Relational colour constancy from invariant cone-excitation ratios

DAVID H. FOSTER AND SÉRGIO M. C. NASCIMENTO

Department of Communication and Neuroscience, Keele University, Staffordshire ST5 5BG, U.K.

SUMMARY

Quantitative measurements of perceptual colour constancy show that human observers have a limited and variable ability to match coloured surfaces in scenes illuminated by different light sources. Observers can, however, make fast and reliable discriminations between changes in illuminant and changes in the reflecting properties of scenes, a discriminative ability that might be based on a visual coding of spatial colour relations. This coding could be provided by the ratios of cone-photoreceptor excitations produced by light from different surfaces: for a large class of pigmented surfaces and for surfaces with random spectral reflectances, these ratios are statistically almost invariant under changes in illumination by light from the sun and sky or from a planckian radiator. Cone-excitation ratios offer a possible, although not necessarily unique, basis for perceptual colour constancy in so far as it concerns colour relations.

1. INTRODUCTION

Colour constancy refers to the invariance of the perceived colour of a surface despite changes in the intensity and spectral composition of the light source. The phenomenon has been the subject of many studies. It was considered, notably, by Young (1807) and von Helmholtz (1866), the latter proposing that we form judgements about the colours of bodies by 'eliminating the differences of illumination by which a body is revealed to us' (p. 287); it was also considered by Hering (1878) and von Kries (1905), and, more recently, by Helson & Jeffers (1940), Judd (1940) and Land (1959*a*, *b*), the last formulating the 'Retinex' theory of colour vision, founded on the principle of colour constancy.

Yet, when quantitative measurements of human colour constancy have been made, it has been found that human observers have a limited and variable ability to make matches between coloured surfaces in scenes illuminated by different light sources (Arend & Reeves 1986; Tiplitz Blackwell & Buchsbaum 1988; Valberg & Lange-Malecki 1990; Troost & de Weert 1991; Reeves 1992). Performance depends on the prior knowledge and expectations of the observer, the nature of the instructions, and the spatial context in which the stimulus is viewed.

There is another approach to assessing human colour constancy, which considers the ability of observers to discriminate coloured scenes that have undergone changes in the spectral composition of the illuminant from coloured scenes that have undergone changes in their spectral reflecting properties. This operationally oriented approach is, in a sense, equivalent to the traditional one based on assessing invariant colour percepts (Appendix 1). In practice, observers can make discriminations between illuminant and non-illuminant changes, both highly reliably (Craven & Foster 1992) and extremely quickly (Foster *et al.* 1992); in a less formal context, the effect of manipulating apparent illuminant and surface reflecting properties has been compellingly demonstrated by R. W. G. Hunt to several audiences (see Brill & West 1986; Hunt 1991).

There is, of course, more to the phenomenon of colour constancy than making discriminations between illuminant and non-illuminant changes; for example, cognitive components are clearly important. Nevertheless, making these discriminations rather than extracting a constant colour percept may be the more appropriate activity for an organism that needs both to distinguish objects in its natural habitat and to detect, rather than to discount, changes in light from the sun and sky (Jameson & Hurvich 1989; Reeves 1992).

How, then, can human observers make discriminations between illuminant and non-illuminant changes? This ability might be based on a visual coding of spatial colour relations within a scene. One of the simplest visual quantities that might serve in this way is the ratio of cone-photoreceptor excitations produced by light from different surfaces (excitations within rather than between cone classes). These ratios have been assumed to be computed in adaptational (von Kries) models of colour constancy, as part of a scaling mechanism, usually with respect to an explicit or inferred reference 'white' (Ives 1912; West & Brill 1982; Worthey & Brill 1986); they have been incorporated into Retinex models as sequential products relating each surface to one or more high-reflectance surfaces (Land & McCann 1971; McCann *et al.* 1976); and they have been used in theoretical studies (see, for example, Brill & West 1981; West & Brill 1982; Worthey & Brill 1986) to set constraints on illuminant and surface reflectance spectra for colour

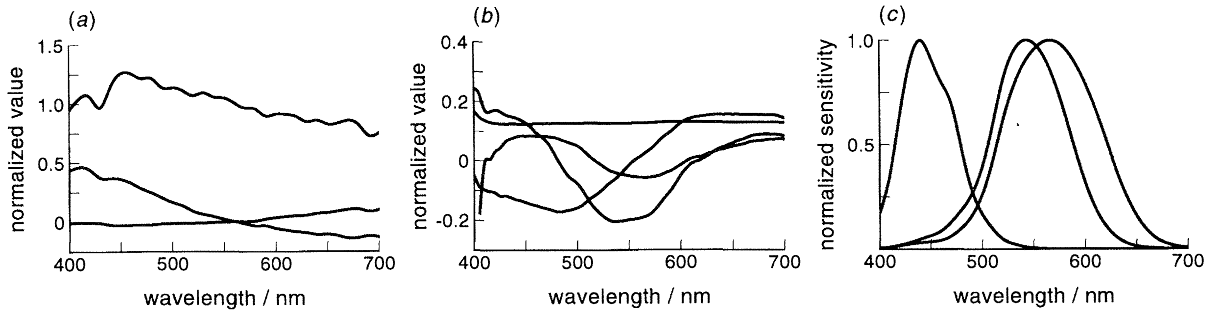

Figure 1. (*a*) Three characteristic vectors for skylight and sunlight-with-skylight from Judd *et al.* (1964); data normalized to unity at 560 nm. (*b*) Four of the eight characteristic vectors for the Munsell set from Parkkinen *et al.* (1989); integrated data normalized to unity. (*c*) Spectral sensitivities of short-, medium-, and long-wavelength sensitive cones from Smith & Pokorny (1972, 1975) and Wyszecki & Stiles (1982); data normalized to unity at maxima.

