

A computational analysis of colour constancy

Donald I.A. MacLeod
Psychology
UCSD
La Jolla, CA 92093
USA

Jürgen Golz
Psychology
University of Kiel
Kiel,
Germany

The visual stimulus associated with an object surface varies with the illumination falling on the object. Constancy of perceived surface colour under changing illumination implies that the *illumination-dependent mapping* from surface reflectance to stimulus is cancelled by a compensatory variation in the mapping from retinal stimulus to perceived surface colour. To accomplish this compensation, the visual system must in some sense internalize, or at least exploit, the regularities of the environmental interaction between observed surfaces and the incident light as these jointly determine the visual stimulus.

In this chapter we begin by asking: What kinds of mappings from cone excitation triplets to surface colour appearance are required to make the latter constant under changes of illumination? This is a question not about the visual process itself, but about the environmental interplay between lights and surfaces, for which the visual system must compensate if it is to achieve colour constancy. In modern parlance (Marr, 1982; Hurlbert, 1998), it is a computational issue. We discuss it first for various idealized worlds that are more or less mathematically tractable (§1, and Appendix), and next for the real world (§2). These investigations reveal the nature of the compensation that a colour-constant visual system should perform in each world. We next consider (§3) what mappings the visual system actually makes in the interests of colour constancy, and how these differ from the optimal ones.

Appropriate compensation for the illuminant presupposes that the illuminant is appropriately estimated, and in the last section (§4) we turn from the compensation problem to the equally critical (and perhaps more challenging) *estimation* problem that a colour-constant visual system must surmount. What stimulus information is available to estimate the illuminant and thereby select the appropriate compensatory mapping? Among many visible effects of changing illumination we concentrate here on one class of cue: the information inherent in the illumination-dependent mapping itself. We use the results of §1 and §2 to investigate how the statistics of the cone excitations associated with individual scene elements can help in estimating the illuminant, and we present experimental results that suggest that certain statistical sources of information are indeed exploited in solving the estimation problem.

1. The Compensation Problem: Ideal Worlds and their Illumination-dependent Mappings

Vision generally depends on light that is incident on object surfaces, gets reflected toward the eye, and is absorbed in the retinal photoreceptors. At each wavelength in the visible

spectrum, the light that gets to the eye is the product of (a) what is incident on the surface, represented by the spectral power distribution $I(\lambda)$ of the incident illumination (which is constrained, but not strictly determined, by the intensity and colour of the light source) and (b) the fraction reflected by the surface, given by the surface's spectral reflectance function $S(\lambda)$. Thus the mapping from local surface reflectance to the associated retinal stimulus (a triplet of cone excitations) is illumination-dependent. To achieve an illumination-invariant representation of surface colour, the visual system must introduce a compensatory (hence also illumination-dependent) mapping from retinal stimulus to perceived colour.

Can this be done—and if so, what is the nature of the required compensatory mapping? We investigate this question for four quite different idealizations of the visual environment, discussing these in turn in the context of a very simple proposal about the nature of the visual system's compensatory mapping.

Normalization

The effects of changes in the *intensity* of the illuminant could in principle be compensated by a reciprocal adjustment of sensitivity, for instance at the photoreceptor level. It has been attractive to extend this idea (e.g. Helmholtz, 1896; Land and McCann, 1971) to changes in the colour balance of the illuminant. A change in the colour balance of the illuminant might be effectively signaled by its effect on the predominant colour of a scene, since it affects the chromaticity of every surface lit by the source, and it might be possible to compensate for these effects by applying reciprocal sensitivity adjustments selectively to different regions of the spectrum. The idea that constancy is implemented by operations equivalent to a scaling of sensitivity to incident light is a recurring one in both current and earlier discussions of the topic, and it will recur in our discussion here. To avoid premature implications about the physiological mechanisms responsible, we refer to it as the normalization model¹. Normalization, then, could be implemented either by a scaling of sensitivity to light at the photoreceptor level, or by later and perhaps much more complex processes, so long as these processes finally change colour appearance in the same way that photoreceptor sensitivity scaling would.

