



Does colour constancy exist?

David H. Foster

Visual and Computational Neuroscience Research Group, University of Manchester Institute of Science and Technology, Manchester, M60 1QD, UK

For a stable visual world, the colours of objects should appear the same under different lights. This property of colour constancy has been assumed to be fundamental to vision, and many experimental attempts have been made to quantify it. I contend here, however, that the usual methods of measurement are either too coarse or concentrate not on colour constancy itself, but on other, complementary aspects of scene perception. Whether colour constancy exists other than in nominal terms remains unclear.

If we look at a patch of green grass under a blue sky and then later at sunset, the colour of the grass seems unchanged. But the reflected light reaching the eye has a very different spectrum in the two situations: in the first more blue and the second more red. Colour constancy refers to the constant appearance of object or surface colour despite changes in the colour of the illumination, and (in modern usage) scene composition and configuration [1]. Colour constancy is assumed to be fundamental to colour vision, allowing functions ranging from the primitive foraging for ripe fruit in a tree canopy, to the modern use of textiles, cosmetics, and packaging. In theory, it is necessary for producing a stable visual world. And yet, introspection aside, how do we know that colour constancy really exists?

Persuasive demonstrations of effects related to colour constancy by Gaspard Monge [2] in the 18th century, and later by Edwin Land [3] in the mid-20th century, have since given way to progressively more sophisticated experiments aimed at quantitatively determining the degree to which colour constancy holds, in both the adapted and unadapted eye. Nevertheless, I shall argue here that what most of these experiments have measured is not strictly colour constancy – that is, the constant appearance of surface colour – but other aspects of our scene-perception, such as the relationship between surface colours, or illumination colour. Some of these aspects have been quantified precisely, but they leave open the question of whether colour constancy itself exists, other than in nominal terms.

What follows is a summary of the evidence for colour constancy gathered to date. First, I give a brief explanation of what is needed for colour constancy to be achieved. Then, I describe the three main methods of testing it and identify their intrinsic limitations. Although these limitations are also relevant to testing colour constancy under changes in scene composition and configuration, for

simplicity the emphasis is on the effects of illuminant changes. Finally, I suggest some possible ways forward.

What does colour constancy need?

The perceived colour of a surface depends on its spectral-reflectance properties – the proportion of incident light reflected at each wavelength of the spectrum – mediated by the long-, medium-, and short-wavelength-sensitive cone receptors of the eye (see Box 1 and [4]). But, if a surface is uniform and presented in isolation in a dark field, it is impossible to tell whether its perceived colour is due to its own reflecting properties or to the spectrum of the illuminating light: a red paper in white light can look the same as a white paper in red light.

In general, the challenge for the visual system is to estimate or eliminate the effect of the illuminant [5], so that surface reflectance can be deduced and, therefore, its true colour perceived. When several different surfaces are present, the task becomes more feasible. For example, in Land's Retinex colour-constancy schemes [6,7], receptor responses to light from an individual surface were thought to be normalized by certain combinations of responses to light from other surfaces, so as to reduce the influence of the illuminant. However, the visual system need not implement these particular calculations (e.g. [8]), and in fact many other constancy algorithms have since been proposed [1]. Even so, extracting information about surface reflectance remains a necessary condition for colour constancy. Without this information, colour constancy loses much of its rationale as a manifestation of a stable visual world.

To determine the extent to which colour constancy holds, therefore, perceived surface colour needs to be measured by methods sufficiently sensitive to changes in surface reflectance, as opposed to other scene properties. If reflectance can be significantly altered without affecting the measurement, then the measurement is only partial and does not properly limit colour constancy.

Colour naming

The most direct approach to measuring perceived surface colour tests whether colour names are used appropriately [9–11]. Although the vocabulary is normally restricted to certain basic colour terms or categories, the principle is general. For example, subjects given a free choice of names might label a surface with a strong light blue colour under one illuminant as 'copen blue'. The degree to which colour constancy holds can then be determined by measuring how accurately they use the label 'copen blue' for the same surface under a different illuminant.

Corresponding author: David H. Foster (d.h.foster@umist.ac.uk).