constancy, although results from the last have yielded unrealistic data (Brill & West 1986, p. 198). In this study, rather than the assumed constancy of cone-excitation ratios being used to derive theoretical spectra, a converse approach was taken: observed spectra were used to derive estimates of the constancy of the ratios; as only colour relations were of interest, the problem of computing a reference white was obviated. It is shown by numerical simulation that a coding of colour relations in terms of ratios could provide an observer with all the information necessary to make the required discriminations.

2. METHODS

Daylight spectral power distributions were drawn from published data in which 622 samples of skylight and sunlight-with-skylight were subjected to a principal components analysis (Judd *et al.* 1964). A spectral power distribution giving spectral irradiance $E(\lambda)$ at wavelength λ could be expressed in terms of three characteristic vectors (basis functions) with values $E_j(\lambda)$, $j = 1, 2, 3$, at each λ; thus, $E(\lambda) = \Sigma_{j=1}^{3} a_j E_j(\lambda)$, where the a_j are numerical weighting coefficients. The characteristic vectors are plotted in figure 1*a*. With these three functions, the average standard deviation of the reconstituted daylight spectra is about 0.03–0.04 on a scale normalized to unity at 560 nm (Judd *et al.* 1964). The a_j were selected here so that spectra were generated with correlated colour temperatures over the range 4300–25000 K (Judd *et al.* 1964).

The spectral power distribution of a planckian (blackbody) radiator at temperature T produces a spectral irradiance $E(\lambda)$ at wavelength λ according to Planck's formula; thus, $E(\lambda) = c_1 \lambda^{-5}\{\exp[c_2/(\lambda T)]-1\}^{-1}$, where c_1 and c_2 are constants (see, for example, Wyszecki & Stiles 1982). Spectra were generated for temperatures over the range 2000–100000 K. The CIE (Commission Internationale de l'Eclairage) Standard Illuminant A, a tungsten-filament source, has the relative spectral power distribution of a planckian radiator at a temperature of 2854 K. Daylight and planckian spectra are distinct; but, in the CIE 1931 (x,y)-chromaticity diagram, the daylight coordinates fall on a smooth curve almost parallel to and slightly on the green side of the planckian locus (Judd *et al.* 1964).

Random spectral power distributions were generated as follows. On each of n consecutive wavelength intervals $\lambda_l \leqslant \lambda < \lambda_{l+1}$, where $0 \leqslant l < n$, the spectral irradiance $E(\lambda)$ was constant with λ. The $E(\lambda_l)$ were independent, identically distributed, random variables, each with a uniform distribution on the unit interval [0, 1]. The number n of wavelength intervals was set so that the lengths $\Delta\lambda = \lambda_{l+1} - \lambda_l$ were fixed at 10, 20, 30, 50, 100 and 150 nm.

Spectral reflectances of pigmented surfaces were taken from a principal components analysis (Parkkinen *et al.* 1989) of 1257 samples from the *Munsell Book of Color* (Munsell Color 1976). A spectral reflectance with value $R(\lambda)$ at wavelength λ could be expressed in terms of eight characteristic vectors with values $R_k(\lambda)$, $k = 1, 2, \ldots, 8$, at each λ; thus, $R(\lambda) = \Sigma_{k=1}^{8} b_k R_k(\lambda)$, where the b_k are numerical weighting coefficients. Four of these characteristic vectors are plotted in figure 1*b*. With eight characteristic vectors, 98.4% of the Munsell set can be reconstructed to within a reflectance error of 0.02, on unit scale, averaged over the wavelength band (Parkkinen *et al.* 1989). This set of characteristic vectors reproduced accurately the spectral reflectances of a large set of natural coloured objects, such as flowers, flower clusters, leaves, and berries (Jaaskelainen *et al.* 1990), and it is based on a larger sample size than in previously published sets. The b_k were selected so that spectra were generated from the set of 1257 Munsell samples. Random reflectances were generated in the same way as for random illuminants.

