFs, local analysis	0.98	0.97	0.97	0.98
Sony camera sensors, West and Brill	0.88	0.98	0.96	0.94
Sony camera sensors, local analysis	0.92	0.97	1.00	0.96
Nikon camera sensors, West and Brill	0.81	0.99	0.86	0.89
Nikon camera sensors, local analysis	0.97	1.00	0.96	0.98
Average (over all sensors), West and Brill	0.86	0.98	0.94	0.93
Average (over all sensors), local analysis	0.97	0.98	0.98	0.98

that predicted by the analysis of West and Brill. That is, the new theory predicts better von Kries invariance than the original theory.

The final step is to look at how these spectral errors correlate with the diagonal model errors given by Eq. (35). That is, we wish to determine which of the two sets of errors summarized in Table II is better correlated with the actual errors obtained by using a diagonal mapping to correct illumination. We do this by looking at the correlation coefficient,

$$R = \frac{\sum_j (\mathrm{Error}^j_{\mathrm{Diag}} - \overline{\mathrm{Error}}_{\mathrm{Diag}})(\mathrm{Error}^j_{\mathrm{ill}} - \overline{\mathrm{Error}}_{\mathrm{ill}})}{\sum_j (\mathrm{Error}^j_{\mathrm{Diag}} - \overline{\mathrm{Error}}_{\mathrm{Diag}})^2 \sum_j (\mathrm{Error}^j_{\mathrm{ill}} - \overline{\mathrm{Error}}_{\mathrm{ill}})^2}, \tag{37}$$

where the summations are taken over a number of different test illuminants. R takes a value between zero and 1: a value of 1 implies that the two errors are perfectly correlated whereas a value of zero implies the errors are completely uncorrelated. In the context of our experiment a correlation of 1 implies that the distance of all test illuminants from the von Kries invariant set is perfectly correlated with the diagonal model errors.

To quantify the correlation of errors for each of the two theories we calculated diagonal model errors for the same sets of sensors and illuminants as described above. We also calculated the von Kries invariant set using each of the two theories and for each test light calculated the closest von Kries invariant light. We then used these data to calculate the correlation coefficients according to Eq. (37) on a per-sensor basis. That is, we looked at diagonal model errors for each sensor separately and compared it to the error in illuminant spectra over the range for which the sensor had significant sensitivity.

Table III shows the correlations for the different sensor sets. Results are shown for the local linear analysis (using 3D local bases) and for West and Brill's analysis using a 3D global basis. The first three columns show the correlation coefficients calculated over the wavelength interval for which a given sensor has nonzero sensitivity. The average correlation coefficient calculated over the wavelength interval for which one or more sensors has significant sensitivity is shown in column 4.

A number of interesting points are revealed by these results. First, if we look at the results for both analyses we see that, in general, the errors between the illuminant that each method recovers and the actual illuminant correlate quite well with the diagonal model errors. Indeed, in the case of the West and Brill analysis, the degree of correlation is perhaps better than we might predict based on an informal visual analysis of the recovered von Kries invariant illuminant (Fig. 1). Of equal interest is the comparative performance of the two different methods we have considered in this article. Here, we see that in general the errors predicted using the local theory have a correlation that is as good as, or significantly better than, the correlation for the West and Brill case. In particular, where West and Brill's analysis has poorer correlation with the diagonal model errors, the new theory is well correlated. On average, over all devices and sensors the new method has a correlation of 0.98 compared with 0.93 for West and Brill's analysis—a significant improvement. Overall, the results support the view that the local based analysis we have presented in this article is a better explanation as to why a von Kries-like model is a better model of illumination change than West and Brill's original analysis.

CONCLUSIONS

In this article we set out to explore the conditions under which a von Kries-like model of adaptation is a perfect model of illumination change. We began by reformulating the previous analysis of West and Brill[9] to derive a set of von Kries illuminants with respect to a set of sensors, surfaces, and a reference light. Application of this theory using lights, surfaces, and sensors that occur in the world predicts that von Kries-like adaptation is a poor model of illumination change. This prediction contradicts empirical evidence that suggests that model is actually quite accurate.