Box 1. Limitations of human colour perception

The eye's ability to extract information about spectral distributions of light is fundamentally constrained by the three classes of cone receptors (Fig. 1). In general, a spectrum needs an infinite number of values to specify its energy at each wavelength, but, when sampled by the cones, this information is reduced to just three values, a triplet specifying excitations in each of the three classes. Consequently, there are theoretically an infinite number of spectra that the visual system confuses (the phenomenon of 'metamerism'). Because a spectrum cannot be recovered uniquely from a triplet of numbers, or some recoded version of this, ideal colour constancy in the sense of extracting a surface spectral reflectance is physically impossible.

Therefore, in the literature, as well as in this article, references to visual estimates of an illuminant or surface reflectance are usually taken to mean within the sampling limits of the cones. In this reduced sense, colour constancy can then be interpreted as the same reflecting surface producing somewhere in the visual system the same triplet of numbers, independent of the illuminant.

Nevertheless, the problem of metamerism persists. In practice, spectra do not vary arbitrarily with wavelength, but they still need more than three numbers to specify them, so two real reflecting materials that match under one illuminant need not match under another (which has troubled dyers and colorists, and clothes shoppers, for many years). As a result, even this reduced version of colour constancy cannot be fully realised.

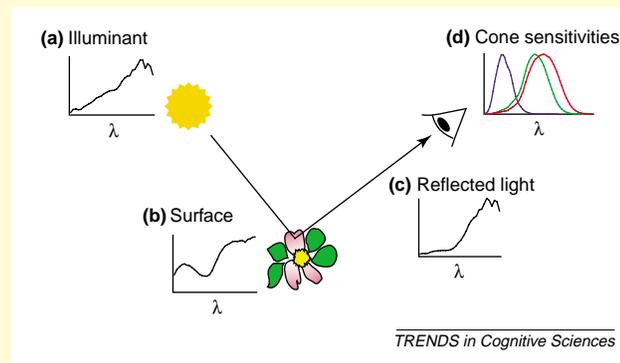


Fig. 1. Sampling spectra. Light from an illuminant with spectrum (a) is reflected by a surface with spectral reflectance (b) (λ = wavelength). The reflected light with new spectrum (c) is absorbed by the eye's cone receptors with spectral sensitivities (d). From this activity in the cones, the visual system has to estimate the surface reflectance (b). Note the change in spectral shape from (b) to (c).

Unfortunately, even when the illuminant is constant, subjects are likely to use the same name for other shades of strong light blue. The basic limitation of the method is that there are many more distinguishable surface colours – perhaps more than two million [12] – than subjects can normally name accurately or consistently [9]. As a measure, it is insufficiently sensitive to differences in surface reflectance, although it has been used to investigate subjects having inherited or acquired defects of vision affecting their colour-discrimination ability [13–15].

It is possible to increase the precision of colour naming by adding ratings [16] ('How good an example of this particular colour is that stimulus?'), but this modified task is more demanding and complications arise in analysing the resulting data (J.M. Speigle, PhD thesis, University of California Santa Barbara, 1997). A similar problem occurs with scaling methods applied to judgments of hue,

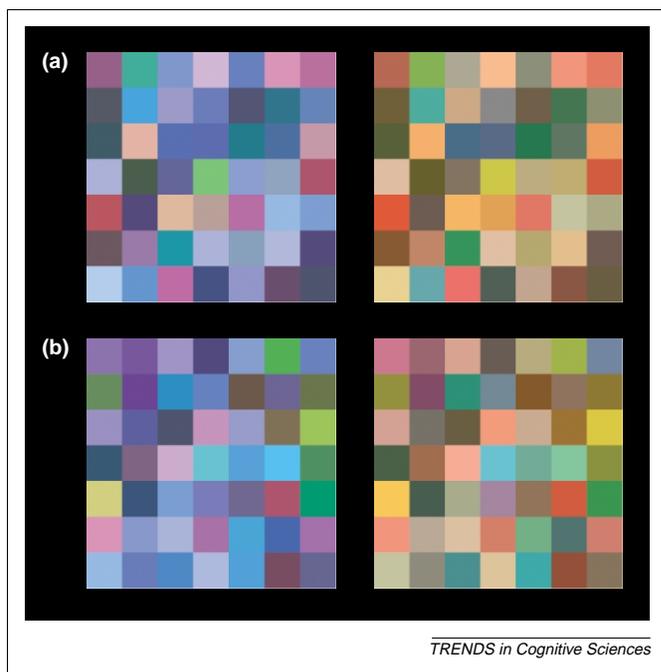


Fig. 1. Matching coloured patterns under different lights. The patterns on the left and right consist of the same Munsell papers illuminated by blue sky and the setting sun, respectively [49]. The subject's task is to adjust the colour of the central patch of the right-hand pattern so that it appears to be cut from the same paper as the central patch of the left-hand pattern: (a) and (b) are successful and unsuccessful surface-colour matches, respectively. (Note, it is unlikely that either will look like a match to the reader because they are not seen in the experimental conditions of controlled illumination.)

saturation, brightness and whiteness under illuminant changes (J.M. Speigle, op. cit.).