Data for the spectral sensitivities of cones with light incident at the cornea were based on transformations of Judd's modification of the colour-matching functions for the CIE 1931 Standard Colorimetric Observer (Smith & Pokorny 1972, 1975; see Wyszecki & Stiles 1982). Figure 1*c* shows the three spectral sensitivities. Cone excitations were calculated routinely as follows. Let $S_i(\lambda)$, $i = 1, 2, 3$, be the spectral sensitivities at wavelength λ of short-, medium-, and long-wavelength sensitive cones, respectively; and let q_i be the excitation in cone class i produced by light from a surface with spectral reflectance $R(\lambda)$ at λ illuminated by a source e giving spectral irradiance $E(\lambda)$ at λ. Then

$$q_i = c \int S_i(\lambda)\, R(\lambda)\, E(\lambda)\, \mathrm{d}\lambda, \tag{1}$$

where the constant c depends on the conditions of viewing; its value is unimportant since only correlations and ratios of the q_i were used here. Integration was performed numerically over the range 400–700 nm according to Simpson's rule with step size 10 nm, corresponding to the sampling interval for the available daylight spectra.

3. CONE EXCITATIONS FOR A SINGLE SURFACE

It is first shown that cone excitations are themselves not invariant under changes in natural illuminants, although the extent of the departures from invariance

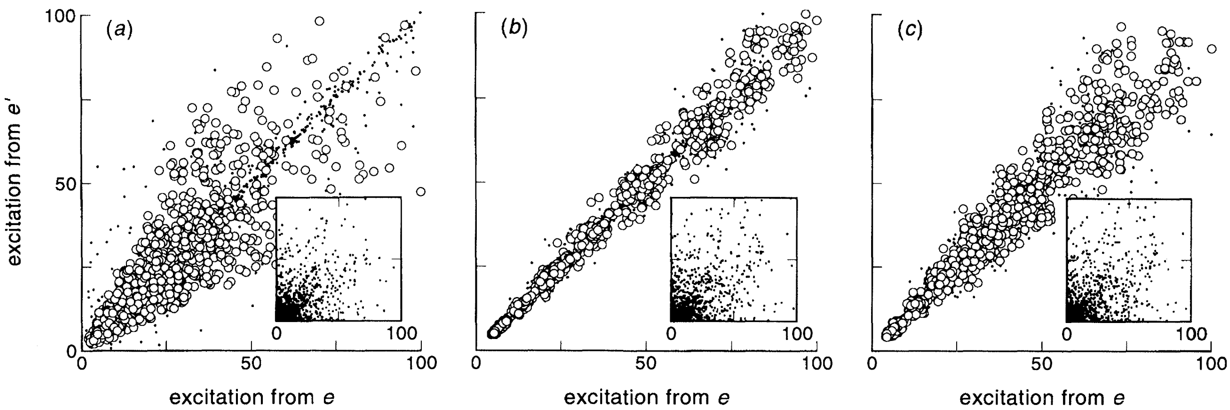

Figure 2. Scatterplot of cone excitations for each of the three cone classes: (*a*) short-, (*b*) medium-, (*c*) long-wavelength sensitive. Each point represents a pair of excitations q_i, q_i' in cone class i produced by light from a single surface drawn at random from the Munsell set, illuminated in turn by two illuminants e, e' drawn at random from the set of daylight spectra at constant power, correlated colour temperatures 4300–25 000 K (open circles) or from a planckian radiator at constant power, temperatures 2000–100 000 K (small solid circles). (The insets show corresponding data from a planckian radiator without the constant-power constraint.) Cone excitations have been scaled by an arbitrary factor; true values depend on the conditions of viewing. Based on 1000 random samples of surfaces and pairs of illuminants in each condition.

Table 1. *Proportion of pairs of cone excitations within 10% of each other and proportion of pairs of ratios of cone excitations within 10% of each other obtained with various combinations of illuminant and reflecting surfaces for each of the three cone classes*

(Daylight illuminants were generated with correlated colour temperatures 4300–25 000 K; planckian radiators with temperatures 2000–100 000 K; Munsell surfaces from a set of 1257 Munsell samples; and random spectra with wavelength intervals $\Delta\lambda$ of 10 nm and 50 nm. Based on 1000 random samples of surfaces or pairs of surfaces and pairs of illuminants in each condition.)

	illuminant	surface	short-wave	medium-wave	long-wave
excitations	(*a*) daylight[a]	Munsell	0.212	0.829	0.552
	(*b*) planckian[a]	Munsell	0.728	0.838	0.746
	(*c*) planckian	Munsell	0.059	0.060	0.060
excitation ratios	(*d*) daylight	Munsell	0.992	0.974	0.961
	(*e*) planckian	Munsell	0.982	0.971	0.964
	(*f*) planckian	random (10 nm)	0.976	0.987	0.987
	(*g*) planckian	random (50 nm)	0.935	0.950	0.953
	(*h*) random (10 nm)	Munsell	0.926	0.855	0.849
	(*i*) random (50 nm)	Munsell	0.817	0.625	0.652
	(*j*) random (50 nm)	random (50 nm)	0.330	0.310	0.326

[a] At constant power.

depends on the class of cones, on whether the light is from the sun and sky or from a planckian radiator, and on how the power of the illuminant is constrained.