The term normalization might suggest that the change in effective sensitivity is reciprocal with some measure of the cone excitation generated by the observed scene—a measure which

¹ Hurlbert (1998) calls such models "lightness models". The normalization factors associated with each cone excitation are often referred to as "von Kries coefficients".

defines an implicit estimate of the illuminant whose effects the normalizing compensation cancels. But a compensatory normalization could be associated with any illuminant estimate, no matter how the estimate is derived, and no matter whether it is accurate or not. Thus the normalization factors need not be reciprocally related to stimulus intensity measures, nor need they each depend only on the excitation of the photoreceptor to which they apply.

The normalization model has a clear rationale if the photoreceptors each respond to a band within the visible spectrum so narrow that illuminant power or surface reflectance are always effectively constant across the band (a situation we consider below). In the real visual system, however, the spectral bands to which the three types of cone photoreceptor are sensitive are not very narrow (for good reason, since making them narrow would severely impair visual sensitivity by neglecting all incident light energy outside the sensitive band). It has been increasingly emphasized (e.g. Brill, 1978; Buchsbaum, 1980; Worthey and Brill, 1986; Hurlbert, 1998; Maloney, this volume) that determining the colour of a surface under changing or unknown illumination is a difficult problem for a visual system with broad-band photoreceptors—so difficult, in fact, as to be in general unsolvable.

The Chaotic World

The nature of the problem can be appreciated by considering two metameric yellow illuminants, visually similar but physically different. Suppose that each supplies light in only two narrow spectral bands, one red and one green, but that the bands supplied by the two light sources are slightly offset from one another, by enough that they do not overlap. Denote the bands by R1 and G1 for illuminant 1 and R2 and G2 for illuminant 2. Imagine each of these sources in turn lighting the following two yellow surfaces. One surface, (which we may call surface R1G2), has a spectral reflectance that takes nonzero values only in the bands R1 and G2. The other yellow surface, surface R2G1, reflects only in the other two bands. Surface R1G2 is physically red under source 1: it only reflects the red band R1 of radiations to the eye, and not the second, green band G1 supplied by source 1.

Surface R2G1, on the other hand, turns green under source 1, because only its green reflection band G1, and not the red one R2, is present in the spectrum of source 1. The situation is completely reversed under source 2. Now surface R2G1 reflects red light only, and surface R1G2 reflects only green. This example is enough to show how there are, in general, no

constraints at all on how a change of light source will change the colour of a surface: any set of colours can be transformed into any other set. In this example the two surfaces, as well as the two illuminants, can be indistinguishable yellows under a white light source (because the reflection bands can be arbitrarily narrow and correspondingly closely spaced), but they can assume any two chromaticities, limited only by the chromaticities of their reflection bands, when a suitably contrived illuminant is switched on to light them. The chromaticities they assume adopt can span the physically realizable gamut, and the changes due to changing the illuminant are independent for different surfaces, in the sense that any set of surface colours can be transformed into any other set. Such chaotic colour changes would defeat any attempt to achieve colour constancy by any systematic remapping of retinal stimuli onto subjective appearances. If a change of light source could be relied on to transform the intensity and chromaticity of retinal stimuli in a more or less orderly way, the visual system might be able to keep the surface colours constant by reversing that orderly transformation or mapping. But you can't create order by unscrambling total chaos. So colour constancy is not in general possible! But of course, colour constancy does occur, because the chaotic world of this scenario is not the real world we live in. This raises the question: under what real world constraints would a simple compensating process, such as normalization through sensitivity adjustments, give us colour constancy?

