PHILOSOPHICAL TRANSACTIONS OF THE ROYAL SOCIETY B

Phil. Trans. R. Soc. B (2005) **360**, 1329–1346
doi:10.1098/rstb.2005.1633
Published online 28 June 2005

Review

Sensory, computational and cognitive components of human colour constancy

H. E. Smithson*

Institute of Ophthalmology, University College London, 11–43 Bath Street, London EC1V 9EL, UK

When the illumination on a scene changes, so do the visual signals elicited by that scene. In spite of these changes, the objects within a scene tend to remain constant in their apparent colour. We start this review by discussing the psychophysical procedures that have been used to quantify colour constancy. The transformation imposed on the visual signals by a change in illumination dictates what the visual system must 'undo' to achieve constancy. The problem is mathematically underdetermined, and can be solved only by exploiting regularities of the visual world. The last decade has seen a substantial increase in our knowledge of such regularities as technical advances have made it possible to make empirical measurements of large numbers of environmental scenes and illuminants. This review provides a taxonomy of models of human colour constancy based first on the assumptions they make about how the inverse transformation might be simplified, and second, on how the parameters of the inverse transformation might be set by elements of a complex scene. Candidate algorithms for human colour constancy are represented graphically and pictorially, and the availability and utility of an accurate estimate of the illuminant is discussed. Throughout this review, we consider both the information that is, in principle, available and empirical assessments of what information the visual system actually uses. In the final section we discuss where in our visual systems these computations might be implemented.

Keywords: colour vision; colour constancy; von Kries transformation; cone-excitation ratios; illuminant estimation

1. INTRODUCTION

(a) *What is colour constancy?*

Through our senses, we recover the information that we need in order to interact effectively with the world around us, but there is not a simple relationship between the properties of objects in the world, and the information collected by our sense organs. For example, the size of the retinal image of an object depends both on the size of the object and on the viewing distance; its shape depends both on the shape of the object and on the viewing angle; and its intensity depends both on the reflectance of the object and on the light illuminating it. The term *perceptual constancy* describes the extent to which objects appear unchanging despite changes in the conditions of observing. The level of 'constancy' differs according to whether an observer is asked to make accurate judgements about how things are in the world, or is asked about his sensations. This is the distinction between *performance constancy* and *phenomenal regression to the real object* (Thouless 1931).

Colour constancy describes the extent to which an observer can stably recognize the spectral reflectance of an object's surface (or in perceptual terms, its colour), despite changes in conditions of observing that change the spectral properties of the light reaching the eye. This paper presents a short review of the sensory, computational and cognitive aspects of human colour constancy under a change in the spectral composition of the illumination.[1]

(b) *The nature of the problem*

To understand colour constancy we must consider both the physical properties of the world, and the biological and psychological properties of the observer (see figure 1). The light that reaches a local area of retina depends both on the spectral reflectance of the object or surface in view, and on the spectral composition of the illuminant.[2] An illuminant is characterized by its spectral composition, $E(\lambda)$, i.e. the energy of light at each wavelength, λ; an object or surface is characterized by its spectral reflectance function, $R(\lambda)$, i.e. the proportion of light reflected as a function of wavelength. Physical measurements of spectral reflectance can be obtained by recording the reflected light with a spectroradiometer and (on a wavelength by wavelength basis) dividing by the spectrum of the illuminant. However, the human visual system, dependent on the photon catches in just three classes of photoreceptor, does not have access to full spectral information about the light that reaches the retina, and in general, cannot obtain direct information about the illuminant.

The starting point for visual experience is the signal from the photoreceptors. The photoreceptors are characterized by their absorption spectra (e.g. for

* (h.smithson@ucl.ac.uk).

Received 14 June 2004
Accepted 22 January 2005

© 2005 The Royal Society

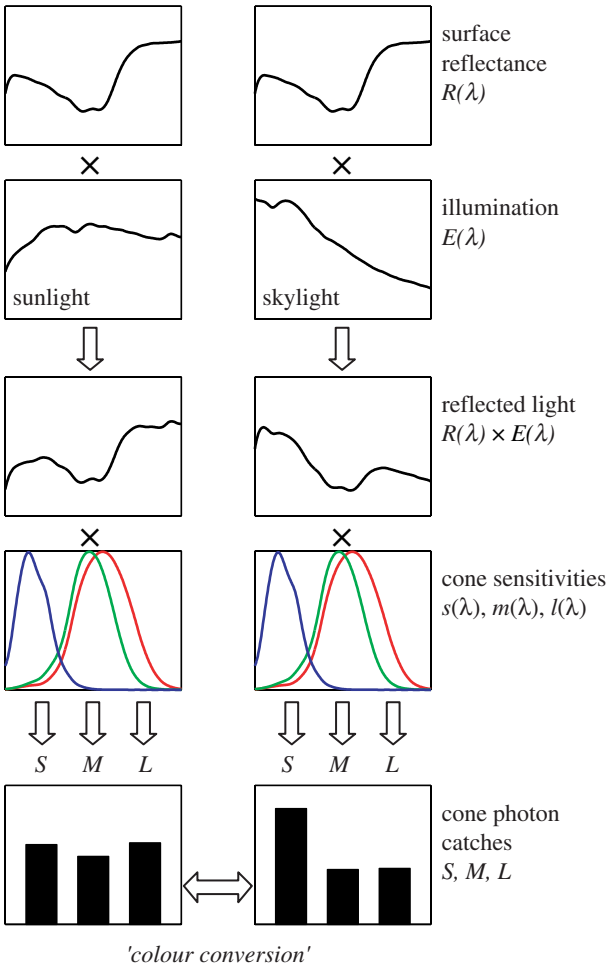

Figure 1. The first row shows $R(\lambda)$, the surface reflectance of an object. Reflectance takes values between 0 (no reflectance) and 1 (perfect reflectance), and is plotted as a function of wavelength from 400 to 700 nm. The second row shows $E_{sun}(\lambda)$ and $E_{sky}(\lambda)$, the energy spectra of two daylight illuminants (Taylor & Kerr 1941), and the third row shows the spectrum of light reflected from the surface, i.e. a wavelength-by-wavelength multiplication of $R(\lambda)$ and $E(\lambda)$. The fourth row shows the spectral sensitivities of the human cones, normalized to equal peak sensitivities. The final row shows the *S*-, *M*- and *L*-cone photon catches for the surface under sunlight (left) and skylight (right).

trichromatic humans, $s(\lambda)$, $m(\lambda)$, $l(\lambda)$ for short-, middle- and long-wave sensitive cones, respectively), and the resultant cone photon catches are given by:

$$S = \int s(\lambda)R(\lambda)E(\lambda)\mathrm{d}\lambda, \tag{1.1a}$$

$$M = \int m(\lambda)R(\lambda)E(\lambda)\mathrm{d}\lambda, \tag{1.1b}$$

$$L = \int l(\lambda)R(\lambda)E(\lambda)\mathrm{d}\lambda. \tag{1.1c}$$

A change in the spectrum of the illumination causes a change in the spectra of lights reflected from objects, and hence a change in the cone signals elicited by those objects, a process we will refer to as 'colour conversion' (Helson 1938). With access only to the cone signals, how does the visual system disentangle the information about objects in the world from the information about the light illuminating them?

(c) *The nature of the solution*

In abstract terms, colour constancy is achieved when a neural process transforms the signals elicited by surfaces under a test illuminant towards the signals that would be elicited by the same surfaces under a reference illuminant. Models of colour constancy must first specify the form of this neural transformation and its associated parameters (see §3), and then identify the information that must be extracted from the retinal image in order to set these parameters (see §4). Similar frameworks have been proposed for models of colour adaptation by Stiles (1961) and for models of colour constancy by Krantz (1968), Maloney (1999) and Brainard (2004). In addition, any complete model of human colour constancy must specify where in our perceptual apparatus the proposed transformations are performed (see §5).

(d) *Evaluating different 'solutions'*

We know that in general, a human observer cannot be perfectly colour constant, for the physical constraints on the problem are insufficient. But regularities of the natural world provide additional constraints, some of which are presumably exploited by the visual system. The computational literature evaluates the plausibility and utility of particular constraints; empirical studies with human observers are required to discover what information human observers actually use. By measuring failures of human colour constancy under situations that satisfy some constraints but not others, we can make critical tests of candidate algorithms for human colour constancy. We start this review by discussing the non-trivial issue of how to measure the level of constancy exhibited by human observers.

2. MEASURING HUMAN COLOUR CONSTANCY

(a) *The observer's task*

(i) *Matching to a standard display*

Perhaps the most popular approach is to ask how well an observer can match the colour of a surface seen under one illuminant to the colour of a test surface seen under a second illuminant (*asymmetric colour matching*). The scenes may be real or artificial, and usually comprise multiple surfaces. The two illuminant-conditions may be presented side-by-side (*simultaneous asymmetric matching*; e.g. Arend & Reeves 1986; Arend *et al.* 1991; Brainard *et al.* 1997), or one after the other (*successive asymmetric matching*; e.g. Brainard & Wandell 1992), or to different eyes (*haploscopic matching*; e.g. McCann *et al.* 1976). Simultaneous matching has the drawback that adaptation to the two scenes will be determined by the pattern of eye movements across the two halves of the scene. Successive matching allows experimental control of adaptation to the two illuminants, but performance will additionally depend on the observer's ability to remember colours. Haploscopic matching allows separate adaptation states in the two eyes, but removes binocular cues to scene geometry.

(ii) *Matching to an internal standard*

An alternative method is to ask an observer to adjust a test patch within a scene to appear white (*achromatic*

setting; e.g. Fairchild & Lennie 1992; Brainard 1998). Importantly for these studies, Morgan *et al.* (2000) have shown that accuracy can be as great with an implicit as with an explicit standard. Achromatic setting provides information about transformations of only a single point in perceptual colour space, and, therefore, cannot provide a general test of colour constancy mechanisms. An extension of the method is to ask observers to make settings of a particular chromatic locus (for example, colours that appear neither 'reddish' nor 'greenish') or to combine measurements of multiple colour loci to track changes in the structure of perceptual colour space (Chichilnisky & Wandell 1999; Smithson & Zaidi 2004).

(iii) *Colour naming*

A further alternative is to measure whether a surface is assigned the same colour name under different illuminants (Troost & de Weert 1991). Our ability to discriminate surface colours far exceeds the number of discriminations represented in our vocabulary, so attempts have been made to improve the precision of colour naming by adding ratings (Speigle & Brainard 1996).

(b) *The question to ask the observer*

As mentioned in §1*a*, it is possible to distinguish between colour constancy based on the ability to recognize the invariance of objects in the world, and colour constancy based on the stability of appearance. Arend & Reeves (1986) have presented a clear demonstration of the influence of instructions in a colour-matching task. When observers were asked to make a match to 'look as if it were cut from the same piece of paper', they showed relatively good constancy compared with conditions where they were asked to match 'hue and saturation' (see also Arend *et al.* 1991; Cornelissen & Brenner 1995). The appearance-based constancy obtained in the second case is a demonstration of phenomenal regression to the real object (Thouless 1931). While we can be sure that, whichever measurement technique is used, less constancy will be achieved as we progressively reduce the contextual cues available, our sensations and our judgements of the outside world can be decoupled.

Perception is an underdetermined problem; multiple physical arrangements can give rise to the same sensory inputs. In order to reconstruct the external world, our perceptual systems use the incoming data in conjunction with constraints imposed by the regularity of the world (see §3) and expectations that arise from recent or concurrent experience (see §4*d*). Different cues, or instructions, may suggest alternative perceptual organizations, which in turn may influence perceived colours (Judd 1940; Adelson 1993; Schirillo & Shevell 2000).

(c) *Representation of surfaces and illuminants*

An interesting issue is whether an observer can represent, simultaneously, the colour of a surface and that of the light illuminating it (Arend 1994; Mausfeld 1998; MacLeod 2003). Rather than discounting the illuminant, would it not be more desirable to recognize that surfaces were being viewed under different illuminants, to infer the relative properties of different illuminants, and the identity of surfaces across illuminant changes (e.g. Zaidi 2001)? Lichtenberg raised just this issue. In a letter to Goethe (7 October 1793), he writes, 'In ordinary life we call white, not what looks white, but what would look white if it was set out in pure sunlight… we believe at every moment that we sense something which we really only conclude' (Joost *et al.* 2002, p. 302).

(i) *Discriminating a change in illumination from a change in surface reflectance*

Foster and colleagues have proposed an *operational* approach to colour constancy in which observers are asked to discriminate between a change in illumination and a change in surface reflectance (Craven & Foster 1992; Nascimento & Foster 1997). Observers are exceedingly sensitive to such differences, and in §3*b*(iv) we discuss the neural signals that might support such discriminations.