Consider the excitations q_i, q_i' in cone class i produced by light from a single surface drawn at random from the Munsell set, illuminated in turn by two illuminants e, e' drawn at random from the set of daylight spectra at constant power. A plot of the points with coordinates (q_i, q_i') over 1000 iterations of this selection procedure is shown in figure 2 (open circles) for each cone class i (illuminants e, e' and surfaces being chosen afresh on each iteration). If the excitation produced by light from each surface was invariant under changes in illuminant, that is, if $q_i = q_i'$, then all the points would lie on the line of unit slope and zero intercept, and R^2, the proportion of variance in the data accounted for by this line, would be 1.0 (R^2 should not be confused with the symbol for spectral reflectance used elsewhere). For short-wavelength sensitive cones (figure 2*a*), there was considerable scatter in the data ($R^2 = 0.68$), and invariance manifestly fails; but, for both medium- and long-wavelength sensitive cones (figure 2*b*, *c*), there was relatively little scatter ($R^2 = 0.98$ and 0.93, respectively). The periodic clustering in the data for medium-wavelength sensitive cones (figure 2*b*) was traceable to the discrete levels taken by the lightness variable, Munsell Value.

The physical relevance of the residual scatter in the data is not immediate from the R^2 values alone, and a more direct quantifier of the invariance of cone excitations is given by the proportion of pairs of excitations q_i, q_i' falling within a certain range of each other, say 10%. The proportion of pairs of excitations within this range ($q_i'/q_i \leqslant 1.1$ if $q_i \leqslant q_i'$, or $q_i/q_i' \leqslant 1.1$ if $q_i' \leqslant q_i$) is listed in table 1, row (*a*), for each cone class i; values ranged from about 0.21 to 0.83, the latter value for medium-wavelength sensitive cones.

The relatively high proportion of pairs of excitations within 10% of each other for medium-wavelength sensitive cones derives from the spectral composition of

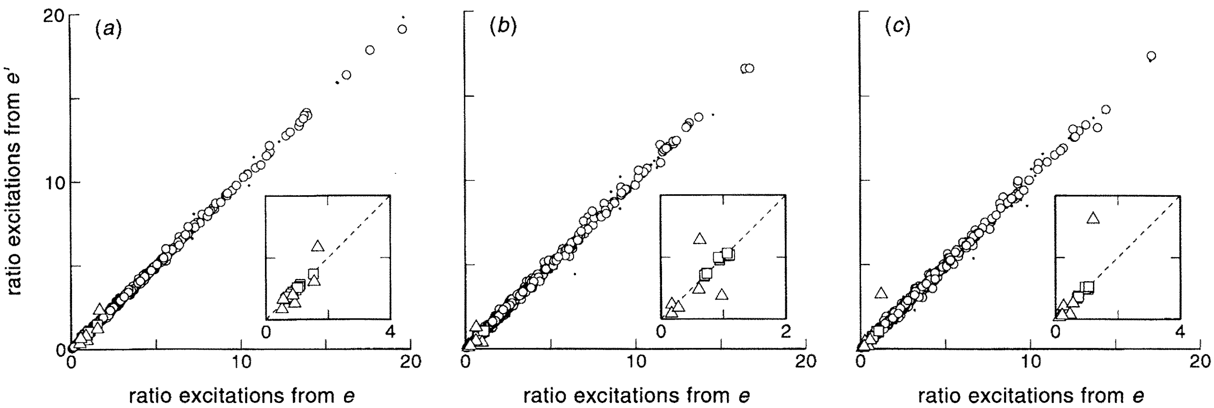

Figure 3. Scatterplot of ratios of cone excitations for each of the three cone classes: (*a*) short-, (*b*) medium-, (*c*) long-wavelength sensitive. Each point represents a pair of ratios r_i, r_i' of excitations in cone class i produced by light from two surfaces drawn at random from the Munsell set, illuminated in turn by two illuminants e, e' drawn at random from the set of daylight spectra, correlated colour temperatures 4300–25000 K (open circles) or from a planckian radiator, temperatures 2000–100000 K (small solid circles, plotted under open circles). Data for particular pairs of surfaces from the Munsell set are shown (on an expanded scale in the insets) for surfaces of intermediate Chroma (5R 5/4, 5G 5/4, 5Y 5/4, 5B 5/4), separated by increments of 2 units in Chroma, 1 in Value, and 2.5 in Hue (open squares); and for surfaces of extreme Chroma (5R 6/14, 5G 7/10, 5Y 7/12, 5B 7/8), in all combinations (open triangles). Based on 1000 random samples of pairs of surfaces and pairs of illuminants in each condition.

daylight illuminants and the constraint of constant power. When daylight spectra were normalized at a fixed wavelength (560 nm, as in Judd *et al.* (1964) and in figure 1*a*), the proportion for medium-wavelength sensitive cones fell to 0.74, and for short-wavelength sensitive cones to 0.16, but, for long-wavelength sensitive cones, it increased to 0.98.

When, instead, illuminants were drawn from a planckian radiator at constant power, scatter in the data was almost uniform across the three cone classes (small solid circles in figure 2*a*–*c*); the proportion of pairs of excitations within 10% of each other ranged from about 0.73 to 0.84 (table 1, row (*b*)). Without the constant-power constraint (the planckian radiator then varying only in its temperature), scatter increased across all cone classes (insets to figure 2*a*–*c*); and the proportion of pairs of excitations within 10% of each other fell to about 0.06 for each cone class (table 1, row (*c*)).