The Three-Band World

The answer to this question depends on whether we require perfect colour constancy, or a merely a useful approximation to constancy. For perfect or nearly perfect constancy through normalization, some extremely artificial conditions must be satisfied (Worthey and Brill, 1986). First, the three cone types must respond to non-overlapping bands in the illuminant spectra. This will happen if the cone spectral sensitivities have no overlap, or if the illuminant energy is always confined to discrete non-overlapping bands, with one cone type sensitive to each such band. Second, illuminants must never vary significantly in their relative distribution of power across any one such band; they can vary only in the amounts of power they radiate in the three bands.² Under this “three-band” scenario, the cone excitations for each cone type vary in proportion to the total illuminant power within the corresponding spectral band, and this variation can be compensated—for all surfaces at once—by a reciprocal adjustment of

² Alternatively but still less plausibly, the surface reflectances, rather than the illuminant power spectral densities, could be uniform in this sense for each band.

sensitivity appropriate to the scene as a whole. Thus the three-band world exemplifies what we call *normalization-compatible mapping*: the possible illumination-dependent mappings from surface reflectance to cone excitations differ only in the scaling factors applied to the three cone excitations.

A convenient way of picturing the effects of changes in intensity and in colour balance of the illuminant is to consider a logarithmic cone excitation space in which three coordinates represent the excitations of the long-, mid- and short-wavelength cones on a log scale. The position of the point representing any surface moves parallel to the positive diagonal under variations in illuminant intensity. If the colour of the illuminant changes, then surfaces move in the two orthogonal, “chromatic” directions as well. But in the three-band world, the constellation of points representing a set of surfaces in a scene moves rigidly under all changes in illuminant. This happens because a change of illuminant changes the excitation of, say, the long-wavelength cones by exactly the same factor for all surfaces in the scene—the factor by which the illuminant power in the red band is changed. Each cone excitation is scaled by the power being supplied within its associated spectral band. The logs of these three cone-specific scaling factors are added to the corresponding coordinate in the logarithmic plot.

If we assume that a change of illuminant not only changes the average cone excitation within the scene but also triggers a reciprocal adjustment of sensitivity in each cone type, then the changes of log sensitivity would exactly cancel the change of log cone excitation for each surface, and the resulting normalized excitations would provide for each surface a representation independent of the illuminant colour. Geometrically, in log cone excitation space, the constellation of points representing the surfaces would be translated back in all three dimensions and would be brought back to where it started.

There are many objections, mainly of a non-computational character, to the adoption of reciprocal adjustments of photoreceptor sensitivity as a model for constancy (e.g. MacLeod, 1985). A couple of these merit at least brief acknowledgement, prior to leaving them aside for the moment. First, cone sensitivity regulation is mainly local (MacLeod, Williams and Makous, 1992) rather than being based on excitation averaged over a scene. Local sensitivity regulation, however, can also produce an illumination-invariant representation—but of the local contrast at edges within the scene, not of local lightness and colour directly. In this way it can form the initial basis of a “retinex” model for colour constancy (Land and McCann, 1971). Second, the model normalizes out variations of lightness and colour balance inherent to the surfaces of a

scene, as well as those due to the illuminant. We return briefly to this deficiency in discussing the “Gray World assumption” in §4 below.

In the context of a computational analysis, a more serious embarrassment is that the three-band scenario is hardly defensible as even a crude approximation to the real world. Obviously we need to consider the possibilities for constancy in a slightly more natural scenario, in which cones, and especially surfaces and illuminants, have fairly broad and smooth spectral characteristics. Any principles of illumination-dependent mapping found to hold in such a world are more likely to be applicable in the real one.

The Linear World

In the search for order in the chromatic universe, the approach usually taken is to approximate naturally occurring spectral reflectance functions and spectral power distributions by a weighted sum of suitably chosen basis functions, the latter forming a fixed set common to all spectra (e.g. Sallstrom, 1973; Brill, 1978; Buchsbaum, 1980; Maloney, this volume). To account for trichromatic perception of surface colour, we need, minimally, three basis functions for surface reflectance. Schemes like this that use linear models of reflectance and power distributions have been investigated quite a lot. And one of the first things to emerge is that the kind of normalization that worked for constancy in the three-band scenario won't do at all, in principle, for the linear world. This is not surprising, since it is only in the hopelessly artificial three-band world that the cone excitations from different surfaces all change by the same factor when the light changes. Yet illumination-dependent mapping in the linear world remains formally fairly simple. Each of the three surface reflectance components generates an illumination-dependent triplet of cone excitations (a cone excitation vector). Each such cone excitation vector is the weighted sum of N_i fixed vectors contributed by the N_i possible illuminant components: the fixed vectors are the cone excitation triplets resulting from the interaction of the illuminant component (in unit quantity) with the surface component under consideration, and the weights are determined by the particular illuminant's composition. The visual system can recover an illumination-invariant representation of surface colour (its composition in terms of the three basis functions for reflectance) if it multiplies the cone excitation triplet by the inverse of an illuminant-specific 3-by-3 matrix that contains the cone excitation vectors that would be generated by unit values of each surface component under the prevailing illumination.