(ii) *Identifying surfaces across illuminants*

Zaidi (1998, 2001) advocates a forced-choice measure of performance colour constancy in which the observer is required to identify like surfaces across illuminants, despite obvious differences in appearance. In a typical experiment, four surfaces are presented, two under each illuminant. Three surfaces have the same reflectance, and one has a different reflectance. The observer's task is to identify the 'odd one out'. To choose the correct surface requires observers first to choose the illuminant condition in which the surfaces differ (a chromatic discrimination task), and then to identify which of those two surfaces is the same as the two standard surfaces under the second illuminant (the constancy task). Khang & Zaidi (2002) found that, in the majority of cases, identification performance was limited only by the limen of discrimination.

3. THE TRANSFORMATION IMPOSED BY A CHANGE IN ILLUMINATION

(a) *An over-simplification?*

We must now turn to the colour conversion imposed on the visual signals by a change in illumination, for it is only through understanding this transformation and its associated parameters that we can design experiments which test how the visual system might 'undo' the colour conversion to achieve constancy.

In a world composed of surfaces that reflect all parts of the spectrum equally, signals representing the level of light reflected from a set of surfaces would retain their relative magnitude under a change in illumination. In such a world, lightness constancy could be achieved by *normalizing* each of the signals by the signal obtained from direct sampling of the illuminant.

Historically, it has been common to suggest that colour constancy could be achieved by a similar 'discounting' of the illuminant (Helmholtz 1866). However, as shown in equations (1.1*a*)–(1.1*c*), once light has been absorbed by the cones, the components of the signal that derive from the illumination, and those that derive from modification by spectrally selective reflection from a surface, are confounded.

Box 1 An over-simplification if, $S_{R,E} = \int s(\lambda)R(\lambda)E(\lambda)\mathrm{d}\lambda$ then, $S_E = \int s(\lambda)E(\lambda)\mathrm{d}\lambda$, when $R(\lambda) = 1$ and, $S_R = \int s(\lambda)R(\lambda)\mathrm{d}\lambda$, when $E(\lambda) = 1$ but, $S_R \neq S_{R,E}/S_E$. $S_{R,E}$ is the *S*-cone signal from a surface, $R(\lambda)$, seen under illuminant $E(\lambda)$. When $R(\lambda) = 1$ the surface has a spectrally uniform reflectance, and thus faithfully reflects the illuminant spectrum. When $E(\lambda) = 1$, the illuminant is a reference light with a constant unit spectral power density. The *S*-cone signal for the surface, S_R cannot generally be recovered by dividing $S_{R,E}$ by the *S*-cone signal for the illuminant, S_E.

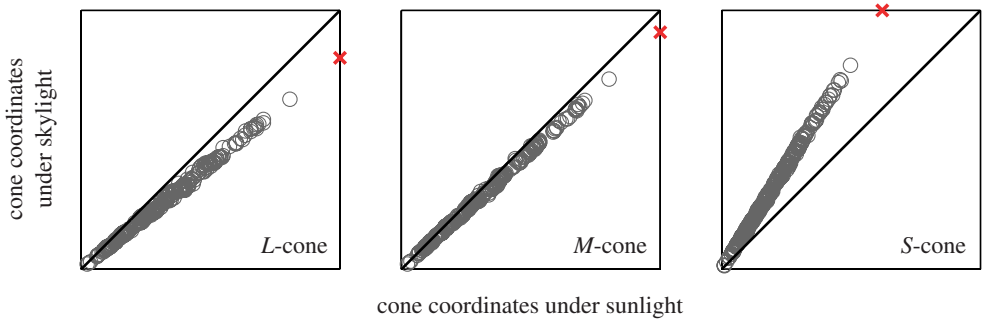

Figure 2. The three panels show *L*-, *M*- and *S*-cone coordinates from each of 280 reflectance spectra (Taylor & Kerr 1941; Chittka *et al.* 1994; Vrhel *et al.* 1994; Hiltunen 1996; Marshall 2000) rendered under two illuminants: zenith skylight on the ordinate and direct sunlight on the abscissa. In each plot, the red cross represents the photon-catch from a surface of perfect, uniform spectral reflectance and is equivalent to the photon-catch obtained from direct sampling of the illuminant. The solid black line is the unit diagonal and represents the case where cone coordinates are the same under the two illuminants. The change from skylight to sunlight causes an approximately multiplicative change in cone coordinates. A neural process that is able to 'undo' this mapping will result in colour constancy. Metamers are represented in these plots as pairs of points that, for all three cone types, lie on vertical or horizontal lines (i.e. match under one illuminant, but not under the other illuminant).

It might seem plausible that in order to recover an illuminant-independent description of the surface, one could simply divide the cone signals for the surface seen under the illuminant by the cone signals for the illuminant itself. However, if $E(\lambda)$, $R(\lambda)$ and $S(\lambda)$, $M(\lambda)$ and $L(\lambda)$ are arbitrary functions of wavelength, there is no *mathematical* reason why this simplification should work, even approximately. Maloney (1999, p. 392) has labelled this false simplification the 'RGB heuristic', and it is summarized by the equations in box 1. Much of the work on algorithms for colour constancy has been directed towards specifying the necessary constraints, for illuminants, reflectances and spectral sensitivity functions, under which equations (1.1*a*)–(1.1*c*) can be simplified. Ultimately, the relevance of such constraints for human colour constancy can be determined only by thorough sampling of environmental reflectances and illuminants (Foster & Nascimento 1994). So, is it reasonable to simplify equations (1.1*a*)–(1.1*c*), and if so, how?

(b) *Colour conversion expressed in cone coordinates*

(i) *What happens to the cone signals?*

One of the most important findings of the last decade has been the empirical result that for environmental surfaces and illuminants, the colour conversion between two illuminant conditions has a simple form when expressed in terms of cone coordinates (Dannemiller 1993; Foster & Nascimento 1994; Zaidi *et al.* 1997; Nascimento *et al.* 2002).[3] The result is illustrated in figure 2. Each data point represents one of a sample of 280 natural objects. The *L*- (or *M*- or *S*-) cone coordinate for each object under zenith skylight is plotted versus the *L*- (or *M*- or *S*-) cone coordinate of the same object under direct sunlight, and the diagonal of unit slope represents the case where signals are equal under the two illuminants.

We should note four properties of the data in these plots: (i) to a first approximation, a change in the spectrum of the illumination produces multiplicative changes in cone coordinates; (ii) the cone coordinates of the illuminant plot at the far end of the line of surfaces; (iii) within each cone class, an illuminant change does not significantly disturb the rank order of signals from a set of surfaces; and (iv) to a large extent, cone coordinates from different surfaces maintain their relative positions under a change in illumination, a property that Foster and colleagues describe as the *invariance of cone-excitation ratios*. Clearly, if a change in the spectrum of the illumination produced exactly multiplicative changes in cone coordinates, the data points in figure 2 would lie perfectly on straight lines through the origin, and the invariance of rank order and ratios of excitation within each cone class would naturally follow. In §3*c* we discuss the perceptual relevance of scatter in plots similar to those in figure 2, but first, we discuss how each of these properties might be exploited in colour constancy.

(ii) *Multiplicative scaling of cone-signals*

Ives (1912) may have been the first to suggest an explicit mechanism for constancy under an illuminant change (Brill 1995). Using arbitrary surface–reflectance functions, and Koenig fundamentals to represent receptor sensitivities, he showed that the multiplicative factors which transform the coordinates of incandescent carbon illumination to those of a reference illuminant, also transform the coordinates of *surfaces* to approximately their coordinates under the reference illuminant. Ives' observations are confirmed for the set of reflectances and illuminants sampled in figure 2, and indeed, for the large samples of environmental surfaces and illuminants

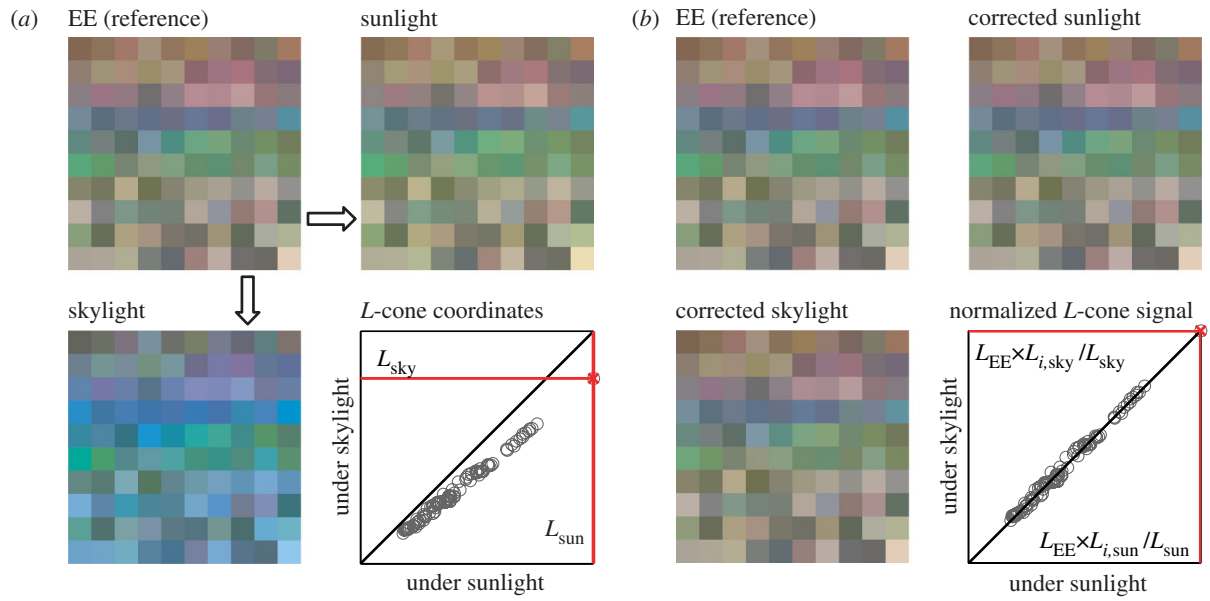

Figure 3. The four panels in (*a*) describe an illuminant change on a synthetic scene. Each of the small squares in the reference image (top left) was coloured by calculating the RGB values (i.e. the levels of the RGB phosphors on a computer monitor) that would produce in the *L*-, *M*- and *S*-cones the same triad of responses as would be produced by one of the reflectance spectra used in figure 2 which was illuminated by an equal energy illuminant ($E_{EE}(\lambda)=1$). Each of the 100 squares corresponds to a different reflectance spectrum. The small square in the bottom right represents a perfectly reflecting surface and thus produces the same cone signals as the illuminant. The 'sunlight' image (top right) shows the same arrangement of reflectance spectra illuminated by sunlight, and the 'skylight' image (bottom left) shows the same scene illuminated by skylight. The plot (bottom right) shows the *L*-cone signals for the same set of reflectance spectra under skylight plotted against the signals obtained under sunlight. The red cross indicates the surface of perfect spectral reflectance and plots at (L_{sun}, L_{sky}). The four panels in (*b*) describe the operation of the Ives transform on these signals. The images 'corrected sunlight' and 'corrected skylight' are transformed versions of images 'sunlight' and 'skylight'. Multiplicative constants for the three cone types have been chosen such that the signals for the perfectly reflecting surface are identical in these images and in the reference image. The Ives transform has successfully generated images that are approximately illuminant-invariant. (NB: There are dangers in performing such simulations. First, owing to the vagaries of colour reproduction, the images presented in your copy of this article will not produce the intended triads of *L*-, *M*- and *S*-cone signals. Furthermore, errors that do not appear large under particular viewing conditions may become significant in other circumstances, and vice versa.) The plot (bottom right) shows quantitative data on the normalization of the *L*-cone signal. The signal from each surface has been normalized to the signal from the test illuminant, and multiplied by the signal from the reference illuminant.

discussed above. Hence, the 'over-simplification' described in §3*a*, and in box 1, is not so wildly wrong for human observers in the physical world.

We refer to Ives' proposal as the Ives transform.[4] Mathematically, it consists of multiplying all cone coordinates by the same diagonal matrix, the elements of which are set by the illuminant's cone coordinates. Under Ives' account, striking failures of colour constancy are predicted if the visual system uses the wrong estimate for the illuminant. In §4, we discuss proposals for how the illuminant's cone coordinates might be estimated from a complex scene. Figure 3 illustrates the operation of the Ives transform on a synthetic scene, assuming perfect knowledge of the illuminant's cone coordinates.

In 1878, von Kries suggested that colour adaptation might be described by independent, multiplicative gain controls within each class of receptor (von Kries 1878; 1905). Von Kries' *coefficient rule* is incorporated in many modern theories of adaptation. However, von Kries' original suggestion was that, at each retinal location, the coefficients were set in inverse proportion to the local stimulation. This operation adjusts cone signals from different surfaces by different amounts, so rather than achieving the required colour-constancy transformation, it shifts cone signals in each class towards a single reference value (Webster 1996). The effect of normalization to the local stimulation is illustrated in figure 6, and discussed in §4*b*(ii).