Data in a format similar to those in figure 2 have been reported previously (Dannemiller 1993) for two fixed illuminants, for which excitations in each cone class then fall very close to a linear regression line.

4. RATIOS OF CONE EXCITATIONS FOR PAIRS OF SURFACES

Consider now the ratio r_i of excitations in cone class i produced by light from two surfaces drawn at random from the Munsell set, illuminated by a single illuminant e drawn at random from the daylight set; and suppose that r_i' is the corresponding ratio for the same two surfaces illuminated by another illuminant e', also drawn at random from the daylight set. Notice that it is immaterial here whether power is constant, since any scaling disappears in the computation of the ratios. A plot of the points (r_i, r_i') over 1000 iterations of this procedure is shown in figure 3 (open circles), for each cone class i (illuminants e, e' and surfaces again being chosen afresh on each iteration). The proportion of pairs of ratios r_i, r_i' within 10% of each other is listed in table 1, row (*d*), for each cone class i; none fell below 0.96. When illuminants were drawn from a planckian radiator, the scatter was little changed (small solid circles in figure 3, plotted under open circles); the corresponding proportions of pairs of ratios within 10% of each other, listed in table 1, row (*e*), did not fall below 0.96.

Some justification for the use of this limiting value of 10% for proportions was inferred from simulations of a task requiring discriminations between illuminant and non-illuminant changes (Craven & Foster 1992) in which values of 5%, 10% and 20% were each used to predict false-alarm rates: the 10% value gave the most accurate predictions of the experimental performances. An indication of the colorimetric significance of the departures from the line $r_i' = r_i$ was obtained by computing ratios for representative pairs of surfaces from the Munsell set, illuminated by planckian radiators with extreme temperatures 2000 K and 100000 K. These points are shown in figure 3 for surfaces of intermediate Chroma (in the Munsell system of notation, 5R 5/4, 5G 5/4, 5Y 5/4, 5B 5/4), separated by increments of perceptually similar magnitude of 2 units in Chroma, 1 in Value, and 2.5 in Hue (open squares); and for surfaces of extreme Chroma (5R 6/14, 5G 7/10, 5Y 7/12, 5B 7/8), in all six combinations (open triangles). The expanded scale of the insets to the figure shows more clearly that for only a few of the extreme-Chroma surfaces are there large departures from invariance.

The ratio of cone excitations produced by light from pairs of surfaces is therefore almost invariant under changes in natural illuminant. Does this invariance depend on using spectral reflectances drawn from the Munsell set? In fact, a small, net increase in the proportion of pairs of ratios falling within 10% of each other was found when the illuminant was planckian and spectral reflectances were random, with wavelength interval $\Delta\lambda$ equal to 10 nm (table 1, row (*f*)).

It might be hypothesized that this invariance with random spectral reflectances was the result of the small value of $\Delta\lambda$: given the broad spectral sensitivities of cones, random spectra would appear as approximately constant functions. Yet the proportion of pairs of ratios within 10% of each other decreased only a little when $\Delta\lambda$ was increased to 50 nm (table 1, row (*g*)), and was still better than 0.93 for all cone classes. For all other values of $\Delta\lambda$ tested (see Methods), the proportion was never less than 0.93.

A less complete invariance was obtained with the opposite arrangement: when the reflecting surfaces were drawn from the Munsell set and illuminant spectral power distributions were random with $\Delta\lambda$ equal to 10 nm, the proportion of pairs of ratios within 10% of each other ranged from about 0.86 to 0.93 (table 1, row (*h*)); when $\Delta\lambda$ was increased to 50 nm, the proportion fell rather further, ranging from about 0.63 to 0.82 (table 1, row (*i*)). For all other values of $\Delta\lambda$ tested, the proportion was never less than 0.61. The asymmetry between the effects of random reflectance spectra and random illuminant spectra is not unreasonable: the ratios are asymmetric with respect to surface and illuminant, and Munsell spectra are different from planckian spectra (figure 1).

Finally, when both reflectance and illuminant spectra were random, with $\Delta\lambda$ equal to 50 nm, the proportion of pairs of ratios within 10% of each other fell within the range 0.31–0.33 (table 1, row (*j*)).

5. COMMENT

In everyday life, rapid changes in the spectral content of the light illuminating a scene occur naturally: the sun may pass behind a cloud, or an incandescent lamp may be switched on in a room already illuminated by daylight; although the colour of the scene might be perceived to change, the relations between the apparent colours of objects are, in general, preserved.