The approximation of natural reflectance functions with a combination of three basis functions might not be considered satisfactory (it accounts for only 60% of the chromatic variance in a haphazard sample of natural surface reflectance functions spectro-radiometrically measured by Brown, this volume), but the linear world is at least a great improvement over the three-band world in this respect, and as a result it has come to dominate discussions of colour constancy from a computational perspective. Yet the linear world has serious shortcomings as a theoretical framework.

At a purely descriptive level, the linear description both profits by and suffers from an arbitrary element in the choice of the basis functions. Free choice of the basis functions permits more accurate description of a given set of spectra with a small number of weight parameters. On the other hand, optimization of the basis functions for a particular target set of spectra (such as Munsell papers, or natural terrains, or haphazard collections of interesting objects from one environment or another) is not a principled justification for their choice for general application. A more serious shortcoming is that the world of linearly modelled surface and illuminant spectra harbours a fundamental inconsistency: when visually identical stimuli originate from physically different sources, their linear world approximations will generally be different in ways that should make them visually distinguishable. This inconsistency has its origin in an awkward general feature of the linear world: even with the minimal three degrees of freedom each for light source and surface, the stimulus spectra have six degrees of freedom; this is more than is necessary or desirable for characterizing them if three are sufficient for the functions that generate them (unless, of course, the linear model happens to be an accurate idealization of reality in this respect—a point that has not been established).

Moreover, the simplicity and order implied in the linear mapping principle are not as complete as might appear. The framework of the linear model places no constraint on the individual components of the reflectance and illumination spectra, so their interactions could in principle be as arbitrary as the ones in the chaotic world. To guarantee order in this sense, the basis functions should be smooth. But as we will see, the same objective can be secured by modeling the surface and illuminant spectra themselves with smooth functions.

The theoretical value of the linear world scenario is also limited. The formal simplicity of the compensation operation (matrix division) is not supported by any known or plausible basis in visual processing. In particular, there is a lack of evidence that the visual system does in any sense internalize the matrices of coefficients specifying cone excitations for different

combinations of surface and illumination components, or that its colour corrections are consistent with such a computational scheme. (We return to this point in §3 below). In general, despite extensive work employing the linear world as an analytical tool, it does not seem to have produced much insight either into the actual principles of illumination-dependent mapping in the environment, or into the visual processes subserving colour constancy.

The fondness of theorists for the linear idealization of the chromatic universe may in part reflect the lack of an alternative. In the *ideal* ideal world, the appropriate illumination-dependent or compensatory transformations would be more intuitively understandable and mechanistically plausible, as is the case with simple normalization, and regularity would be guaranteed by restricting the model to smooth (and not arbitrary) functions throughout. We next introduce one such alternative, which we have found illuminating in the computational analysis of colour constancy.