(iii) *Rank ordering*

To the extent that the rank order of cone signals is not disturbed by a change in illumination, reasonable constancy could be achieved by storing only the ranked position of each surface. Under this proposal (see figure 4), signals are mapped to the unit diagonal at the cost of losing information about the relative differences between surfaces. Primate photoreceptors in a steady state of adaptation respond reasonably linearly to small increases in light, but large step changes in light level produce saturating nonlinear responses (Schnapf *et al.* 1990). A benefit of considering only the rank order of signals is that this order remains unchanged under any monotonic transformation of the photoreceptor photon catches.

The idea that constancy may be based on rank order is paralleled in a large body of psychological literature on human judgements. Helson (1938, 1947) proposed the existence of a centrally stored *level of reference* that represents past and present environments, against which all new stimuli are judged. Rank ordering fails to achieve constancy if the reflectances sampled under

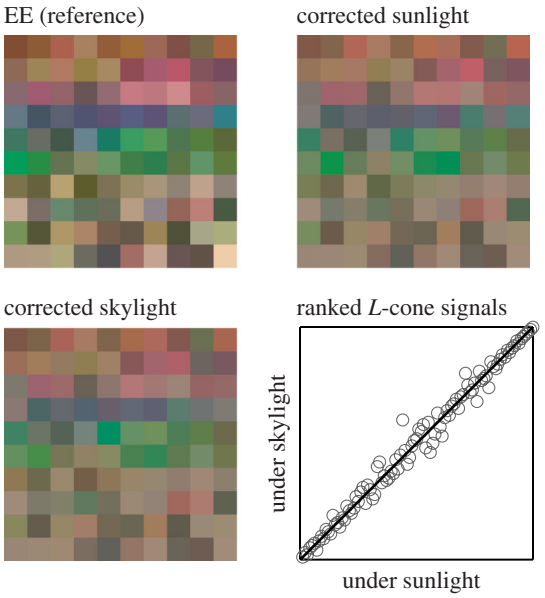

Figure 4. The four panels describe the effect of recoding the *L*-, *M*- and *S*-signals by their rank position, rather than using actual values. The images 'corrected sunlight' and 'corrected skylight' are transformed versions of images 'sunlight' and 'skylight' from figure 3*a*. The reference image is replotted for comparison. The plot (bottom right) shows that rank ordering achieves reasonable colour constancy since the transformed signals plot close to the unit diagonal, but information about the relative differences between surfaces has been lost.

the one illuminant are spectrally biased compared with those sampled under other illuminants (see also §4*b*(ii)).

(iv) *Coding colour relations by ratios*

The invariance of cone-excitation ratios is shown explicitly by Foster & Nascimento (1994, fig. 3). For randomly chosen pairs of surfaces, they plot the ratio of cone excitations for the two surfaces under one illuminant against the ratio obtained for the same two surfaces under a second illuminant (and, in fact, a different two illuminants are drawn at random from the daylight set for each pair of surfaces). In these plots, all data points fall very close to the unit diagonal.

If, for each cone class, the visual system encoded the spatial ratios of signals from different surfaces, this code could be used by observers to discriminate between scenes that changed in illumination and scenes that changed in reflectance: the code would be virtually unchanged by a change in illumination but would be disturbed by a change in the surfaces comprising the scene. It has been suggested that this signal might support *operational colour constancy*, i.e. the ability to distinguish between a change in illumination and a change in surface reflectance (Craven & Foster 1992; see §2*c*(i)). For ideal observers viewing uniformly illuminated Mondrian worlds, operational colour constancy is formally equivalent to colour constancy based on invariant colour percepts (see Foster & Nascimento 1994, Appendix 1; Foster *et al.* 1997). However, as indicated in §§1*a* and 2*b*, the two abilities may be dissociated in practice (Foster *et al.* 1997).

(v) *Post-receptoral combination of cone signals*

The preceding analysis of cone signals can be extended to subsequent stages of processing in the visual system. The MacLeod–Boynton (1979) chromaticity axes ($L/(L+M)$, $S/(L+M)$) provide a good representation of the post-receptoral colour signals that are transmitted to the cortex (Derrington *et al.* 1984). Zaidi *et al.* (1997) showed that when the effects of changes in illuminant spectrum are expressed in MacLeod–Boynton coordinates, the transformation of the $L/(L+M)$ chromaticities is approximately an additive transformation, whereas the transformation of the $S/(L+M)$ chromaticities is approximately a multiplicative transformation.

(c) *These schemes are not perfect. Are they good enough?*

(i) *Statistical measures*

Cone excitations for collections of real-world surfaces under different illuminants are not perfectly described by any of the simplifications above, and therefore, none of the suggested constancy transforms will result in perfect constancy. We are sometimes aware of violations of the invariance of cone-excitation ratios in the phenomenon of *metamerism*: two surfaces may match under one illuminant but not under another.[5] The analysis by Worthey (1985) suggests that metameric surfaces are in fact rare, but Foster & Nascimento (1994) show other violations of the invariance of cone-excitation ratios (for example, with some Munsell surfaces of extreme chroma). For two extreme illuminants, Nascimento *et al.* (2002) found mean relative deviations of cone-excitation ratios of around 4% for distributions of reflectances encountered in natural scenes, and deviations of around 9% for random sampling of the Munsell set. Dannemiller (1993) reports average absolute shifts in rank of 3.7, 3.1 and 2.9 for the *L*-, *M*- and *S*-cones, respectively, and a maximum shift of 26 positions out of the 337 materials he analysed from the Krinov set. However, the perceptual relevance of such statistics is unclear. Do the residual violations set the limits of human colour constancy? Or can we do better than these analyses suggest?

(ii) *Psychophysical measures*

Brainard & Wandell (1992) tested how well different linear transformations of cone coordinates accounted for observers' performance in an asymmetric colour-matching task. The mapping between the cone-coordinates of the test and the cone-coordinates of the match was well described by multiplicative scaling within cone classes (a diagonal linear transform). Moreover, the elements of the diagonal transform were linearly related to the change in simulated illuminant (i.e. the mapping obtained under an illuminant change that was the sum of two independent illuminant changes was predicted by the sum of the mappings measured for each illuminant change alone).

Nascimento & Foster (1997) required their observers to discriminate between simulations of real changes in illumination, and simulations of changes in illumination that were modified such that spatial cone-excitation ratios were preserved exactly.

The modified changes were identified as changes in illumination, even though they corresponded to illuminant transformations that are highly unlikely in the natural world. Furthermore, the greater the violations of invariance in the real transform, the higher the likelihood of misidentification. It is not totally clear that observers in this experiment actually equate the appearance of a modified illuminant change with the appearance of a natural illuminant change,[6] but what is clear is that observers are highly sensitive to violations of the invariance of spatial cone-excitation ratios, at least when the two images are presented in quick succession (Linnell & Foster 1996). Indeed, there is evidence that spatial cone-excitation ratios might be an elementary feature extracted from the visual scene (Westland & Ripamonti 2000; Foster *et al.* 2001; Hurlbert & Wolf 2004).

Spatially invariant cone-excitation ratios are strong predictors of the measured perceptual constancy of the *relations* between the colours of surfaces under changes in illumination (Foster *et al.* 1997). However, to what extent they support colour constancy in the stronger sense is still unknown. A difficulty here is that the task of asymmetric colour-matching requires only relational judgements. This is clearly the case when the spatial layout of surfaces is held constant under the two illuminants, but can be extended to cases in which the set of surfaces, or their spatial configuration, may change. Amano & Foster (2004) provide evidence for the use of cone-excitation ratios when the surfaces are rearranged, and they suggest that the most likely cue in this case is the ratio of cone excitations between the test surfaces and a spatial average over the whole pattern, a suggestion we will return to in §4*b*(ii).

If it could be demonstrated that human observers are more colour constant than predicted from the residual errors following multiplicative scaling of cone signals, or if failures of constancy are not consistent with errors in estimation of the parameters of such a transform, we must consider alternative simplifications of equations (1.1*a*)–(1.1*c*).

(d) *Colour conversion expressed in terms of linear basis functions*

(i) *A new transformation matrix*

A popular alternative is based on the claim that environmental illuminants and reflectances can be characterized as the weighted sum of a small number of basis functions (Sällström 1973). Measurements of real illuminants suggest that three basis functions may be sufficient (e.g. Judd *et al.* 1964; Wyszecki & Stiles 1982), and likewise for spectral reflectance functions (e.g. Cohen 1964; Dannemiller 1992). Maloney (1986) argues, however, that although the first three functions in Cohen's analysis accounted for 0.992 of the variance of the overall goodness of fit, there is considerable variation in how well individual surfaces are described, with the worst fits having notable patterned deviations. Maloney repeats the analysis for a larger dataset and finds that the number of parameters required to model the spectral reflectances is five to seven, not three. However, if the residuals used in deriving the fits are weighted by the photopic sensitivity curve, three or four parameters are sufficient.

If a surface reflectance, $R(\lambda)$, can be represented as $[b_1 r_1(\lambda) + b_2 r_2(\lambda) + b_3 r_3(\lambda)]$, where, $r_1(\lambda)$, $r_2(\lambda)$ and $r_3(\lambda)$ are the three basis functions, and b_1, b_2 and b_3 are relative weights, equation (1.1*a*)–(1.1*c*) become:

$$\begin{aligned} S = &\left[\int s(\lambda)E(\lambda)r_1(\lambda)\mathrm{d}\lambda\right]b_1 \\ &+ \left[\int s(\lambda)E(\lambda)r_2(\lambda)\mathrm{d}\lambda\right]b_2 \\ &+ \left[\int s(\lambda)E(\lambda)r_3(\lambda)\mathrm{d}\lambda\right]b_3, \end{aligned} \tag{3.1a}$$

$$\begin{aligned} M = &\left[\int m(\lambda)E(\lambda)r_1(\lambda)\mathrm{d}\lambda\right]b_1 \\ &+ \left[\int m(\lambda)E(\lambda)r_2(\lambda)\mathrm{d}\lambda\right]b_2 \\ &+ \left[\int m(\lambda)E(\lambda)r_3(\lambda)\mathrm{d}\lambda\right]b_3, \end{aligned} \tag{3.1b}$$

$$\begin{aligned} L = &\left[\int l(\lambda)E(\lambda)r_1(\lambda)\mathrm{d}\lambda\right]b_1 \\ &+ \left[\int l(\lambda)E(\lambda)r_2(\lambda)\mathrm{d}\lambda\right]b_2 \\ &+ \left[\int l(\lambda)E(\lambda)r_3(\lambda)\mathrm{d}\lambda\right]b_3. \end{aligned} \tag{3.1c}$$

And this set of equations can be rewritten in matrix form as:

$$\begin{bmatrix} S \\ M \\ L \end{bmatrix} = \begin{bmatrix} \varepsilon_{s,1} & \varepsilon_{s,2} & \varepsilon_{s,3} \\ \varepsilon_{m,1} & \varepsilon_{m,2} & \varepsilon_{m,3} \\ \varepsilon_{l,1} & \varepsilon_{l,2} & \varepsilon_{l,3} \end{bmatrix} \begin{bmatrix} b_1 \\ b_2 \\ b_3 \end{bmatrix}. \tag{3.2}$$

To the extent to which the basis functions capture the properties of surface reflectances in the world, colour constancy is thus reduced to recovering b_1, b_2 and b_3, from the cone signals by undoing the transformation imposed by the illuminant (represented by matrix ε, whose elements are given by the expressions in square brackets in equations (3.1*a*)–(3.1*c*). This is certainly a simplification of equations (1.1*a*)–(1.1*c*), although the inverse of matrix ε can be recovered only by imposing further constraints.

(ii) *Additional constraints*

If, for example, the illuminant is a weighted sum of three known spectral power distributions, an object of known spectral reflectance present in the scene would provide all the information necessary to find the three unknown weights of the illuminant (see equation (3.1*a*)–(3.1*c*). If additionally, the reflectance of each object could be specified as the weighted sum of three known spectral reflectances, the three cone signals provide simultaneous equations that can be solved to recover the three unknown weights of the reflectance (Buchsbaum 1980). If one assumes that the average spectral reflectance over all objects in the scene has a known distribution (a proposal we will return to in §4*b*(ii)), then no single reference standard is required.

Alternative linear basis models of colour constancy relax one constraint at the expense of others. For example, the illuminant can be any arbitrary light (not just a weighted sum of three spectral power distributions), provided that there are three reference

standards rather than one (Brill 1978; Brill & West 1986). Maloney & Wandell (1986) have shown that, with trichromatic vision, the need for a reference standard can be eliminated if the surface reflectance of each object in the scene is a weighted sum of two rather than three basis functions, and if information can be collected from three or more surfaces. Their 'sub-space' model (Maloney & Wandell 1986) obtains an estimate of the illuminant from the way the set of sampled surfaces clusters in three-dimensional cone-excitation space. If surfaces are described by only two basis functions, they will lie exactly in a plane, whose location within the space is defined by the illuminant (see Brainard *et al.* 1993 for an intuitive description of why this will work). Second-order statistics of the cluster, like the best-fitting plane, are more reliable than the mean as cues to the illuminant, both with respect to changes in the set of surfaces sampled, and to surfaces that are not well described by a limited number of basis functions.