One indirect way that this invariance might be achieved is by first recovering (possibly partially) the spectral reflectances of the surfaces in the scene or their colours under a fixed illuminant, a task that is computationally feasible for variously constrained sets of illuminants, reflecting surfaces, and views (see, for example, West 1979; Buchsbaum 1980; Worthey 1985; Brill & West 1986; D'Zmura & Lennie 1986; Hurlbert 1986; Maloney & Wandell 1986; Hurlbert & Poggio 1988; Forsyth 1990; D'Zmura & Iverson 1993*a*, *b*; Maloney 1993); spatial colour relations could then be expressed in terms of these recovered quantities. It is not clear, however, whether models that accurately recover reflectance data are relevant to human psychophysical performance in traditional colour-constancy tasks (see Introduction, and, for single-cell and animal behavioural data, see, for example, Zeki 1983*a*, *b*; Carden *et al.* 1992). Another way this invariance might be achieved, without implying colour constancy in the traditional sense, is by directly coding spatial colour relations so that they are illuminant invariant. As has been shown here, cone-excitation ratios could provide the appropriate signals: for a large class of pigmented surfaces, they are almost invariant under changes in daylight or planckian illuminants; moreover, this invariance is maintained when the spectral reflectances of the surfaces are random, but only partly maintained when the spectral power distributions of the illuminant are random. In the language of geometric-invariance theory (see, for example, Mundy & Zisserman 1992), cone excitations are relative invariants with respect to natural illuminant changes.

There are two caveats to this analysis. First, cone-excitation ratios need not be unique in providing the visual system with invariant signals: other post-receptoral signals could have a related role (see, for example, Dannemiller 1989; van Hateren 1993). Second, the relational colour constancy provided by cone-excitation ratios is a statistical phenomenon: it holds for almost all surfaces, but there are particular instances for which the ratio of excitations produced by light from two surfaces fails markedly to be illuminant invariant; in addition to the extreme-Chroma Munsell samples of different hue considered earlier, there is the set of metameric materials, that is, those materials that by definition appear identical under one illuminant and different under another. Such materials are actually rare, both in the natural world and in the samples considered here (Worthey 1985).

Why should cone-excitation ratios provide such invariant signals? Previous theoretical analyses have considered the role of receptor bandwidth in relation to reflectance and illuminant spectra (Buchsbaum & Gottschalk 1984; Maloney 1986; Worthey & Brill 1986); the argument, in the limit, is as follows. Consider the profiles of the cone spectral sensitivities (figure 1*c*) and imagine their being narrowed until they are of infinitesimal width. Suppose that each cone class i has zero sensitivity at all wavelengths except μ_i and an integrated sensitivity s_i. Then, from equation (1), the ratio r_i of excitations in each such cone class produced by light from two surfaces with arbitrary (non-zero) spectral reflectances $R_1(\lambda)$, $R_2(\lambda)$ at wavelength λ illuminated by a source giving spectral irradiance $E(\lambda)$ at λ is simply $[s_i R_1(\mu_i) E(\mu_i)]/[s_i R_2(\mu_i) E(\mu_i)]$; that is, $r_i = R_1(\mu_i)/R_2(\mu_i)$, which is independent of the illuminant. Thus, in a visual system with these hypothetical spectral sensitivities, the preservation of spatial colour relations between surfaces would, at each of these wavelengths μ_i, be equivalent to the preservation of spatial luminance contrast under changes in luminance level in an achromatic or monochromatic scene. As the width of these hypothetical spectral sensitivities increases, the ratios of cone excitations produced by light from different surfaces become more and more illuminant dependent (Worthey & Brill 1986), an effect that may be considered from the viewpoint of sampling theory (see Barlow 1982; Buchsbaum & Gottschalk 1984; Maloney 1986). Informal theoretical estimates of the role of receptor spectral width on colour constancy may, however, be uncertain, since the present analysis has shown that cone-excitation ratios are almost illuminant invariant for random spectral reflectances with wavelength intervals of 10–50 nm and larger. The level at which departures from invariance become perceptible with given illuminant and reflectance spectra remains an empirical problem.

We thank B. B. Boycott, B. J. Craven, P. Fletcher, T. M. Goodman, R. W. G. Hunt, K. J. Linnell, T. Moons, J. D. Moreland, E. J. Pauwels and S. Westland for helpful advice; and K. J. Linnell, A. Reeves, M. G. A. Thomson and S. Westland for critical reading of the manuscript. We also thank J. P. S. Parkkinen, J. Hallikainen and T. Jaaskelainen for supplying their data for the Munsell set. This work was supported by the Wellcome Trust, the Junta Nacional de Investigação Científica e Tecnológica (grant no. BD/1328/91-RM), and the Universidade do Minho, Braga, Portugal.

APPENDIX 1.

The following group-theoretic argument shows that the traditional interpretation of colour constancy is formally equivalent to one based on discriminating illuminant from non-illuminant changes (the operational approach). The argument is technically non-constructive, and therefore does not deal with particular algorithms for generating colour-constant functions.