To address this apparent contradiction we presented a modified theory that considers sensors independently of one another. Considering a single sensor we showed how to derive the set of illuminants for which a diagonal model is a perfect model of illumination change. To help us in this task we introduced the idea of a local linear model that models reflectance data only within the restricted region to which a sensor has significant sensitivity. We applied this modified theory to the same sets of sensors, surfaces, and illuminants and were able to show that in this case the predicted von Kries illuminants were close to actual illuminants. We further showed that errors in these predicted illuminants were very well correlated with diagonal model errors, thus supporting the empirical evidence that a von Kries-like model accurately accounts for illumination change.

1. Finlayson GD, Hordley SD, Hubel PM. Color by correlation: a simple, unifying framework for color constancy. IEEE Trans Pattern Anal Mach Intell 2001;23:1209–1221.
2. Forsyth DA. A novel algorithm for colour constancy. Int J Comput Vis 1990;5:5–36.
3. Land EH. The retinex theory of color vision. Sci Am 1977;108–129.
4. Maloney LT, Wandell BA. Color constancy: a method for recovering surface spectral reflectance. J Opt Soc Am A 1986;3:29–33.
5. Arend L, Reeves A. Simultaneous color constancy. J Opt Soc Am A 1986;3:1743–1751.
6. Brainard DH. Color constancy in the nearly natural image. 2. Achromatic loci. J Opt Soc Am A 1998;15:307–325.
7. Brainard DH, Brunt WA, Speigle JM. Color constancy in the nearly natural image. I. Asymmetric matches. J Opt Soc Am A 1997;14: 2091–2110.
8. Jameson D, Hurvich LM. Essay concerning color constancy. Ann Rev Psychol 1989;40:1–22.
9. West G, Brill MH. Necessary and sufficient conditions for von Kries chromatic adaptation to give color constancy. J Math Biol 1982;15: 249–258.
10. Finlayson G. Coefficient colour constancy. Ph. D. thesis, Simon Fraser University, 1995.
11. Brainard DH, Wandell BA, Cowan WB. Black light: how sensors filter spectral variation of the illuminant. IEEE Trans Biomed Eng 1989;36: 140–149.
12. Maloney LT. Evaluation of linear models of surface spectral reflectance with small numbers of parameters. J Opt Soc Am A 1986;3: 1673–1683.
13. Marimont DH, Wandell BA. Linear models of surface and illuminant spectra. J Opt Soc Am A 1992;9:1905–1913.
14. Parkkinen J, Jaaskelainen T, Kuittinen M. Spectral representation of color images. In IEEE 9th International Conference on Pattern Recognition, vol 2, November 1998. p 933–935.
15. Horn BKP. Robot vision. MIT Press; 1986.
16. Smith B, Spiekermann C, Sember R. Numerical methods for colorimetric calculations: sampling density requirements. Color Res Appl 1992;17:394–401.
17. Judd DB, MacAdam DL, Wyszecki GW. Spectral distribution of typical daylight as a function of correlated colour temperature. J Opt Soc Am A 1964;54:1031–1040.
18. von Kries J. Influence of adaptation on the effects produced by luminous stimuli. In MacAdam DL, editor. Sources of color science. Cambridge, MA: MIT Press; 1970. p 109–119.
19. Bloj MG, Kersten D, Hurlbert AC. Preception of three-dimensional shape influences colour perception through mutual illumination. Nature 1999;402:877–879.
20. Finlayson GD, Drew MS, Funt BV. Color constancy: generalized diagonal transforms suffice. J Opt Soc Am A 1994;11:3011–3019.
21. Finlayson GD, Drew MS, Funt BV. Spectral sharpening: sensor transformations for improved color constancy. J Opt Soc Am A 1994;11: 1553–1563.
22. Hunt RWG. Measuring colour, 2nd ed. London: Ellis Horwood; 1991.
23. Jolliffe IT. Principal component analysis. New York: Springer-Verlag; 1986.
24. Newhall SM, Nickerson D, Judd DB. Final report of the OSA subcommittee on the spacing of the Munsell colours. J Opt Soc Am A 1943;33:385–418.
25. Vrhel MJ, Gershon R, Iwan LS. Measurement and analysis of object reflectance spectra. Color Res Appl 1994;19:4–9.
26. Golub GH, van Loan CF. Matrix computations. John Hopkins University Press; 1983.