(iii) *Sensor sharpening*

The matrix equation above could be expressed in terms of the spectral sensitivities of any arbitrary system of three sensors. A common goal in computational studies of colour constancy has been to determine the conditions under which matrix ε might be diagonal (for review, see Hurlbert 1998). This is appealing since constancy could then be implemented by independent scaling within each sensor class. Equations (3.1*a*)–(3.1*c*) show that a sufficient condition for such a constancy mechanism would be that the spectral sensitivities of the three sensors were perfectly matched filters for the basis functions of surface reflectance. However, spectral sensitivity functions of human cones are not matched to the first three basis functions required for typical sets of naturally occurring surfaces.

An alternative is that the visual system applies two transformations (D'Zmura & Lennie 1986; Hurlbert 1986; Finlayson *et al.* 1994). Finlayson *et al.* (1994) generate 'sharpened' spectral sensitivities by opponent combination of the *L*- and *M*-cone spectral sensitivities prior to the operation of a diagonal transformation. D'Zmura & Lennie (1986) propose a diagonal transformation followed by a linear transformation that recodes the scaled cone signals into channels that correspond to the basis functions for surface spectral reflectance. In their scheme, the output of the combined transformation provides a description of surface reflectance and is represented at the colour-opponent stage of the primate visual system.

(e) *What are the constraints on illuminants, reflectances and spectral sensitivity functions that predict simple receptor scaling?*

From the above discussion, it is clear that the empirical result of approximate multiplicative scaling of cone coordinates under an illuminant change (which effectively reduces the colour conversion to a diagonal matrix, and predicts that the triad of cone signals provides an adequate description of surface reflectance) is not directly predicted from the assumption that surfaces and illuminants can be described by a small number of linear basis functions; to quote D'Zmura & Lennie, 'the three numbers used in scaling cone signals cannot undo what it takes nine numbers to describe' (D'Zmura & Lennie 1986, p. 1667).

So why does the empirical result hold? Detailed analyses of the way this result depends on the relationships between reflectance spectra, illuminant spectra and receptoral sensitivities are provided by several authors (e.g. Maloney 1986; Worthey & Brill 1986; Dannemiller 1993; Zaidi 2001). But their conclusions differ. While we can be certain that the human visual system makes use of regularities in the world to simplify equations (1.1*a*)–(1.1*c*), we do not yet know the exact form of the simplification, or the constraints that make it possible.

4. ESTIMATING THE ILLUMINANT

(a) *What is estimated?*

The Ives-transform relies on the visual system's ability to estimate the cone coordinates of the illuminant. Similarly, solutions to equation (3.2) require, at some stage, an estimate of the illuminant, although here it is specified in terms of basis function coefficients. However, the illuminant itself is often not in the field of view so its coordinates must be estimated from cues within the visual scene (e.g. Maloney & Yang 2003). Suggestions for how this might be achieved have a very long history (for review, see Mollon 2003). Their popularity may at times have derived from the appeal of the 'over-simplification' described in §3*a*, but the work described in §§3*b*–3*d* provides both a justification for the role of illuminant estimation in human colour constancy, and shows its limitations.

In a study of lightness constancy, Rutherford & Brainard (2002) tested a version of the illuminant estimation hypothesis in which the illuminant estimate is associated with the explicitly perceived illuminant, and found it to be false. This raises the intriguing possibility that the same physical quantity has multiple psychological representations (see Bhalla & Proffitt 1999; Gilchrist *et al.* 1999; Rutherford & Brainard 2002).

The term *estimating the illuminant* is, in many senses, misleading. Since the colour conversion between two conditions of illumination is determined by the properties of the two illuminants, the inverse transformation that achieves constancy must be related to them. However, the illuminant estimate need not be explicitly available to the observer, or coded explicitly during any stage of the computation. This distinction is seen most easily for transformations that rely on multiplicative scaling of cone signals. To perform the Ives transform, the multiplicative constants for each cone class should be set in the ratios $L_E : L_{E'}$, $M_E : M_{E'}$, $S_E : S_{E'}$, respectively, where E and E' are the test and reference illuminants (see figure 3). These quantities can be derived from the cone signals from a perfectly reflecting surface, but approximately the same ratio could be extracted by using the cone signals from any other surface that was identified under the two illuminants. In this case, neither surface would provide an estimate of the illuminant, but the computation would reveal the parameter required to undo the illuminant transform.

(b) *Cues to the illuminant*

(i) *Reference surfaces*

The diagonal transformation of cone coordinates imposed by a change in the illuminant can be identified by finding the mapping between corresponding sets of objects sampled under the two illuminants. Zaidi (1998) presents two template-matching algorithms to achieve this: one that requires the set of surfaces to remain the same (and which essentially reduces to normalization to the global mean) and one that requires only a subset of surfaces to remain the same. If the correspondence of just a single reference surface is known, the correct transformation can be estimated directly. Similarly, in the computational literature, a frequent suggestion is that the unknown elements in matrix ε might be recovered by using information from standard reflectances within the scene (Brill 1978; Buchsbaum 1980; Brill & West 1986).

This is a viable operation for the human visual system only if observers can identify reference surfaces within a complex image (see §4*d*(iii)) and store and access descriptions of reflectance (Bramwell & Hurlbert 1993; Fairchild 1993). Jin & Shevell (1996) have presented empirical evidence that observers tend to remember the 'reflectance' (or at least an illuminant-invariant description) of surfaces presented under different illuminants, rather than the chromaticity of the light that reaches the eye.

(ii) *Grey world*

A very common suggestion is that the illuminant is estimated from the mean cone coordinates of the global scene (Helson 1964; Buchsbaum 1980; Land 1983, 1986). Land's retinex model (Land & McCann 1971; Land 1983, 1986) extends von Kries adaptation to include spatially extended information. For each element in the cone mosaic, the model computes 'lightness values' (adjusted cone signals) from paths within the image that intersect only cones of the same class. Both the number and the length of the paths are parameters in the model. In the limiting case of arbitrarily many and arbitrarily long paths, the normalization factor approximates the geometric mean of responses from all cones of the same type (Brainard & Wandell 1986).

Normalizing the cone signals for all surfaces to the mean along each axis in figure 2 would rescale the points to the unit diagonal,[7] and other population statistics, such as the geometric mean, or the standard deviation could also be used. The 'grey world hypothesis' formally identifies a stable mean reflectance as spectrally uniform, and thus becomes an algorithm for extracting an unbiased estimate of the illuminant. This assumption is unlikely to be true for most scenes (Brown 1994; Webster & Mollon 1997). However, *constancy* requires only that the average reflectance remains constant, not that it be spectrally uniform. Furthermore, the average need not include the whole scene: the ground-plane or other potentially less variable portion of the scene could serve as the reference (Maloney 1999). Morgan *et al.* (2000) have shown that observers in a discrimination task can decide on the basis of a symbolic cue, which of several implicit reference standards is the correct one, with only a moderate loss of precision.

If the average reflectance was not stable and was different for the two illuminants, normalization to the mean would not achieve the required result. An inappropriate normalization of this type is illustrated in figure 5. Similarly, if the mean is collected over an area that is small compared with the relevant spatial variation in the scene, differences between regions will be lost. The effect of normalizing each region to the local mean is illustrated in figure 6 (see also the discussion of the von Kries transform in §3*b*(ii)).

(iii) *Second-order statistics*

Golz & MacLeod (2002) have suggested that luminance–chromaticity correlations within an image may provide illuminant estimates that are less influenced by the set of reflectances available. Moreover, they present data which suggest that human observers do use such a cue and give it a weighting that is statistically appropriate for the natural environment. One drawback of using the complete range of samples in the scene is that darker surfaces may contribute more noise than signal to the estimate. Tominaga *et al.* (2001) argue that for this reason it is better to use only the brightest objects.

(iv) *Brightest is white*

Early versions of the retinex algorithm (Land 1964) assumed that the brightest patch in each spectral channel has 100% reflectance in that channel (although a different patch might be used for each channel). In natural scenes, and in many experimental situations, the space-averaged *L*-, *M*- and *S*-cone signals covary with the maximum in each cone class, and with the illumination. In these cases, observers could use either cue to set the parameters of the colour-constancy transform.

McCann and colleagues have performed several experiments to determine empirically which cue dominates. In general, the experiments have been performed with a small number of real surfaces illuminated with narrowband lights. McCann (1997, 1992) describes a series of 'destroy the match' experiments. Two sets of five surfaces (colour 'tautomi' displays) that had the same relative reflectances but different absolute values for each waveband were observed in a restricted field of view and illuminated such that the two displays produced equivalent cone signals. At this point, since the displays were the same, they appeared the same. However, when a new white surface was introduced in one field, the match was destroyed. Furthermore, a new maximum was found to reset the appearance even when an equal area of low reflectance was introduced to hold the average constant.

Linnell & Foster (1997, 2002) performed similar experiments. Observers were asked to make matches of illumination across patterns (7° of visual angle) in which the global mean and the brightest patch were chosen to predict conflicting illuminants. With very small patches (0.03° of visual angle), illuminant estimates were set by the global mean, as expected. The brightest patch had an effect only for the largest

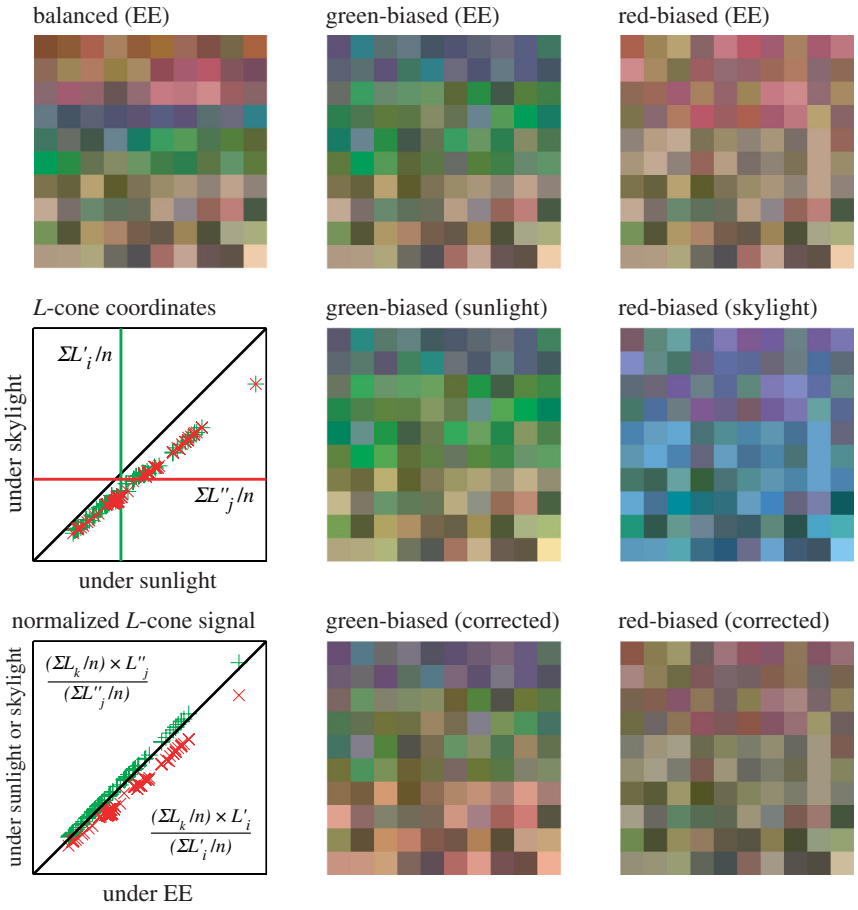

Figure 5. The three images in the top row show three scenes under an equal energy illuminant. From left to right, the first image is a chromatically balanced reference (with a mean that is approximately spectrally uniform); the upper portion of the second image is green biased; and the upper portion of the third image is red biased. The lower portion of all three images is identical. Images in the second row show the green-biased scene under sunlight, and the red-biased scene under skylight. The plot on the left shows the *L*-cone signals for the green-biased set of reflectances (green plus symbols), and for the red-biased set of reflectances (red crosses). The signals obtained under skylight are plotted against the signals obtained under sunlight. The horizontal red line marks the mean *L*-cone signal from the red-biased scene (*n* reflectances indexed with subscript *j*) under skylight (indicated with a double prime symbol), and the vertical green line marks the mean *L*-cone signal from the green-biased scene (*n* reflectances indexed with subscript *i*) under sunlight (indicated with a signal prime symbol). The bottom row shows images corrected by the global mean of these scenes compared with the global mean of the reference scene. In each image, the relative difference between the top and bottom halves is reasonably well preserved, but the previously neutral bottom halves are too red in the green-biased scene, and too green in the red-biased scene. The plot on the left shows the normalized *L*-cone signals. The signal from each surface in the red-biased scene under skylight (L''_j) has been normalized by the mean from that scene ($\Sigma L''_j/n$), multiplied by the mean from the reference scene ($\Sigma L_k/N$), and plotted against the corresponding signal from the reference scene. The signal from each surface in the green-biased scene under sunlight (L'_i) has been normalized by the mean from that scene ($\Sigma L'_i/n$), multiplied by the mean from the reference scene ($\Sigma L_k/N$), and plotted against the corresponding signal from the reference scene. The signals from the red-biased scene are under-corrected. The signals from the green-biased scene are over-corrected.

patches (1° of visual angle). Linnell & Foster conclude that, for flat, richly sampled, Mondrian-type stimuli the global mean is the dominant cue to the illuminant. Had the patterns contained cues to three-dimensional shape, the results may have been different.