The Gaussian World

Special cases of the linear world may involve a principled choice of basis functions and guarantee smooth spectra. One such scheme approximates natural spectral reflectance functions (cf. Moon and Spencer, 1945) and spectral power distributions by polynomials in wavelength or some other spectral variable. Trichromacy requires at least second-order polynomials for the surfaces, so a second-order polynomial constraint on the surfaces and illuminants yields a minimal trichromatic world. The Gaussian World we introduce here differs from this in just the following way: it is the *logarithms* of the spectral distributions that are built up from additive components. That is, the log of the radiant power or of the surface reflectance is always describable as a second order polynomial. Depending on the sign of the quadratic term these idealized spectral functions are either Gaussians (with exponentials as a special case), or the reciprocals of Gaussians.³

Like the linear world with three degrees of freedom for surfaces and illuminants, the Gaussian World is a fully (but minimally) trichromatic world. It is an equally crude approximation to reality. But the benefits of abandoning the linear world by taking logs are many:

³With integration over a finite spectral range, we need not be troubled by the fact that the latter functions approach infinity in the limit of infinite wavelength.

- We don't have to worry about negative light emissions or negative reflectances, since the numbers represented by their logarithms can't go negative. And because of their more natural approach to zero the Gaussian descriptions fit natural spectra better than do the comparable linear ones.
- We can include monochromatic radiations in the chromatic universe as a limiting case, something not possible in the linear world.
- There is an important gain in mathematical simplicity and tractability: the proliferation of degrees of freedom that occurs in the linear world in the progression from light and surface to the visual stimulus is here entirely avoided. When incident light is multiplied by surface reflectance, the second-order polynomials descriptive of the log power and the log reflectance simply add to form the log of the retinal stimulus, making this too a second-order polynomial. (In non-logarithmic terms, the product of the Gaussians for surface and illuminant yields another Gaussian.) Thus the function specifying the proximal stimulus still has the same form as the surface spectral reflectance and the illuminant SPD, and the same three degrees of freedom (each of its three parameters being a function of the surface and illuminant parameters), whereas in the linear world it has six.
- Each stimulus chromaticity is associated with the same spectral energy distribution no matter what the light and surface reflectance that combined to generate it, so indistinguishable stimuli are now consistently represented as indistinguishable.
- Most important, the cone spectral sensitivities themselves can be approximated by Gaussians,⁴ and this gives us something not available in the other theoretical worlds: a simple equation (Appendix, Equation (1)) for the cone excitations generated by any surface and illuminant, where the only variables are the parameters specifying cone sensitivity, reflectance and illuminant.

⁴ The Gaussian description of cone sensitivity fails in the tails, but that doesn't matter much if we are interested in broad-band stimuli. This shortcoming can be alleviated by choosing a suitable nonlinear function of wavelength as the spectral variable. We have not explored this, since (for natural surfaces) the errors introduced by approximating the cone sensitivities are much smaller than those inherent in the idealization of the reflectances and illuminants.

The Appendix is an algebraic analysis of illumination-dependent mapping in the Gaussian World. In Equation (1) of the Appendix, cone excitation is a product of three Gaussian factors. One factor represents surface colour, independent of the illuminant. It is a Gaussian function of the spectral distance of the surface's spectral centroid from the lambda max of the cone photoreceptor. This is what we want to preserve in a colour constant visual system. A second factor similarly depends on the illuminant colour. If surface bandwidths are not too narrow, this factor is approximately the same for all surfaces, as it depends only on the relation between the illuminant spectral centroid and the cone's wavelength of peak sensitivity. This means it can be successfully normalized out, in just the way that we saw with the three band world. The equation for cone excitation is completed by another factor that depends both on the illuminant and the surface parameters. In a sense this factor precludes full colour constancy based on normalization, since the normalization factor required to remove it is different for different surfaces within a scene. Luckily though, this factor is independent of the cone's preferred wavelength, which means that it generates variations in effective intensity only, not colour. What this factor means is simply that *when the lighting gets red, the reds get lighter*, relative to the blues and greens.

Illumination-dependent mapping in the Gaussian World therefore has the following characteristics, which are specified quantitatively by the equations in the Appendix.

- (1) *Normalization-compatible mapping for colour.* Independent normalization of the cone photoreceptor excitations can recover the chromatic aspects of surface reflectance accurately for sufficiently broadband surfaces and illuminants. In that sense, normalization could yield colour constancy in this idealized world.