By contrast, the second and third experimental approaches, described and analysed in the following sections, are capable of delivering the required precision, or close to it, without modification. These are the methods that have most often been used to determine the degree to which colour constancy holds.

Matching surfaces

The second main approach to measuring perceived surface colour tests how well subjects can make matches between coloured surfaces under different lights ('asymmetric colour matching') [17–20]. The scenes usually comprise multiple, uniformly reflecting surfaces (as in Fig. 1), which are presented simultaneously, side by side, or sequentially. Subjects are asked to adjust 'red', 'green' and 'blue' controls (or their equivalent) to make the colour of a particular surface (e.g. the centre surface) in one scene under one light the same as the colour of a particular surface, usually in a similar scene, under another light, as if it were 'cut from the same piece of paper' [17]. Instructions of this kind can be important, as subjects might concentrate on the colour of the light – hue, saturation, brightness – rather than on the colour of the surface [10,17,18,21–23]. To make this distinction clearer, look at a surface half in shadow: the shadowed region has the same surface colour as the unshadowed region but it simultaneously appears less bright and generally more blue [22].

In making surface-colour matches, subjects may be

encouraged to make comparisons with other surfaces in the scene that by chance have similar colours [17,22,23]. For example, in the left-hand pattern of Fig. 1a, they might compare the centre patch with the second patch along from the top-left corner and then do the same in the right-hand pattern of Fig. 1a. But other matching strategies may be used, and matching itself can also be replaced by a more natural, forced-choice procedure [24]. Levels of matching performance with computer-simulated scenes can reach 79–87% of the ideal value corresponding to perfect colour constancy [20,23].

Asymmetric colour matching is often treated as measuring colour constancy, but what it really measures is something weaker. What this might be is explained next.

Relative versus absolute judgments

The problem with colour matching is that it can guarantee only the equivalence of two stimuli, not their identical colour perceptions. This is because to make a surface-colour match, subjects need merely to judge how the colour of one surface *relates* to the colour of one or more other surfaces or, indeed, to the scene as a whole, first under one illuminant and then under another, or under changes in composition or configuration of the scene.

To see this, consider Fig. 1a again. Suppose that the reflectance of the centre paper in the left-hand pattern is changed by placing a moderate bluish or yellowish filter over it. Providing the same change is made to the other papers in the pattern, there should be little effect on the match represented by the centre patch of the right-hand pattern, if matches depend only on relational judgments. Because information about individual surface reflectance is not needed to succeed at asymmetric colour matching, it cannot measure colour constancy (in extreme circumstances, a perfect surface-colour match can be made in the complete absence of colour constancy [1]). The same principle would apply to any relative judgment – we can decide whether the angle between a pair of lines is the same as that between another pair, without knowing the absolute orientation of any individual line.

Arguably, this manipulation of reflectances reveals only an inherent confounding of illuminant and reflectance spectra. But the problem is not to do with the stimuli; rather, it is the task, as can be seen by changing it to a colour-naming one. Despite a perfect surface-colour match in Fig. 1a, many subjects would give the centre patches

different surface-colour names. (Note that subjects were matching the surface colour under different illuminants, not matching hue, saturation and brightness, so Fig. 1a will not appear as a perfect match to the reader, who sees the images on the page under a single illuminant.)

Although colour constancy may not be needed for successful surface-colour matching, might subjects use it anyway? If there are enough surfaces in the scene, then, in theory, the illuminant can be reliably estimated, and therefore, in turn, can surface reflectance. For example, the scene could be assumed to be colorimetrically unbiased, that is, no particular colour predominates (the ‘grey-world’ assumption) [7,25]. The spatial average colour then coincides with the illuminant colour. With natural scenes, higher-order statistics could also be exploited [26].