(v) *Chromaticity convergence*

Pure specular highlights are the extreme example of bright, 'white' elements in an image, and as such, they provide a direct glimpse of the illuminant. However, specular reflections also give rise to a more subtle cue. The light reflected from an inhomogeneous material is the sum of two components: the 'interface reflection', which has the same spectrum as the illuminant, and the 'body reflection', which is a wavelength-by-wavelength multiplication of the object and illuminant spectra (Shafer 1985). The relative contribution of the two components depends on the viewing geometry, so the chromaticities of light reflected from different regions of a glossy surface will fall on a line joining the body reflection and the illuminant colour. Lines corresponding to several surfaces in the same scene (illuminated by a single source) will therefore intersect at the illuminant chromaticity (D'Zmura & Lennie 1986; Lee 1986; Lehmann & Palm 2001). The suggestion has been named the *chromaticity convergence* algorithm by Hurlbert (1998).

Chromaticity convergence was anticipated by Monge (1789; Mollon 2003). He presented a seeming paradox: if a complex scene is viewed through a red

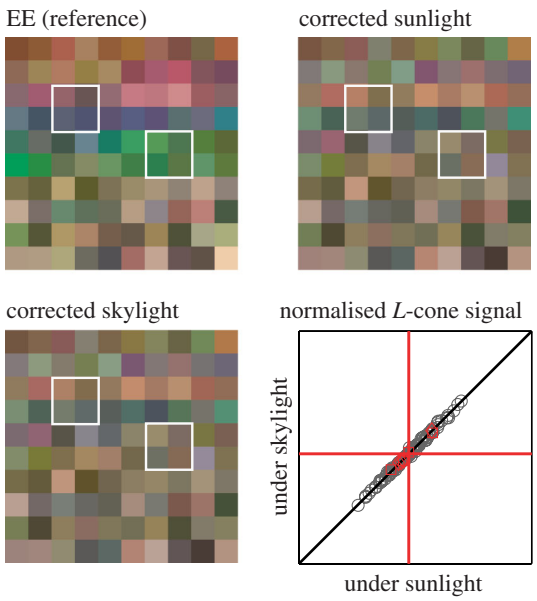

Figure 6. The four panels describe the effect of normalizing the cone signals by their local mean (i.e. the original von Kries transform). The *L*- (or *M*- or *S*-) signals from each group of four small squares were normalized to the mean *L*- (or *M*- or *S*-) value of those four small squares. The images 'corrected sunlight' and 'corrected skylight' are transformed versions of images 'sunlight' and 'skylight' from figure 3*a*. The reference image is replotted for comparison. The differences between the images from the two illuminant conditions have been substantially reduced, but the variations within each scene that occur at a larger spatial scale than the normalization (e.g. the orange, red, blue and green bands in the upper portion of the reference image) have been lost. The plot (bottom right) confirms that the normalization achieves approximate constancy. The red symbols (diamonds and squares) in this plot correspond to the surfaces in the local patches outlined in white in the two images. In the limiting case where the normalization is performed over a local area that contains only a single surface, this transformation maps cone signals for all surfaces to the point (1,1), and thus achieves constancy only by eliminating differences between surfaces.

filter, red objects appear, not a saturated red, but desaturated or white. In a stunning piece of logical argument, Monge relates this paradox to specular reflections. He argues first that all objects send to the eye a component of light that is determined by their surface colour, and a component that is determined by the illuminant. White objects are unique in sending to the eye the same spectrum of light from every point on their surface. When viewed through a red filter, red objects acquire this characteristic, and thus appear white.

(vi) *Mutual reflections*

A surface with reflectance $R_1(\lambda)$ that is illuminated by $E(\lambda)$ will re-emit light with a modified spectrum $\int E(\lambda)R_1(\lambda)\mathrm{d}\lambda$, which may in turn be incident on a second surface $R_2(\lambda)$, and be re-emitted with a modified spectrum $\int E(\lambda)R_1(\lambda)R_2\mathrm{d}\lambda$, and so on. Drew & Funt (1990) have demonstrated, at least for the 'single-bounce' case, how such a complication may be exploited in estimating the illuminant (see also Funt *et al.* 1991; Bloj *et al.* 1999).

(vii) *Shadows*

D'Zmura (1992) presents a generalization of the 'subspace' model (Maloney & Wandell 1986) in which multiple views of the scene under different illuminants are used to provide additional information to solve generic forms of equation (3.2), in which illuminants and surfaces are represented by m and n linear-basis functions, respectively (D'Zmura & Iverson 1993*a*,*b*, 1994*a*). A shadow boundary across a surface will provide the multiple views the algorithm requires. An alternative proposal is that observers could exploit the spatial distribution of macular pigment across the retina. This spectrally selective pre-receptoral filter peaks in density around the fovea, falling to very low levels beyond approximately 8° of visual angle, and will thus provide multiple views of the scene under differently filtered illumination (Broackes 1992).

(viii) *Colour by correlation*

Finlayson *et al.* (2001) present a correlation framework within which to consider illuminant estimation algorithms. Like other authors (e.g. Forsyth 1990; D'Zmura *et al.* 1995; Brainard & Freeman 1997), they recognize that the problem of illuminant estimation may not be sufficiently constrained to provide a unique solution, and that a pragmatic approach (for the visual system, and for computational algorithms) is to find a set of possible solutions and to search for the best one. In their scheme, the problem of illuminant estimation is represented as finding the correlation between the colours in an image and prior knowledge about the probability with which different colours are observed under different lights. A thresholding procedure can return the most likely illuminant or set of illuminants. The strength of Finlayson *et al.*'s approach is that it allows the assumptions of different constancy algorithms to be compared. Different algorithms amount to different methods of computing the correlation matrix. The grey-world hypothesis, for example, states that each colour that is observed in the image is equally likely to have been the illuminant, and that the best estimate is obtained by taking the average of the observed values. This clearly does not capture our (implicit) knowledge of the interaction between surfaces reflectance and illumination (see also §4*b*(iii)). For machine vision, the correlation matrix is constructed during an initial process of sensor calibration. Recent data from infant monkeys (Sugita 2004) suggest that experience in early infancy with broadband lights and surfaces is vital for the development of normal colour constancy. It is appealing to interpret this as evidence for a calibration phase in human colour constancy.

(c) *Limited availability of cues*

(i) *How many surfaces?*

With only a single uniformly illuminated surface in view, it is impossible to disambiguate illumination and surface reflectance; with multiple surfaces, estimation of the illuminant becomes possible. Interestingly, observers in a surface-matching task perform nearly as well with spatially congruent 'scenes' containing only two samples of spectral reflectance, as they do with scenes containing many samples (Blackwell &

Buchsbaum 1988; Arend *et al.* 1991). This result suggests that performance in such tasks is mediated by purely relational constancy, rather than by a process of illuminant estimation and subsequent discounting (Foster *et al.* 2001; Zaidi 2001; Foster 2003).

Linnell & Foster (2002) have asked how operational colour constancy depends on the number of surfaces in a scene. Observers' ability to detect a change in illumination over two scenes, containing different random samples of reflectance, improved as a function of the number of patches in the scene (from 9 to 49).

Smithson & Zaidi (2004) measured boundaries between colour categories (red versus green, and yellow versus blue) as a function of the illuminant on a sequence of *single* test patches, with a conflicting illuminant on surrounding patches. In a single trial there was no information about the test illuminant, since this fell only on a single surface. However, colour boundaries were more closely predicted by the test illuminant (estimated over time) than by the surround illuminant (estimated across space). The recent history of reflectances sampled by the observer is a primary determinant of colour appearance. The history of reflectances might be taken from successive presentations (as in this study), or from successive fixations within a steady image. Spatially distributed cues to the illuminant are carefully specified in studies and models of colour constancy. This study highlights the importance of also considering temporally distributed cues.

(ii) *'Flat-world', 'shape-world' or real-world?*

The simplified physical world ('flat-world'; Maloney 1999) in which the illuminant is spatially uniform, and surfaces are flat and matte and engage in no mutual reflections (and for which equations (1.1*a*)–(1.1*c*) describe the cone signals from surfaces) has been used extensively in colour-constancy experiments. However, this world is impoverished compared with a world that includes cues from specularity, mutual reflections and shadows ('shape-world'; Maloney 1999), which in turn is severely impoverished compared with the world in which we live (real-world) that additionally includes multiple light sources and transparency. The flat-world has an important place in colour constancy research since scenes from such a world are straightforward to simulate, manipulate and control. Recently, the visual world of colour-constancy experiments has been enriched through the use of sophisticated physics-based rendering software (Yang & Maloney 2001; Maloney 2002) and through the use of real objects and computer-controlled lighting systems (Brainard *et al.* 1997; Brainard 1998).

(iii) *Real or simulated?*

The level of colour constancy achieved by human observers is typically less for simulated scenes than for real scenes. Brainard (1998) found that with real scenes, observers can compensate for 84% of the change in illumination (assessed via achromatic settings), while typical performance with scenes presented on computer monitors suggests only 50% compensation. Yang & Maloney (2001) took several steps to ensure that their simulated scenes were as real as possible, and their observers achieved achromatic settings that compensated for 65% of the change in illumination. It is clear that there are subtle cues in the real world that we do not yet understand, and cannot yet simulate.

(d) *Perceptual and cognitive factors*

(i) *Combination of cues*

It is likely that the visual system makes use of several sources of information. Much of the current work on colour constancy aims to discover what weights the visual system gives to different cues under different natural settings. Kraft & Brainard (1999) used real objects (rather than rendered images of objects) and 'silenced' some of the individual cues in a semi-naturalistic setting. Their subjects exhibited poorer and poorer constancy as cues were successively reduced.

Maloney (2002) highlights a possible complication, for the human visual system may dynamically assign different weights to different cues, depending on which cues are available, or on the basis of task demands, or prior knowledge. When the scene is rich in reliable cues, eliminating one of the cohort may have little effect on the illuminant estimate since the shortfall may be taken up by the remaining cues. Maloney suggests that to understand cue-weighting we might use the method of perturbation analysis, developed in studies of depth and shape perception. In this method, cues are not removed, but instead are perturbed to signal a conflicting illuminant. Although this is difficult to achieve in real scenes, Yang & Maloney (2001) describe an experiment in which they perturbed cues in a simulated scene that was rendered with physics-based graphics software.

(ii) *Segmentation of the image*

In order to use shadows, mutual reflections or specular highlights to identify the illuminant, the visual system must first identify these components within a complex image. This is the problem of *image segmentation*. The visual system must parse the two-dimensional image on the retina into regions that correspond to distinct three-dimensional objects that cast shadows, reflect light on to one another and offer specular reflections. Hurlbert (1998) provides a full discussion of the relationship between algorithms for colour constancy and algorithms for colour segmentation.

The process can fail when the visual system chooses the wrong perceptual hypothesis to interpret a given input signal. Such illusions have been analysed more frequently in the lightness literature than in the colour constancy literature, perhaps because with lightness, the form of the illuminant transformation is well known (see §3*a*). Striking demonstrations of the role of perceptual organization in lightness constancy are provided by Adelson (1993; see also Adelson & Pentland 1996; Gilchrist *et al.* 1983 for discussion).

(iii) *Memory colour*

Some objects have a characteristic colour. Bananas are yellow, grass is green, and blood is red. Hering (1874, 1964) has suggested that the characteristic colour of an object is an important factor in colour constancy, but empirical tests of this claim have yielded conflicting

results (see Beck 1972 for discussion). Clearly, to use the characteristic colour of an object to anchor colour percepts, observers must have extracted an illuminant-independent description of the colour of an object, remembered that description, and subsequently be able to identify examples of such objects in complex scenes under different illuminations.

5. WHERE ARE THE COLOUR CONSTANCY MECHANISMS?

(a) *Scope*

The highly systematic nature of transformations of cone coordinates under an illuminant change implies that simple neural transformations could support colour constancy. The types of neural mechanisms that could, in principle, contribute to such processing range from automatic to volitional and from peripheral to central.

Naturally, this discussion is intrinsically linked to the types of computation that must be performed (see §3), and to the way in which the parameters of the transform might be set by the image (see §4). However, when asking which neural mechanisms are responsible for colour constancy we must additionally consider the conditions under which colour constancy has been assessed. Colour constancy is influenced by the task at hand, and by the instructions given to observers (see §2). Furthermore, different neural mechanisms are required for colour constancy exhibited over different time-scales (seconds, hours, days or even years).