Normalization-compatible mapping of chromatic values has a simple geometric expression. As we noted in discussing the three-band world, normalization-compatible mapping of both colour and intensity implies rigid translation in $\log(L, M, S)$ space under changes of illumination. We can abandon the requirement for correct recovery of intensity by considering Figure 1, where chromatic values of stimuli are represented in a planar cross-section of that space by projecting along parallel lines of variable intensity but constant chromaticity. The axes are the differences in log cone excitation between different pairs of cones. Normalization-compatible mapping of chromatic values requires that under different illuminants, the constellation of points representing a set of surfaces must be translated rigidly in this plane. For broadband colours, the logarithmic plot of the (r, b) coordinates of

MacLeod and Boynton (1979) is practically equivalent to such a plane, and is used below to evaluate the rigidity of shifts for sets of natural colours.

But normalization-compatibility is only approximate, and the equations in the Appendix also show how the Gaussian World violates normalization-compatibility in three distinct and quantitatively specifiable (but generally minor) ways:

- (2) *When the lighting gets red, the reds get lighter.* First, as mentioned above, and as the Appendix demonstrates, there remains (after normalization) a deviation from constancy, occurring for both broadband and narrowband colours, that can be corrected by a change of stimulus intensity alone: surfaces similar to the illuminant in colour are rendered as lighter than surfaces dissimilar from the illuminant in colour. While it precludes rigid translation in log cone excitation space, this does not affect chromatic values and thus preserves rigid translation in Figure 1.

For narrow-band colours, normalization fails to restore even the chromatic values correctly. The Appendix identifies two such deviations from normalization-compatibility:

- (3) *Resistance of narrow-band surfaces to chromaticity shift.* Stimuli from narrow-band surfaces are resistant to shifts in log cone excitation space, with illumination-invariance as a limit approached for monochromatic reflectances (these remain at the spectrum locus).
- (4) *Gamut compression with narrow-band illuminants.* For narrow-band illuminants, the spectrum locus is again approached as a limit, but this time for all surfaces, so the gamut is compressed.

Figure 1 illustrates these principles of the mapping of chromatic values in the Gaussian World. The “kite” represents a set of surface colours, centred on a non-selective white and with a dispersion roughly representative of natural scenes (see §2). For the centre kite, the surfaces are viewed under equal-energy illumination; at left and right, under an idealization of bluish and reddish daylights.

The kite is formed from lines of constant surface-reflectance spectral centroid (its radii) or of constant surface-reflectance bandwidth or spectral curvature (its crossbars).⁵ Under chromatic illumination, the kite is translated almost rigidly, illustrating the normalization-compatible mapping of (1) above. But it also undergoes a slight skewing because the more narrow-band reflectances (e.g. the yellows forming the bottom edge of the kite) are more

resistant to the illuminant-dependent chromaticity shift. Figure 1(b) clarifies the geometry of this skewing effect ((3) above): the kite radii pivot around the fixed points obtained by extrapolating them to the spectrum locus. In addition to skewing, the kite undergoes a less obvious illuminant-dependent compression ((4) above). Because the illuminants considered are all broad-band (or more precisely, have close to exponential spectra, giving them low spectral curvatures in the sense defined in the Appendix), the compression here is minimal.

Since discussions framed in the linear world have emphasized that normalization is in principle inadequate, the approximate validity of normalization in the Gaussian World is somewhat surprising. This raises the question: which ideal world is more pertinent to real vision in the real world? In §3 we briefly review evidence about the types of compensation actually implemented by the visual system. But first we consider the illumination-dependent mapping found in the real world, and compare it with what happens in the ideal worlds—with particular reference to the Gaussian one, since that model of the environment is particularly tightly constrained and provides a rich set of easily tested predictions.