Even so, showing that the illuminant could be estimated is not the same as showing that it is estimated, or, more importantly, that perceived surface colour is constant. A critical test is whether matching is worse when the illuminant cannot be reliably estimated. The minimalist ‘scenes’ of just two surfaces in Figs. 2a and b provide little information about the illuminant. Yet surface-colour matches can certainly be made with such scenes (Fig. 2a), and they do not seem much worse than with scenes containing many surfaces [21,22] (as in Fig. 1a, the illuminants on the left and right are different). In an operational version of this matching task, described in the next section, performance was essentially the same [27].

Relational colour constancy

The ability of subjects to make accurate relational judgments independent of the illuminant, as with these minimalist scenes, has been attributed to ‘relational colour constancy’ [28]; that is, the constancy of perceived colour relations under different illuminants. It was initially conceived for situations in which scenes differ only in illuminant, so that one surface could be related to another – the simplest possible comparison. But it can be extended to situations in which scene composition and configuration change by relating one surface to multiple other surfaces or to some average over the scene as a whole (without estimating the illuminant, as with colour constancy).

Relational colour constancy can be interpreted operationally: subjects are simply asked to discriminate between illuminant and surface-reflectance changes [28,29]. Performance in one such task has been found to be fast, accurate and effortless, suggesting that some aspects of surface-colour information are processed in parallel over the visual field [27].

Relational colour constancy has a plausible physiological substrate: the ratios of cone-receptor excitations generated in response to light reflected from pairs of surfaces or groups of surfaces. Such ratios, which can also be calculated across post-receptoral combinations and spatial averages of cone signals, have the remarkable property of being almost exactly invariant under changes in illuminant, both with natural scenes [30] and with artificial scenes of coloured papers [28], and might explain performance in several colour-vision tasks [27,31,32], including asymmetric colour matching [21]. Significantly, they provide a compelling cue to subjects trying to

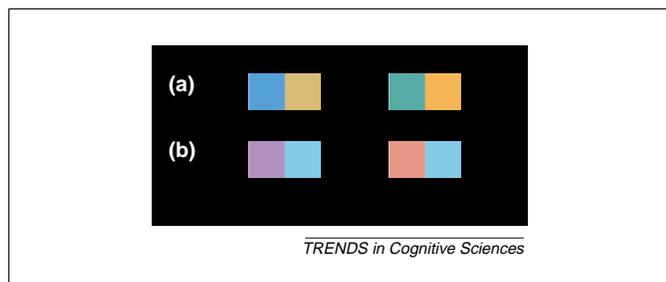


Fig. 2. Matching ‘minimalist’ coloured patterns under different lights. The patterns on the left and right consist of the same Munsell papers illuminated by blue sky and the setting sun, respectively. The subject’s task is to adjust the colour of the right-hand patch of the right-hand pattern so that it appears to be cut from the same paper as the right-hand patch of the left-hand pattern: (a) and (b) are successful and unsuccessful surface-colour matches, respectively.

distinguish between illuminant and reflectance changes in scenes, even when they sometimes correspond to highly unlikely natural events [33]. Consistent with their proposed role in relational colour constancy, cone-excitation ratios are not informative about individual surface reflectances [21].

The perceived relations between the colours of surfaces should not be confused with what are sometimes technically called ‘related colours’, exemplified by brown and grey, which require other colours for their perception (in isolation, brown becomes orange and grey becomes white). As already emphasized, judging the relationship between the colours of surfaces is not the same as judging surface colours themselves, whether they are ‘related colours’ or not. For example, in both the left- and right-hand patterns of Fig. 1a, it is possible to decide that the centre patch appears lighter and more yellow than the patch to its immediate left, even though that latter patch appears dark blue in the left-hand pattern and brown in the right-hand pattern. This is then an illustration of relational colour constancy in the absence of colour constancy. The converse should be impossible.

Judging white

The third main approach to measuring perceived surface colour tests how well subjects can adjust the colour of a surface within a scene so that it appears white (‘achromatic adjustment’) [34–37]. Although sometimes presented as measuring colour constancy, this method in fact provides data only on the appearance of neutral (i.e. greyscale) surfaces: indirect arguments can be made about its implications for asymmetric colour matching, but they require significant additional assumptions [38]. Insofar as subjects regard the method as adjusting surface colour, it effectively records their estimate of illuminant colour at that point in the scene (Figs. 3a and b). The chromatic context of the surface might bias this estimate, but it is still correlated with the colour of the illuminant. The extent of this bias can be exploited as a quantitative probe of the effects of scene structure [34,37]. If subjects compare the surface with an imaginary standard, then the task is similar to asymmetric colour matching, but is less

precise [16]. Levels of performance can reach 83% of the ideal unbiased value with nearly natural scenes [37].