(b) *Early sensory mechanisms*

(i) *Processes of adaptation*

How do the types of transformation imposed on cone coordinates by a change in the illumination correspond to early adaptation mechanisms? In his coefficient law, von Kries (1878, 1905) identified receptor scaling as the mechanism that maintains the constancy of colour appearance under adaptation. Receptor scaling is implicit in Stiles' two-colour increment-threshold technique (Stiles 1939, 1949), and is discussed explicitly by Rushton (1972). It is widely accepted as part of the process of visual adaptation.

Current models of the early stages of human vision propose that the signals from the cones are recombined into two chromatically opponent channels and one or more achromatic channels. In addition to first-site modification of signals within cone classes (e.g. von Kries 1905) a second-site modification of the opponent signals (e.g. Pugh & Mollon 1979) is also identified. Processes of adaptation are likely to be nonlinear if considered over a wide dynamic range, but given only moderate excursions from a constant level, both single-cell responses and psychophysical sensitivity are consistent with transformations of the cone signals that have a similar form to matrix ε (Maloney (1999) discusses the relationship in more detail). D'Zmura & Lennie (1986) explicitly link their model of colour constancy to the two sites of adaptation identified above. Importantly though, the work discussed in §§3*b* and 3*d* does not *require* the constancy computations to occur in the receptors themselves or at any particular point in the visual pathways.

(ii) *Spatial and temporal extent*

The spatial extent of colour constancy mechanisms is not known. The grey-world hypothesis clearly requires mechanisms that collect information from an extended area. Kraft & Brainard (1999) have shown that remote surfaces have a significant effect on appearance even when pitted against local contrast. Shevell & Wei (1998) have shown that contrast in a remote region markedly influences colour induction in a chromatic matching task. On the other hand, when observers are asked to null or discriminate the colour change induced during a short (500–800 ms) presentation of a surround stimulus, remote fields contribute relatively little to appearance (Wachtler *et al.* 2001; Wolf & Hurlbert 2003; Barbur *et al.* 2004).

Published psychophysical measurements indicate that early adaptation mechanisms are extremely local in their spatial properties (MacLeod *et al.* 1992; MacLeod & He 1993; He & MacLeod 1998). This limitation might at first suggest a role for central mechanisms with large receptive fields. However, there are other ways in which spatially extended information could be collected. Eye movements convert spatial variations into temporal variations, so local mechanisms do in fact sample information from spatially dispersed points in a scene (D'Zmura & Lennie 1986; Fairchild & Lennie 1992). Obtaining illuminant information from such samples would require these local mechanisms to respond with long time-constants. Although a significant change in sensitivity occurs within approximately 200 ms of a change in background level (Crawford 1947), adaptation may not be complete for several seconds, and this is especially true for adaptation at the second site (Pugh & Mollon 1979).

A very different suggestion for how early sensory mechanisms might integrate information from the entire scene was made by Mollon *et al.* (1998). The human retina is encircled at the ora serrata by a cone-rich rim. The function of these cones is uncertain but it is possible that they may integrate light scattered within the globe of the eye, or passing through the sclera.

(iii) *Information about the illuminant is not discarded*

The main argument against a theory of constancy based only on distal gain adjustments has been eloquently described by Katz, 'Paradoxical as it may at first sound, such a thoroughgoing efficiency on the part of the adaptive mechanisms as Hering postulates is not even to be considered as desirable; for it would partially or wholly compensate for any change in illumination, and thereby make it imperceptible' (Katz 1935, p. 265). Under some circumstances, observers do seem to have access to both the surface colour and the illuminant colour, although colour *scission* of this sort is particularly dependent on the geometrical properties of the scene (D'Zmura & Iverson 1993*a*,*b*, 1994*a*; Hagedorn & D'Zmura 2000). Further evidence against automatic and complete peripheral adjustments is that colour-appearance judgements can be influenced by the set of surfaces requiring judgement, and not only by the temporal and spatial statistics of the stimuli (Smithson & Zaidi 2004).

(c) *Cortical mechanisms*

Lesion studies, and more recently, neuroimaging studies, have suggested functional segregation and specialization of the human cortex. The search for a 'colour area' or even a 'colour constancy area' has attracted considerable effort. In an influential study, Land *et al.* (1983) tested the colour constancy of a patient with complete resection of his corpus collosum. Since the patient's speech centre was in his left hemisphere, he was able verbally to describe things in his right visual field, but not in his left visual field. When a Mondrian was presented in his right visual field, his reports of the appearance of a centrally presented test-patch were consistent with those made by normal observers. However, when the Mondrian was presented in his left visual field, his reports of the same centrally presented test-patch paralleled those that a normal observer would give if the test-patch were seen in isolation. Land *et al.* argue that the long-range constancy computations that presumably set the appearance of the centrally presented test-patch for normal observers, could not occur in the retina, and must occur in the cortex where regions representing the two halves of the visual field are joined by the corpus collosum. Ruttiger *et al.* (1999) presented further evidence that cortical computations are essential for colour constancy. Five patients with circumscribed unilateral lesions in parieto-temporal cortex of the left or right hemisphere exhibited a selective loss of colour constancy (assessed via achromatic settings), while their colour discrimination thresholds and colour associations for familiar objects were preserved.

Zeki (1983*a*,*b*) has shown that area V4 of macaque visual cortex contains many cells whose response parallels more closely 'perceived colour' rather than the wavelength composition of the stimulus. He and his collaborators have also demonstrated that damage to human cortical area V4 (see Wandell & Wade 2003 for discussion of the distinction between macaque V4, human V4 and V8) can result in a failure of constancy in a colour-naming task (Zeki *et al.* 1999). The receptive fields of V4 neurons are well suited to a constancy computation since they have a large, silent, suppressive surround around the classic receptive field (see Hurlbert & Poggio (1988) and Schein & Desimone (1990), for modelling and electrophysiology, respectively).

However, V4 is not wholly devoted to computing surface colour (Schiller & Lee 1991), nor is it the only neural area required for such computations. Areas V1 and V2 contain cells that respond not directly to the hue of a stimulus but to combinations of hue and surrounding luminance (Yoshioka *et al.* 1996). Hurlbert *et al.* (1998) found that a cerebrally achromatic observer (with bilateral damage to the ventral and temporo-occipital regions including the lingual and fusiform gyri; Heywood *et al.* 1991) could discriminate changes in cone-excitation ratios for simple but not complex scenes and was thus able to display colour constancy in an asymmetric matching paradigm. Barbur *et al.* (2004) used a dynamic matching technique to quantify changes in the appearance of a central test-patch as a function of rapid changes in the illumination of a surrounding Mondrian pattern. Data from a patient with unilateral damage to V1 suggested that retinal mechanisms could not support constancy in this situation. However, data from two patients with cerebral achromatopsia showed evidence for constancy mechanisms (revealed as a modification of chromatic discrimination thresholds measured along the line joining the two illuminants). The magnitude of the simultaneous constancy effect was smaller than the observer's chromatic discrimination threshold, and would thus remain hidden in other paradigms. Barbur *et al.* suggest that mechanisms that support colour constancy under instantaneous changes in illumination may occur in V1, with only a small contribution from extrastriate areas.

Colour constancy is not a unitary process, achieved by all or nothing computation. Wandell *et al.* (1999) argue that rather than looking for functional segregation of a 'colour constancy area', imaging techniques might be most usefully employed to track the transformation of visual information along specialized pathways. It seems most reasonable to say that processing for colour constancy starts in the retina, is enhanced in V1/V2 and continues in V4 (see also Walsh 1999).

6. SUMMARY

Our current understanding of colour constancy is patchy, at best. A full understanding would require detailed knowledge of the physical world of illuminants and surfaces, and of the biological and psychological worlds of our sensory and cognitive processes. Sophisticated analyses of the colour conversions imposed by a change in illumination show that reasonable constancy could be obtained from a variety of neural transformations. These range from independent multiplicative scaling, or rank-ordering of receptor responses, to more complex transforms of cone signals that require interactions between receptor classes. Observers are exquisitely sensitive to violations of the invariance of cone-excitation ratios, at least when the two images are presented in quick succession (Foster *et al.* 2001), but where and how our neural systems compute these ratios we do not know, and their role in predicting appearance-based phenomenological colour constancy remains uncertain. There are two significant hurdles in constructing critical experiments to distinguish the hypotheses.

First, there is the problem of generating the right stimuli. As we have seen, 'illuminant estimation' can proceed via several cues. Given the superiority of constancy measured with real scenes compared with simulated scenes, there are presumably additional cues that we do not yet understand. Brainard and colleagues have suggested that when confronted with impoverished stimuli, the colour constancy mechanisms that normally operate in the real world may produce unstable or conflicting results, and they therefore advocate using real scenes to study constancy under nearly natural viewing (Brainard *et al.* 2003). Maloney, on the other hand, argues that our understanding would be better advanced by accurate simulations of (unnatural) environments that are defined according to

the mathematical idealizations which form the basis of competing theories of constancy (Maloney 2003).

The second problem is to choose the appropriate task for the observer. If asymmetric matches can be solved via relational judgements, the stability of colour appearance must be assessed via alternative techniques. However, the options are limited. Colour naming is of limited resolution and achromatic settings track only a single point in colour space (Foster 2003). Measurements of boundaries between colour categories (e.g. red versus green, and yellow versus blue), rather than a single achromatic point, may provide a means of tracking the transformations of perceptual colour space under an illuminant change (Chichilnisky & Wandell 1999; Smithson & Zaidi 2004).

It is likely that multiple constancy computations are performed in parallel. The different computations may perform different transformations and use different subsets of cues to set the parameters of those transformations. The observer's perceptual experience may depend on what they are doing and their report may depend on what they think we want to know.

I would like to thank Qasim Zaidi and John Mollon for many helpful discussions, and for their comments on an earlier version of the manuscript.

ENDNOTES

[1]A second kind of colour constancy, which is only touched upon here, acts to discount the effects of simultaneous colour contrast and thus preserve the constancy of colour appearance in the presence of different surrounding colours or different backgrounds. For discussion, see Whittle (2003).

[2]Initially, we will consider a simplified physical world in which the illuminant is spatially uniform, and in which surfaces are flat, matte and engage in no mutual reflections (although see §§4*b*(v)–4*b*(vii)). In such a world, the fraction of incident light reflected at each wavelength is not altered by the geometry of the surfaces, light source or observer.

[3]Cone coordinates (L, M and S in equations (1.1*a*)–(1.1*c*)) are inferred from multiplication of the cone spectral sensitivities and the reflected light. Primate photoreceptors in a steady state of adaptation respond reasonably linearly to small increases in light intensity (Schnapf *et al.* 1990) so, within limits, these coordinates are assumed to be linearly related to the cone signal. This issue is discussed further in §3*b*(iii).

[4]Ives' diagonal transform has been widely analysed in the computer vision literature where it is frequently misnamed the 'von Kries transform'.

[5]In general, a match will be obtained if the triad of cone signals for one surface is the same as the triad of signals for a second surface ($L_1=L_2$; $M_1=M_2$; $S_1=S_2$). If the effect of an illuminant change were simply to modify the signals in each class of cone by a multiplicative constant, and to preserve cone-excitation ratios for the pair of surfaces, the match should be undisturbed (e.g. if $L'_1=kL_1$ and $L'_2=kL_2$ then $L'_1=L'_2$ etc., or if $L_1/L_2=1=L'_1/L'_2$ then $L'_1=L'_2$ etc.).

[6]Changes that preserve cone-excitation ratios appear like a 'wash of colour' over the scene, while natural changes with large violations of invariance appear spatially non-uniform. In Foster & Nascimento's experiment, some discriminations were easy and others not, but the difficult discriminations were always between changes that had only small violations of invariance. It would not be unreasonable for observers to assume that these presentations were the ones that most interested the experimenters, and thus to select these for response.

[7]In the case where the effect of an illuminant change is not a simple multiplication of cone signals, and surfaces do not plot on perfect straight lines, Khang and Zaidi (2002) show that this normalization will work provided that deviations from perfect correlation sum to zero (e.g. $\Sigma e_i=0$, where $e_i=S'_i-(kS_i)$, where S_i is the S-cone signal for the ith surface under the first illuminant; S'_i is the S-cone signal for the ith surface under the second illuminant, and k is a constant).

REFERENCES

Adelson, E. H. 1993 Perceptual organization and the judgment of brightness. *Science* **262**, 2042–2044.

Adelson, E. H. & Pentland, A. P. 1996 The perception of shading and reflectance. In *Perception as Bayesian inference* (ed. D. C. Knill & W. Richards), pp. 409–423. Cambridge University Press.

Amano, K. 2004 Colour constancy under simultaneous changes in surface position and illuminant. *Proc. R. Soc. B* **271**, 2319–2326. (doi:10.1098/rspb.2004.2884.)

Arend, L. 1994 Surface colors, illumination and surface geometry: intrinsic image models of human vision. In *Lightness, brightness and transparency* (ed. A. L. Gilchrist), pp. 159–214. Hillside, NJ: Lawrence Erlbaum Associates.