First, some basic definitions and properties. Let Σ be a family of spatially disjoint, uniformly coloured surfaces (an example of such a family being a Mondriaan painting), illuminated by a spatially uniform illuminant giving spectral irradiance $E(\lambda) > 0$ at each wavelength λ over the range Λ. The family Σ is assumed to comprise more than one surface σ, and each such surface is assumed to be non-fluorescent and to have spectral reflectance $R(\lambda, \sigma) \geqslant 0$ at each λ in Λ (this notation for spectral reflectance is more useful here than $R_\sigma(\lambda)$). There are no other constraints on the functions E, R. Given the pair E, R, the colour signal C (Buchsbaum & Gottschalk 1984) presented to the eye is uniquely determined with respect to wavelength and surface: $C(\lambda, \sigma) = E(\lambda)\, R(\lambda, \sigma)$ for each λ in Λ, σ in Σ. In contrast, given an arbitrary colour signal C presented to the eye, there are an infinite number of pairs E_1, R_1 such that $C(\lambda, \sigma) = E_1(\lambda)\, R_1(\lambda, \sigma)$ for each λ in Λ, σ in Σ (e.g. the same Mondriaan painting with yellower pigments viewed under a bluer illuminant).

Let $\mathscr{C}$ be the set of all colour signals. Changes in the spectral power distribution of an illuminant may be treated as illuminant transformations of $\mathscr{C}$ into itself. Explicitly, for C_1, C_2 in $\mathscr{C}$, suppose that E_1, E_2, R are such that $C_1(\lambda, \sigma) = E_1(\lambda)\, R(\lambda, \sigma)$ and $C_2(\lambda, \sigma) = E_2(\lambda)\, R(\lambda, \sigma)$ for each λ in Λ, σ in Σ; as E_1 is nowhere zero, it follows that $C_2(\lambda, \sigma) = (E_2(\lambda)/E_1(\lambda))\, C_1(\lambda, \sigma)$ for each λ in Λ, σ in Σ. This relation between C_1, C_2 represents a surface-independent, spectral scaling; it is independent of the choice of E_1, E_2, R giving rise to C_1, C_2, and therefore defines a transformation T of C_1 into C_2; namely, $C_2(\lambda, \sigma) = K_T(\lambda)\, C_1(\lambda, \sigma)$ for an everywhere-positive scalar function K_T. This transformation is an illuminant transformation. It is equivalent to the insertion of a coloured filter between illuminant and surface, and filters and transformations are identified in the following. Because the set $\mathscr{T}$ of all illuminant transformations T is a one-to-one copy of the multiplicative group of everywhere-positive functions K_T on Λ, it inherits the group structure of the latter. The group operation is the composition of transformations: for T_1, T_2 in $\mathscr{T}$, the composite $T_1 \circ T_2$ is given by the usual rule $(T_1 \circ T_2)\ (C) = T_1(T_2(C))$ for all C in $\mathscr{C}$.

Consider the operational approach to colour constancy. It assumes that, for each pair of colour signals C_1, C_2 in $\mathscr{C}$, it is possible to determine whether they are related by an illuminant transformation; that is, whether $T(C_1) = C_2$ for some T in $\mathscr{T}$. Because $\mathscr{T}$ is a group, this relation is reflexive, symmetric, and transitive, and is therefore an equivalence relation on $\mathscr{C}$, partitioning it into disjoint subsets or equivalence classes of illuminant-related colour signals. For each C in $\mathscr{C}$, let $[C]$ be its equivalence class.

If $\mathscr{T}$ is not a group, then the foregoing relation ceases to be an equivalence relation, and $\mathscr{C}$ cannot be partitioned by equivalence classes. Physically, this problem could occur if the illuminant transformations were effected with just two filters T_1, T_2 such that T_1 is not the inverse of T_2. Suppose that T_1, T_2 take, respectively, colour signals C_1, C_2 into colour signal C_0. Then C_1, C_2 belong to the same class of illuminant-related colour signals, namely, that containing C_0. But C_1 is not illuminant related to C_2, since there is no filter that takes C_1 into C_2 (if $\mathscr{T}$ were a group, it would be the composite $T_2^{-1} \circ T_1$). Therefore C_1, C_2 belong to different classes of illuminant-related colour signals, which is a contradiction.

Now consider the traditional approach to colour constancy. It assumes that it is possible to find some 'colour-constant' function ϕ that associates with each colour signal C in $\mathscr{C}$ a percept $\phi(C)$ that is invariant under illuminant transformations; more precisely, given C_1, C_2 in $\mathscr{C}$, then $\phi(C_1) = \phi(C_2)$ if, and only if, $T(C_1) = C_2$ for some T in $\mathscr{T}$. The type of the values of ϕ need not be made explicit.

The function ϕ thus associates a unique value $\phi(C)$ with each equivalence class $[C]$. The association is well defined precisely because $\mathscr{T}$ is a group: if C_1, C_2 are in $[C]$, then $\phi(C_1) = \phi(C_2)$, by the same reasoning as was used to show the consistent partitioning of $\mathscr{C}$ by $\mathscr{T}$. It is one-to-one because if two values $\phi(C_1)$, $\phi(C_2)$ coincide, then $T(C_1) = C_2$ for some T in $\mathscr{T}$, and therefore $[C_1] = [C_2]$. This one-to-one correspondence between colour-constant percepts and classes of illuminant-related colour signals establishes the formal equivalence of the two approaches.

REFERENCES

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

Barlow, H. B. 1982 What causes trichromacy? A theoretical analysis using comb-filtered spectra. *Vision Res.* **22**, 635–643.

Brill, M. & West, G. 1981 Contributions to the theory of invariance of color under the condition of varying illumination. *J. math. Biol.* **11**, 337–350.