By design, achromatic adjustment provides no direct data on the appearance of individual coloured surfaces in the scene, for it is generally possible to alter the reflectance of a particular surface and keep the achromatic setting constant by adjusting the reflectances of other surfaces. Again, because individual surface reflectance is not needed to succeed at achromatic adjustment, it also cannot measure colour constancy. It can, however, be used to make inferences about the appearance of neutral surfaces themselves, as can other adjustments, for example, of a particular yellow that is ‘neither reddish nor greenish’ [39], or a red (‘neither yellowish nor bluish’), and so on [40]. But, apart from white, the task is limited to the four unique hues (red, green, blue and yellow).

Crucially, judgments of illuminant colour can depend on many sources of information [41]. Assumptions about a grey world and other, higher-order, statistical properties of scenes [26] have already been mentioned. Another common assumption is that the surface with the highest luminance is white [6]. When pitted against each other, information about a spatial average seems to take priority over that from the highest-luminance region [42], even though the grey-world assumption can prove misleading [43]. Other cues such as mutual illumination [44] and specularities [45] can also be used to infer the illuminant. This complex combination of cues suggests that there is probably no single physiological mechanism underlying illuminant estimation.

Multiple kinds of surface-colour information

These second and third experimental approaches – asymmetric colour matching and achromatic adjustment – are evidently intrinsically limited. They each give partial but complementary descriptions: neither alone adequately determines perceived surface colour.

These limitations could, however, be a simple consequence of the fact that colour constancy is not a unitary phenomenon [46,47]. On the one hand, perceived colour relations provide precise information about the similarities and differences between surfaces, while indicating little about the illuminant [48]. On the other hand, inferences about the illuminant allow a more complete perceptual representation, but one that might not be very accurate or essential to every task. To take an earlier example, judging the degree of ripeness of fruit in a tree canopy need not require the extraction of information about its surface reflectance, merely a comparison of its perceived colour in relation to that of other fruit or the surrounding foliage. In other circumstances, a coarse categorization that the fruit is red might suffice.

Conclusion

So, does colour constancy exist? Progress has been made in quantifying the extent to which it might hold, but current measurement methods remain incomplete, and their limitations need to be made more explicit. The most direct method, colour naming, might be improved by training subjects to use a larger vocabulary, but, without additional scales or ratings, it is unlikely to achieve the precision of

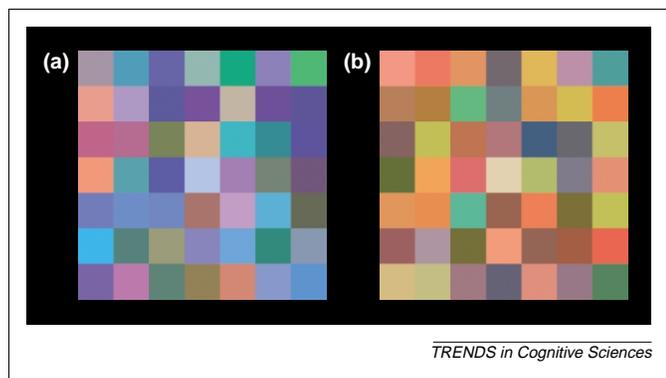


Fig. 3. Adjusting to white under different lights. The patterns on the left and right consist of different Munsell papers illuminated by blue sky and the setting sun, respectively. The subject’s task is to adjust the colour of the central patch of each pattern so that it appears white: (a) and (b) are both successful adjustments to the colour of the illuminant; that is, they record the subject’s estimate of illuminant colour at that point in the scene.

matching procedures. Individually, the other two main methods, of asymmetric colour matching and achromatic adjustment, are not sufficiently specific to detecting changes in surface reflectance, but they might be when taken together. Is there a way, therefore, to combine them naturally into a single measure? This represents an interesting problem – although the illuminant can only be estimated reliably in scenes with many surfaces, it is not clear that surface-colour perception is different in scenes with just a few surfaces. Until that problem is resolved or other specific measurement methods are devised, then whether colour constancy exists, other than in nominal terms, will remain unproven.