Arend, L. & Reeves, A. 1986 Simultaneous color constancy. *J. Opt. Soc. Am. A* **3**, 1743–1751.

Arend, L., Reeves, A., Schirillo, J. & Goldstein, R. 1991 Simultaneous color constancy: papers with diverse Munsell values. *J. Opt. Soc. Am. A* **8**, 661–672.

Barbur, J. L., de Cunha, D. & Williams, C. B. 2004 Study of instantaneous color constancy mechanisms in human vision. *J. Electronic Imaging* **13**, 15–28.

Beck, J. 1972 *Surface color perception*. New York: Cornell University Press.

Bhalla, M. & Proffitt, D. R. 1999 Visual-motor recalibration in geographical slant perception. *J. Exp. Psychol. (Hum. Percept. Perform.)* **25**, 1076–1096.

Blackwell, K. T. & Buchsbaum, G. 1988 Quantitative studies of color constancy. *J. Opt. Soc. Am. A* **5**, 1772–1780.

Bloj, M. G., Kersten, D. & Hurlbert, A. C. 1999 Perception of three-dimensional shape influences colour perception through mutual illumination. *Nature* **402**, 877–879.

Brainard, D. H. 1998 Color constancy in the nearly natural image. 2. Achromatic loci. *J. Opt. Soc. Am. A* **15**, 307–325.

Brainard, D. H. 2004 Color constancy. In *The visual neurosciences* (ed. L. M. Chalupa & J. S. Werner), pp. 948–961. Cambridge, MA/London: MIT Press.

Brainard, D. H. & Freeman, W. T. 1997 Bayesian color constancy. *J. Opt. Soc. Am. A* **14**, 1393–1411.

Brainard, D. H. & Wandell, B. A. 1986 Analysis of the retinex theory of color vision. *J. Opt. Soc. Am. A* **3**, 1651–1661.

Brainard, D. H. & Wandell, B. A. 1992 Asymmetric color matching: how color appearance depends on the illuminant. *J. Opt. Soc. Am. A* **9**, 1433–1448.

Brainard, D. H., Wandell, B. A. & Chichilnisky, E. J. 1993 Color constancy: from physics to appearance. *Curr. Dir. Psychol. Sci.* **2**, 165–170.

Brainard, D. H., Brunt, W. A. & Speigle, J. M. 1997 Color constancy in the nearly natural image. I. Asymmetric matches. *J. Opt. Soc. Am. A* **14**, 2091–2110.

Brainard, D. H., Kraft, J. M. & Longère, P. 2003 Colour constancy: developing empirical tests of computational models. In *Colour perception: mind and the physical world* (ed. R. Mausfeld & D. Heyer), pp. 307–328. Oxford University Press.

Bramwell, D. I. & Hurlbert, A. 1993 The role of object recognition in colour constancy. *Perception* **22**, S62.

Brill, M. H. 1978 A device performing illuminant-invariant assessment of chromatic relations. *J. Theor. Biol.* **71**, 473–478.

Brill, M. H. 1995 The relation between the color of the illuminant and the color of the illuminated object—commentary. *Color Res. Appl.* **20**, 70–72.

Brill, M. H. & West, G. 1986 Chromatic adaptation and color constancy: a possible dichotomy. *Color Res. Appl.* **11**, 196–204.

Broackes, J. 1992 The autonomy of color. In *Reduction, explanation and realism* (ed. K. Lennon & D. Charles), pp. 421–465. Oxford University Press.

Brown, R. O. 1994 The world is not gray. *Invest. Ophthalmol. Vis. Sci.* **35**, 2165.

Buchsbaum, G. 1980 A spatial processor model for object colour perception. *J. Franklin Inst.* **310**, 1–26.

Chichilnisky, E. J. & Wandell, B. A. 1999 Trichromatic opponent color classification. *Vis. Res.* **39**, 3444–3458.

Chittka, L., Shmida, A., Troje, N. & Menzel, R. 1994 Ultraviolet as a component of flower reflections, and the colour perception of Hymenoptera. *Vis. Res.* **34**, 1489–1508.

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

Cornelissen, F. W. & Brenner, E. 1995 Simultaneous colour constancy revisited: an analysis of viewing strategies. *Vis. Res.* **35**, 2431–2448.

Craven, B. J. & Foster, D. H. 1992 An operational approach to colour constancy. *Vis. Res.* **32**, 1359–1366.

Crawford, B. H. 1947 Visual adaptation in relation to brief conditioning stimuli. *Proc. R. Soc. B* **134**, 283–302.

Dannemiller, J. L. 1992 Spectral reflectance of natural objects: how many basis functions are necessary? *J. Opt. Soc. Am. A* **9**, 507–515.

Dannemiller, J. L. 1993 Rank orderings of photoreceptor photon catches from natural objects are nearly illuminant-invariant. *Vis. Res.* **33**, 131–140.

Derrington, A. M., Krauskopf, J. & Lennie, P. 1984 Chromatic mechanisms in lateral geniculate nucleus of macaque. *J. Physiol.* **357**, 241–265.

Drew, M. S. & Funt, B. V. 1990 Calculating surface reflectance using a single-bounce model of mutual reflection. In *Proc. IEEE Int. Conf. on Computer Vision, Osaka, Japan, 4–7 December 1990*, pp. 393–399.

D'Zmura, M. 1992 Color constancy: surface color from changing illumination. *J. Opt. Soc. Am. A* **9**, 490–493.

D'Zmura, M. & Iverson, G. 1993*a* Color constancy. I. Basic theory of two-stage linear recovery of spectral descriptions for lights and surfaces. *J. Opt. Soc. Am. A* **10**, 2148–2165.

D'Zmura, M. & Iverson, G. 1993*b* Color constancy. II. Results for two-stage linear recovery of spectral descriptions for lights and surfaces. *J. Opt. Soc. Am. A* **10**, 2166–2180.

D'Zmura, M. & Iverson, G. 1994 Color constancy. III. General linear recovery of spectral descriptions for lights and surfaces. *J. Opt. Soc. Am. A* **11**, 2398–2400.

D'Zmura, M., Iverson, G. & Singer, B. 1995 Probabilistic color constancy. In *Geometric Representations of Perceptual Phenomena: Articles in Honor of Tarow Indow's 70th Birthday* (ed. R. D. Luce, M. D'Zmura, D. Hoffman, G. Iverson & K. Romney), pp. 187–202. Hillsdale, NJ: Laurence Erlbaum Associates.

D'Zmura, M. & Lennie, P. 1986 Mechanisms of color constancy. *J. Opt. Soc. Am. A* **3**, 1662–1672.

Fairchild, M. D. 1993 Chromatic adaptation in hard-copy / soft-copy comparisons. In *Proc. SPIE: Color Hard Copy and Graphic Arts II (ed. J. Bares), vol. 1912, San Jose, CA, 31 January–5 February 1993*, pp. 47–61.

Fairchild, M. D. & Lennie, P. 1992 Chromatic adaptation to natural and incandescent illuminants. *Vis. Res.* **32**, 2077–2085.

Finlayson, G. D., Drew, M. S. & Funt, B. V. 1994 Color constancy: generalized diagonal transforms suffice. *J. Opt. Soc. Am. A* **11**, 3011–3019.

Finlayson, G. D., Hordley, S. D. & Hubel, P. M. 2001 Color by correlation: a simple, unifying framework for colour constancy. *IEEE Trans. Pattern Anal. Mach. Intell.* **23**, 1209–1221.

Forsyth, D. A. 1990 A novel algorithm for colour constancy. *Int. J. Comput. Vis.* **5**, 5–36.

Foster, D. H. 2003 Does colour constancy exist? *Trends Cogn. Sci.* **7**, 439–443.

Foster, D. H. & Nascimento, S. M. 1994 Relational colour constancy from invariant cone-excitation ratios. *Proc. R. Soc. B* **257**, 115–121.

Foster, D. H., Nascimento, S. M., Craven, B. J., Linnell, K. J., Cornelissen, F. W. & Brenner, E. 1997 Four issues concerning colour constancy and relational colour constancy. *Vis. Res.* **37**, 1341–1345.

Foster, D. H., Nascimento, S. M., Amano, K., Arend, L., Linnell, K. J., Nieves, J. L., Plet, S. & Foster, J. S. 2001 Parallel detection of violations of color constancy. *Proc. Natl Acad. Sci. USA* **98**, 8151–8156.

Funt, B. V., Drew, M. S. & Ho, J. 1991 Color constancy from mutual reflection. *J. Comput. Vis.* **6**, 5–24.

Gilchrist, A., Delman, S. & Jacobsen, A. 1983 The classification and integration of edges as critical to the perception of reflectance and illumination. *Percept. Psychophys.* **33**, 425–436.

Gilchrist, A., Kossyfidis, C., Bonato, F., Agostini, T., Cataliotti, J., Li, X., Spehar, B., Annan, V. & Economou, E. 1999 An anchoring theory of lightness perception. *Psychol. Rev.* **106**, 795–834.

Golz, J. & MacLeod, D. I. 2002 Influence of scene statistics on colour constancy. *Nature* **415**, 637–640.

Hagedorn, J. & D'Zmura, M. 2000 Color appearance of surfaces viewed through fog. *Perception* **29**, 1169–1184.

He, S. & MacLeod, D. I. 1998 Local nonlinearity in S-cones and their estimated light-collecting apertures. *Vis. Res.* **38**, 1001–1006.

Helmholtz, H. 1866 *Handbuch der Physiologishen Optik*, 3rd edn. [Transl. J. P. C. Southall 1962 *Helmholtz's treatise on physiological optics*, 2nd edn. New York: Dover.]

Helson, H. 1938 Fundamental problems in color vision. I. The principle governing changes in hue, saturation, and lightness of non-selective samples in chromatic illumination. *J. Exp. Psychol.* **23**, 439–476.

Helson, H. 1947 Adaptation-level as frame of reference for prediction of psychophysical data. *J. Opt. Soc. Am.* **60**, 1–29.

Helson, H. 1964 *Adaptation level theory*. Oxford: Harper and Row.

Hering, E. 1920 *Grundzüge der Lehre vom Lichtsinn*. [Transl. L. M. Hurvich & D. Jameson 1964 *Outlines of a theory of the light sense*. Cambridge, MA: Harvard University Press.]

Heywood, C. A., Cowey, A. & Newcombe, F. 1991 Chromatic discrimination in a cortically colour blind observer. *Eur. J. Neurosci.* **3**, 802–812.

Hiltunen, J. 1996 *Munsell colors mat (spectrophotometer measurements by Hiltunen)*. Retrieved September 10, 1999, from http://www.it.lut.fi/ip/research/color/database/download.html#munsell_spec_matt.

Hurlbert, A. C. 1986 Formal connections between lightness algorithms. *J. Opt. Soc. Am. A* **3**, 1684–1693.

Hurlbert, A. C. 1998 Computational models of colour constancy. In *Perceptual constancy: why things look as they do* (ed. V. Walsh & J. Kulikowski), pp. 283–322. Cambridge University Press.

Hurlbert, A. C. & Poggio, T. A. 1988 Synthesizing a color algorithm from examples. *Science* **239**, 482–485.

Hurlbert, A. C. & Wolf, K. 2004 Color contrast: a contributory mechanism to color constancy. *Progr. Brain Res.* **144**, 147–160.

Hurlbert, A. C., Bramwell, D. I., Heywood, C. & Cowey, A. 1998 Discrimination of cone contrast changes as evidence for colour constancy in cerebral achromatopsia. *Exp. Brain Res.* **123**, 136–144.

Ives, H. E. 1912 The relation between the color of the illuminant and the color of the illuminated object. *Trans. Illuminat. Eng. Soc.* **7**, 62–72. (Reprinted in: *Color Res. Appl.* **20**, 70–75.)

Jin, E. W. & Shevell, S. K. 1996 Color memory and color constancy. *J. Opt. Soc. Am. A* **13**, 1981–1991.

Joost, U., Lee, B. B. & Zaidi, Q. 2002 Lichtenberg's letter to Goethe on färbige schatten. *Color Res. Appl.* **27**, 300–303.

Judd, D. 1940 Hue saturation and lightness of surface colors with chromatic illumination. *J. Opt. Soc. Am.* **30**, 2–32.

Judd, D., MacAdam, D. L. & Wyszecki, G. W. 1964 Spectral distribution of typical daylight as a function of correlated color temperature. *J. Opt. Soc. Am.* **54**, 1031–1040.

Katz, D. 1935 *The world of colour*. London: K. Paul, Trench, Trubner.

Khang, B. G. & Zaidi, Q. 2002 Cues and strategies for color constancy: perceptual scission, image junctions and transformational color matching. *Vis. Res.* **42**, 211–226.

Kraft, J. M. & Brainard, D. H. 1999 Mechanisms of color constancy under nearly natural viewing. *Proc. Natl Acad. Sci. USA* **96**, 307–312.

Krantz, D. 1968 A theory of context effects based on cross-context matching. *J. Math. Psychol.* **5**, 1–48.