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

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

Buchsbaum, G. & Gottschalk, A. 1984 Chromaticity coordinates of frequency-limited functions. *J. opt. Soc. Am. A* **1**, 885–887.

Carden, D., Hilken, H., Butler, S. R. & Kulikowski, J. J. 1992 Lesions of primate visual area V4 produce long-lasting deficits to colour constancy. *Irish J. Psychol.* **13**, 455–472.

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

Dannemiller, J. L. 1989 Computational approaches to color constancy: adaptive and ontogenetic considerations. *Psychol. Rev.* **96**, 255–266.

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

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. & Lennie, P. 1986 Mechanisms of color constancy. *J. opt. Soc. Am. A* **3**, 1662–1672.

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

Foster, D. H., Craven, B. J. & Sale, E. R. H. 1992 Immediate colour constancy. *Ophthal. physiol. Optics* **12**, 157–160.

van Hateren, J. H. 1993 Spatial, temporal and spectral preprocessing for colour vision. *Proc. R. Soc. Lond.* B **251**, 61–68.

von Helmholtz, H. 1866 *Handbuch der Physiologischen Optik*, vol. 2, 1st edn. (Translation 3rd edn: *Helmholtz's treatise on physiological optics* (ed. J. P. C. Southall), pp. 286–287. Optical Society of America (1924). Republished New York: Dover Publications (1962).)

Helson, H. & Jeffers, V. B. 1940 Fundamental problems in color vision. II. Hue, lightness, and saturation of selective samples in chromatic illumination. *J exp. Psychol.* **26**, 1–27.

Hering, E. 1878 Collected works. In *Outlines of a theory of the light sense* Cambridge, Massachusetts: Harvard University Press (1920). (Translated by L. M. Hurvich & D. Jameson (1964), pp. 7–8.)

Hunt, R. W. G. 1991 Revised colour-appearance model for related and unrelated colours. *Color Res. Appl.* **16**, 146–165.

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

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

Ives, H. E. 1912 The relation between the color of the illuminant and the color of the illuminated object. *Trans. illum. Eng. Soc.* **7**, 62–72.

Jaaskelainen, T., Parkkinen, J. & Toyooka, S. 1990 Vector-subspace model for color representation. *J. opt. Soc. Am. A* **7**, 725–730.

Jameson, D. & Hurvich, L. M. 1989 Essay concerning color constancy. *A. Rev. Psychol.* **40**, 1–22.

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

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

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

Land, E. H. 1959*a* Color vision and the natural image. Part I. *Proc. natn. Acad. Sci. U.S.A.* **45**, 115–129.

Land, E. H. 1959*b* Color vision and the natural image. Part II. *Proc. natn. Acad. Sci. U.S.A.* **45**, 636–644.

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

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. 1993 Color constancy and color perception: the linear-models framework. *Attent. Perform.* **14**, 59–78.

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

McCann, J. J., McKee, S. P. & Taylor, T. H. 1976 Quantitative studies in Retinex theory. A comparison between theoretical predictions and observer responses to the "Color Mondrian" experiments. *Vision Res.* **16**, 445–458.

Mundy, J. L. & Zisserman, A. (ed.) 1992 *Geometric invariance in computer vision*. Cambridge, Massachusetts: MIT Press.

Munsell Color 1976 *Munsell book of color–matte finish collection*. Baltimore: Munsell Color Corporation.

Parkkinen, J. P. S., Hallikainen, J. & Jaaskelainen, T. 1989 Characteristic spectra of Munsell colors. *J. opt. Soc. Am. A* **6**, 318–322.

Reeves, A. 1992 Areas of ignorance and confusion in color science. *Behav. Brain Sci.* **15**, 49–50.

Smith, V. C. & Pokorny, J. 1972 Spectral sensitivity of color-blind observers and the cone photopigments. *Vision Res.* **12**, 2059–2071.

Smith, V. C. & Pokorny, J. 1975 Spectral sensitivity of the foveal cone photopigments between 400 and 500 nm. *Vision Res.* **15**, 161–171.

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

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

Valberg, A. & Lange-Malecki, B. 1990 "Colour constancy" in Mondrian patterns: a partial cancellation of physical chromaticity shifts by simultaneous contrast. *Vision Res.* **30**, 371–380.

West, G. 1979 Color perception and the limits of color constancy. *J. math. Biol.* **8**, 47–53.

West, G. & Brill, M. H. 1982 Necessary and sufficient conditions for Von Kries chromatic adaptation to give color constancy. *J. math. Biol.* **15**, 249–258.

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. & Stiles, W. S. 1982 *Color science: concepts and methods, quantitative data and formulae*. New York: John Wiley & Sons.

Young, T. 1807 *A course of lectures on natural philosophy and the mechanical arts*, vol. I, Lecture XXXVIII. London: Joseph Johnson.

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* Colour coding in the cerebral cortex: the responses of wavelength-selective and colour-coded cells in monkey visual cortex to changes in wavelength composition. *Neuroscience* **9**, 767–781.

Received 6 April 1994; accepted 26 April 1994