Acknowledgements

I am grateful to many colleagues for discussion of the ideas presented here, and to the BBSRC, EPSRC and Wellcome Trust for support.

References

- Maloney, L.T. (1999) Physics-based approaches to modeling surface color perception. In *Color Vision: From Genes to Perception* (Gegenfurtner, K.R. and Sharpe, L.T., eds), pp. 387–416, Cambridge University Press
- Mollon, J. (1985) Colourful notions: studies in scarlet. *Listener* 10 January 1985, 6–7
- Land, E.H. (1959) Color vision and the natural image. *Part I. Proc. Natl. Acad. Sci. U. S. A.* 45, 115–129
- Wandell, B.A. (1995) *Foundations of Vision*, Sinauer Associates
- von Helmholtz, H. (1867) *Handbuch der Physiologischen Optik* (Vol. 2) (1st edn), Leopold Voss, Leipzig. Translation (1909) *Helmholtz's Treatise on Physiological Optics* (3rd edn) (Southall, J.P.C., ed.), Optical Society of America, Washington, DC, 1924, pp. 286–287. Republished 1962, Dover Publications, New York
- Land, E.H. and McCann, J.J. (1971) Lightness and retinex theory. *J. Opt. Soc. Am.* 61, 1–11
- Land, E.H. (1986) Recent advances in Retinex Theory. *Vis. Res.* 26, 7–21
- Brainard, D.H. and Wandell, B.A. (1986) Analysis of the retinex theory of color vision. *J. Opt. Soc. Am. A Opt. Image Sci. Vis.* 3, 1651–1661
- Boynton, R.M. and Olson, C.X. (1987) Locating basic colors in the OSA space. *Color Res. Appl.* 12, 94–105
- Troost, J.M. and de Weert, C.M.M. (1991) Naming versus matching in color constancy. *Percept. Psychophys.* 50, 591–602
- Uchikawa, K. *et al.* (1998) Measurement of color constancy by color memory matching. *Opt. Rev.* 5, 59–63
- Pointer, M.R. and Attridge, G.G. (1998) The number of discernible colours. *Color Res. Appl.* 23, 52–54
- Montag, E.D. (1994) Surface color naming in dichromats. *Vis. Res.* 34, 2137–2151
- Kennard, C. *et al.* (1995) Colour identification and colour constancy are impaired in a patient with incomplete achromatopsia associated with prestriate cortical lesions. *Proc. R. Soc. Lond. B. Biol. Sci.* 260, 169–175
- Zeki, S. *et al.* (1999) The neurological basis of conscious color perception in a blind patient. *Proc. Natl. Acad. Sci. U. S. A.* 96, 14124–14129
- Speigle, J.M. and Brainard, D.H. (1996) *Is color constancy task independent? Proceedings of The Fourth Color Imaging Conference: Color Science, Systems and Applications*, Soc. Imaging Science & Technology, Springfield, pp. 167–172
- Arend, L. and Reeves, A. (1986) Simultaneous color constancy. *J. Opt. Soc. Am. A Opt. Image Sci. Vis.* 3, 1743–1751
- Cornelissen, F.W. and Brenner, E. (1995) Simultaneous colour constancy revisited: an analysis of viewing strategies. *Vis. Res.* 35, 2431–2448
- Brainard, D.H. *et al.* (1997) Color constancy in the nearly natural image: 1. Asymmetric matches. *J. Opt. Soc. Am. A Opt. Image Sci. Vis.* 14, 2091–2110
- Foster, D.H. *et al.* (2001) Colour constancy from temporal cues: better matches with less variability under fast illuminant changes. *Vis. Res.* 41, 285–293
- Tiplitz Blackwell, K. and Buchsbaum, G. (1988) Quantitative studies of color constancy. *J. Opt. Soc. Am. A Opt. Image Sci. Vis.* 5, 1772–1780
- Arend, L.E. Jr *et al.* (1991) Simultaneous color constancy: papers with diverse Munsell values. *J. Opt. Soc. Am. A Opt. Image Sci. Vis.* 8, 661–672
- Bäumli, K-H. (1999) Simultaneous color constancy: how surface color perception varies with the illuminant. *Vis. Res.* 39, 1531–1550