Land, E. H. 1964 The retinex. *Am. Scientist* **52**, 247–264.

Land, E. H. 1983 Recent advances in retinex theory and some implications for cortical computations: color vision and the natural image. *Proc. Natl Acad. Sci. USA* **80**, 5163–5169.

Land, E. H. 1986 Recent advances in retinex theory. *Vis. Res.* **26**, 7–21.

Land, E. H. & McCann, J. J. 1971 Lightness and retinex theory. *J. Opt. Soc. Am.* **61**, 1–11.

Land, E. H., Hubel, D. H., Livingstone, M. S., Perry, S. H. & Burns, M. M. 1983 Colour-generating interactions across the corpus callosum. *Nature* **303**, 616–618.

Lee, H. C. 1986 Method for computing the scene-illuminant chromaticity from specular highlights. *J. Opt. Soc. Am. A* **3**, 1694–1699.

Lehmann, T. M. & Palm, C. 2001 Color line search for illuminant estimation in real-world scenes. *J. Opt. Soc. Am. A* **18**, 2679–2691.

Linnell, K. J. & Foster, D. H. 1996 Dependence of relational colour constancy on the extraction of a transient signal. *Perception* **25**, 221–228.

Linnell, K. J. & Foster, D. H. 1997 Space-average scene colour used to extract illuminant information. In *John Dalton's colour vision legacy* (ed. C. Dickinson, I. Murray & D. Carden), pp. 501–509. London: Taylor and Francis.

Linnell, K. J. & Foster, D. H. 2002 Scene articulation: dependence of illuminant estimates on number of surfaces. *Perception* **31**, 151–159.

MacLeod, D. I. 2003 New dimensions in color perception. *Trends Cogn. Sci.* **7**, 97–99.

MacLeod, D. I. & Boynton, R. M. 1979 Chromaticity diagram showing cone excitation by stimuli of equal luminance. *J. Opt. Soc. Am.* **69**, 1183–1186.

MacLeod, D. I. & He, S. 1993 Visible flicker from invisible patterns. *Nature* **361**, 256–258.

MacLeod, D. I., Williams, D. R. & Makous, W. 1992 A visual nonlinearity fed by single cones. *Vis. Res.* **32**, 347–363.

Maloney, L. T. 1986 Evaluation of linear models of surface spectral reflectance with small numbers of parameters. *J. Opt. Soc. Am. A* **3**, 1673–1683.

Maloney, L. T. 1999 Physics-based approaches to modeling surface color perception. In *Color vision: from genes to perception* (ed. K. R. Gegenfurtner & L. T. Sharpe), pp. 387–422. Cambridge University Press.

Maloney, L. T. 2002 Illuminant estimation as cue combination. *J. Vis.* **2**, 493–504.

Maloney, L. T. 2003 Commentaries on Brainard, Kraft and Longère: surface colour perception and its environments. In *Colour perception: mind and the physical world* (ed. R. Mausfeld & D. Heyer), pp. 329–330. Oxford University Press.

Maloney, L. T. & Wandell, B. A. 1986 Color constancy: a method for recovering surface spectral reflectance. *J. Opt. Soc. Am. A* **3**, 29–33.

Maloney, L. T. & Yang, J. N. 2003 The illumination estimation hypothesis and surface color perception. In *Colour: connecting the mind to the physical world* (ed. R. Mausfeld & D. Heyer), pp. 329–330. Oxford University Press.

Marshall, N. J. 2000 Communication and camouflage with the same 'bright' colours in reef fishes. *Phil. Trans. R. Soc. B* **355**, 1243–1248.

Mausfeld, R. 1998 Colour perception: from Grassman codes to a dual code for object and illuminant colours. In *Color vision: perspective from different disciplines* (ed. W. Backhaus, R. Kliegl & J. S. Werner), pp. 219–250. Berlin: Walter De Gruyter.

McCann, J. J. 1992 Rules for colour constancy. *Ophthalmic Physiol. Opt.* **12**, 175–177.

McCann, J. J. 1997 Adaptation or contrast: the controlling mechanism for colour constancy. In *John Dalton's colour vision legacy* (ed. C. Dickinson, I. Murray & D. Carden), pp. 469–473. London: Taylor and Francis.

McCann, J. J., McKee, S. P. & Taylor, T. H. 1976 Quantitative studies in retinex theroy. A comparison between theoretical predictions and observer responses to the "color mondrian" experiments. *Vis. Res.* **16**, 445–458.

Mollon, J. D. 2003 Origins of modern color science. In *The science of color* (ed. S. K. Shevell), pp. 1–39, 2nd edn. Amsterdam/London: Elsevier/Optical Society of America.

Mollon, J. D., Regan, B. C. & Bowmaker, J. K. 1998 What is the function of the cone-rich rim of the retina? *Eye* **12**, 548–552.

Monge, G. 1789 Mémoire sur quelques phénomènes de la vision. *Annales de Chimie* **3**, 131–147.

Morgan, M. J., Watamaniuk, S. N. & McKee, S. P. 2000 The use of an implicit standard for measuring discrimination thresholds. *Vis. Res.* **40**, 2341–2349.

Nascimento, S. M. & Foster, D. H. 1997 Detecting natural changes of cone-excitation ratios in simple and complex coloured images. *Proc. R. Soc. B* **264**, 1395–1402.

Nascimento, S. M., Ferreira, F. P. & Foster, D. H. 2002 Statistics of spatial cone-excitation ratios in natural scenes. *J. Opt. Soc. Am. A* **19**, 1484–1490.

Pugh, E. N. & Mollon, J. D. 1979 A theory of the Pi 1 and Pi 3 colour mechanisms of Stiles. *Vis. Res.* **19**, 293–312.

Rushton, W. A. H. 1972 Review lecture: pigments and signals in colour vision. *J. Physiol.* **220**, 1P–31P.

Rutherford, M. D. & Brainard, D. H. 2002 Lightness constancy: a direct test of the illumination-estimation hypothesis. *Psychol. Sci.* **13**, 142–149.

Rüttiger, L., Braun, D. I., Gegenfurtner, K. R., Petersen, D., Schönle, P. & Sharpe, L. T. 1999 Selective color constancy deficits after circumscribed unilateral brain lesions. *J. Neurosci.* **19**, 3094–3106.

Sällström, P. 1973 Colour and physics: some remarks concerning the physical aspects of human colour vision. University of Stockholm, Institute of Physics Report, pp. 73–109.

Schein, S. J. & Desimone, R. 1990 Spectral properties of V4 neurons in the macaque. *J. Neurosci.* **10**, 3369–3389.

Schiller, P. H. & Lee, K. 1991 The role of the primate extrastriate area V4 in vision. *Science* **251**, 1251–1253.

Schirillo, J. A. & Shevell, S. K. 2000 Role of perceptual organization in chromatic induction. *J. Opt. Soc. Am. A* **17**, 244–254.

Schnapf, J. L., Nunn, B. J., Meister, M. & Baylor, D. A. 1990 Visual transduction in cones of the monkey *Macaca fascicularis*. *J. Physiol.* **427**, 681–713.

Shafer, S. A. 1985 Using color to separate reflection components. *Color Res. Appl.* **10**, 210–218.

Shevell, S. K. & Wei, J. 1998 Chromatic induction: border contrast or adaptation to surrounding light? *Vis. Res.* **38**, 1561–1566.

Smithson, H. & Zaidi, Q. 2004 Colour constancy in context: roles for local adaptation and levels of reference. *J. Vis.* **4**, 693–710.

Speigle, J. M. & Brainard, D. H. 1996 Is color constancy task independent? In *Proceedings IS&T/SID 4th Color Imaging Conference, Scottsdale, AZ, November 1995.* pp. 167-172

Stiles, W. S. 1939 The directional sensitivity of the retina and the spectral sensitivities of the rods and cones. *Proc. R. Soc. B* **127**, 64–105.

Stiles, W. S. 1949 Increment thresholds and the mechanisms of colour vision. *Doc. Opthalmol.* **3**, 138–163.

Stiles, W. S. 1961 Adaptation, chromatic adaptation, colour transformation. *Anales de la Real Sociedad Espanola de Fisica y Quimica: Seria A* **57**, 149–175.

Sugita, Y. 2004 Experience in early infancy is indispensable for color perception. *Curr. Biol.* **14**, 1267–1271.

Taylor, A. H. & Kerr, G. P. 1941 The distribution of energy in the visible spectrum of daylight. *J. Opt. Soc. Am.* **31**, 3.

Thouless, R. 1931 Phenomenal regression to the "real" object (I). *Br. J. Psychol.* **21**, 339–359.

Tominaga, S., Ebisui, S. & Wandell, B. A. 2001 Scene illuminant classification: brighter is better. *J. Opt. Soc. Am. A* **18**, 55–64.

Troost, J. M. & de Weert, C. M. 1991 Naming versus matching in color constancy. *Percept. Psychophys.* **50**, 591–602.

von Kries, J. 1878 *Beitrag zur Physiologie der Gesichtsempfindugen*. [Transl. Physiology of visual sensations. In *Sources of color science*, (ed. D. L. MacAdam), pp. 101–108, 1st edn. Cambridge, MA: MIT press.]

von Kries, J. 1905 Die Gesichtsempfindungen. In *Handbuch der Physiologie des Menschen*, vol. 3 (ed. W. Nagel), pp. 109–282. Braunschweig: Vieweg und Sohn.

Vrhel, M., Gershon, R. & Iwan, L. S. 1994 Measurement and analysis of object reflectance spectra. *Color Res. Appl.* **19**, 4–9.

Wachtler, T., Albright, T. D. & Sejnowski, T. J. 2001 Nonlocal interactions in color perception: nonlinear processing of chromatic signals from remote inducers. *Vis. Res.* **41**, 1535–1546.

Walsh, V. 1999 How does the cortex construct color? *Proc. Natl Acad. Sci. USA* **96**, 13 594–13 596.

Wandell, B. A. & Wade, A. R. 2003 Functional imaging of the visual pathways. *Neurol. Clin. North Am.* **21**, 417–443.

Wandell, B. A., Baseler, H. A., Poirson, A. B., Boynton, G. M. & Engel, S. A. 1999 Computational neuroimaging: color tuning in two human cortical areas measured using fMRI. In *Color vision: from genes to perception* (ed. K. R. Gegenfurtner & L. T. Sharpe), pp. 269–282. Cambridge University Press.

Webster, M. A. 1996 Human colour perception and its adaptation. *Netw. Comput. Neural Syst.* **7**, 587–634.

Webster, M. A. & Mollon, J. D. 1997 Adaptation and the color statistics of natural images. *Vis. Res.* **37**, 3283–3298.

Westland, S. & Ripamonti, C. 2000 Invariant cone-excitation ratios may predict transparency. *J. Opt. Soc. Am. A* **17**, 255–264.

Whittle, P. 2003 Contrast colours. In *Colour perception* (ed. R. Mausfeld & D. Heyer). New York: Oxford University Press.

Wolf, K. & Hurlbert, A. 2003 Effect of global contrast distribution on colour appearance. In *Normal and defective colour vision* (ed. J. D. Mollon, J. Pokorny & K. Knoblauch), pp. 147–160. Cambridge University Press.

Worthey, J. A. 1985 Limitations of color constancy. *J. Opt. Soc. Am. A* **2**, 1014–1026.

Worthey, J. A. & Brill, M. H. 1986 Heuristic analysis of von Kries color constancy. *J. Opt. Soc. Am. a* **3**, 1708–1712.

Wyszecki, G. W. & Stiles, W. 1982 *Color science*. New York: Wiley.

Yang, J. N. & Maloney, L. T. 2001 Illuminant cues in surface color perception: tests of three candidate cues. *Vis. Res.* **41**, 2581–2600.

Yoshioka, T., Dow, B. M. & Vautin, R. G. 1996 Neuronal mechanisms of color categorization in areas V1, V2 and V4 of macaque monkey visual cortex. *Behav. Brain Res.* **76**, 51–70.

Zaidi, Q. 1998 Identification of illuminant and object colors: heuristic-based algorithms. *J. Opt. Soc. Am. A* **15**, 1767–1776.

Zaidi, Q. 2001 Color constancy in a rough world. *Color Res. Appl.* **26**, 192–200.

Zaidi, Q., Spehar, B. & DeBonet, J. 1997 Color constancy in variegated scenes: role of low-level mechanisms in discounting illumination changes. *J. Opt. Soc. Am. A* **14**, 2608–2621.

Zeki, S. 1983*a* Colour coding in the cerebral cortex: the reaction of cells in monkey visual cortex to wavelengths and colours. *Neuroscience* **9**, 741–765.

Zeki, S. 1983*b* The relationship between wavelength and color studied in single cells of monkey striate cortex. *Progr. Brain Res.* **58**, 219–227.

Zeki, S., Aglioti, S., McKeefry, D. & Berlucchi, G. 1999 The neurological basis of conscious color perception in a blind patient. *Proc. Natl Acad. Sci. USA* **96**, 14 124–14 129.