- Bramwell, D.I. and Hurlbert, A.C. (1996) Measurements of colour constancy by using a forced-choice matching technique. *Perception* 25, 229–241
- Buchsbaum, G. (1980) A spatial processor model for object colour perception. *J. Franklin Institute* 310, 1–26
- Golz, J. and MacLeod, D.I.A. (2002) Influence of scene statistics on colour constancy. *Nature* 415, 637–640
- Foster, D.H. *et al.* (2001) Parallel detection of violations of color constancy. *Proc. Natl. Acad. Sci. U. S. A.* 98, 8151–8156
- Foster, D.H. and Nascimento, S.M.C. (1994) Relational colour constancy from invariant cone-excitation ratios. *Proc. R. Soc. Lond. B. Biol. Sci.* 257, 115–121
- Craven, B.J. and Foster, D.H. (1992) An operational approach to colour constancy. *Vis. Res.* 32, 1359–1366
- Nascimento, S.M.C. *et al.* (2002) Statistics of spatial cone-excitation ratios in natural scenes. *J. Opt. Soc. Am. A Opt. Image Sci. Vis.* 19, 1484–1490
- Westland, S. and Ripamonti, C. (2000) Invariant cone-excitation ratios may predict transparency. *J. Opt. Soc. Am. A Opt. Image Sci. Vis.* 17, 255–264
- Ripamonti, C. and Westland, S. Prediction of transparency perception based on cone-excitation ratios. *J. Opt. Soc. Am. A Opt. Image Sci. Vis.* (in press)
- Nascimento, S.M.C. and Foster, D.H. (1997) Detecting natural changes of cone-excitation ratios in simple and complex coloured images. *Proc. R. Soc. Lond. B. Biol. Sci.* 264, 1395–1402
- Brainard, D.H. (1998) Color constancy in the nearly natural image: 2. Achromatic loci. *J. Opt. Soc. Am. A Opt. Image Sci. Vis.* 15, 307–325
- Chichilnisky, E.J. and Wandell, B.A. (1999) Trichromatic opponent color classification. *Vis. Res.* 39, 3444–3458
- Bäumli, K-H. (1999) Color constancy: the role of image surfaces in illuminant adjustment. *J. Opt. Soc. Am. A Opt. Image Sci. Vis.* 16, 1521–1530
- Kraft, J.M. and Brainard, D.H. (1999) Mechanisms of color constancy under nearly natural viewing. *Proc. Natl. Acad. Sci. U. S. A.* 96, 307–312
- Speigle, J.M. and Brainard, D.H. (1999) Predicting color from gray: the relationship between achromatic adjustment and asymmetric matching. *J. Opt. Soc. Am. A Opt. Image Sci. Vis.* 16, 2370–2376
- Shevell, S.K. and Wei, J. (1998) Chromatic induction: border contrast or adaptation to surrounding light? *Vis. Res.* 38, 1561–1566
- Arend, L.E. (1993) How much does illuminant color affect unattributed colors? *J. Opt. Soc. Am. A Opt. Image Sci. Vis.* 10, 2134–2147
- Maloney, L.T. (2002) Illuminant estimation as cue combination. *J. Vision* 2, 493–504
- Linnell, K.J. and Foster, D.H. (2002) Scene articulation: dependence of illuminant estimates on number of surfaces. *Perception* 31, 151–159
- Brown, R.O. (1994) The world is not grey. *Invest. Ophthalmol. Vis. Sci.* 35, 2165
- Bloj, M.G. *et al.* (1999) Perception of three-dimensional shape influences colour perception through mutual illumination. *Nature* 402, 877–879
- Yang, J.N. and Maloney, L.T. (2001) Illuminant cues in surface color perception: tests of three candidate cues. *Vis. Res.* 41, 2581–2600
- Walsh, V. (1999) How does the cortex construct color? *Proc. Natl. Acad. Sci. U. S. A.* 96, 13594–13596
- Moutoussis, K. and Zeki, S. (2000) A psychophysical dissection of the brain sites involved in color-generating comparisons. *Proc. Natl. Acad. Sci. U. S. A.* 97, 8069–8074
- MacLeod, D.I.A. (2003) New dimensions in color perception. *Trends Cogn. Sci.* 7, 97–99
- Wyszecki, G. and Stiles, W.S. (1982) *Color Science: Concepts and Methods, Quantitative Data and Formulae*, John Wiley & Sons