2. The Real World

For an initial sample of the real world we have relied mainly on data of Ruderman, Cronin and Chiao (Ruderman et al. 1998). Ruderman, Cronin and Chiao obtained spectral reflectance estimates, pixel by 3 min arc pixel, for 12 entire views of natural environments, creating a data set comprising nearly 200,000 pixels. They characterized each surface by its spectral reflectance relative to a full reflectance white standard measured in the same scene. We derived the cone excitations L, M and S (for long-wavelength, midspectral and short-wavelength cones) from the measured spectral reflectances by integrating their crossproducts with the energy basis cone sensitivities of Stockman et al. (Stockman, MacLeod et al. 1993), assuming four different illuminants from the set of CIE daylight spectra (Wyszecki and Stiles, 1982, p. 145), with correlated colour temperatures 4000K, 5500K, 8500K, and 20000K, covering the extreme range of unimpeded daylight illuminants (Fig. 2). The measure of surface luminance is given by the summed excitations of L and M cones. The intensity-invariant chromatic measures we mainly

⁵ In these logarithmic coordinates, lines of constant bandwidth, including the spectrum locus, are entirely straight. But in a chromaticity diagram whose coordinates are linear with the cone excitations, these straight lines are bent into the more familiar curved form, with a straight region only where S cone excitation is negligible.

consider are the (r,b) chromaticity coordinates of MacLeod and Boynton (1979), defined by $r = L/(L+M)$ and $b = S/(L+M)$, with units for b chosen to make it unity for white. The choice of cone sensitivities or chromaticity measures is not critical. The particular logarithmic colour coordinates adopted by Ruderman et al. in their own analysis, for example, are practically linear with $\log(r)$ and $\log(b)$ in the case of natural broadband stimuli, although the equivalence is not exact.

We discuss in turn the different aspects of illumination-dependent mapping that were listed above.

Normalization-compatible mapping

As noted above, normalization-compatible mapping implies a rigid translation in $\log(L,M,S)$ space under changing illumination. For all individual pixels in the natural scenes of Ruderman et al. we plotted the appropriate logarithmic coordinate for 20000K illumination against its value under the 4000K illuminant. Fig. 3 shows results for one typical scene (their Park4). Despite large chromatic shifts the pixels cluster closely around lines of slope 1, as expected if the change of illumination from 4000K to 20000K makes the stimulus chromaticities all undergo the same shift in $\log(L,M,S)$ space. These results support those of Dannemiller (1993) based on Krinov's average terrain data, and of Foster and Nascimento (1994) based on Munsell papers, and extend them to individual elements of natural scenes.

A novel feature of our analysis is the separate treatment of the chromatic and luminance axes. While the validity of the rigid translation approximation appears better for the luminance axis (3rd panel) than for the chromatic ones, it should be noted that the axis for $\log(r)$ has a very different scale from the other two, due to the very small variance of this quantity in natural images (Ruderman et al., 1998) which is accompanied by a commensurate sensitivity of the visual system to this variable (MacLeod and von der Twer, this volume). Approximately rigid displacements measured on the coarse scale required for displaying the individual cone excitations, or luminance, do not imply (perceptually) rigid displacement along this highly sensitive chromatic axis. Nevertheless, the rigidity principle appears here to be a useful approximation for all three axes. Averaged over 12 scenes, the standard deviation of the changes in the (decimal) log of r (among the pixels of one scene) is only 0.0024; this means that the factor by which r changes is stable across pixels with a standard deviation of 0.55% of its mean.

For b and for luminance, the standard deviation is 1.8%; all three deviations from rigid translation are small enough to be undetectable or barely detectable, hence perceptually inconsequential. Although the deviations imply that a constancy compensation based on normalization could not be strictly perfect, they are so small that it could be practically perfect by comparison with the conspicuously imperfect constancy that characterizes human vision.

The closely proportional variation of the luminance values under the two illuminants (Figure 3(c)) warrants further comment. As mentioned in §1, strict proportionality is not characteristic of the Gaussian World. Nor is it in fact characteristic of normalization-compatible mapping in the sense we consider here, where each cone type has its own normalization factor. This requires proportional variation of the L and the M cone excitations under different illuminants, rather than proportional variation of luminance (which is well modelled by L+M). These are not equivalent (unless the proportionality constant is the same for L and M cones). But if L and M cone excitations are considered separately, the proportionality holds still more precisely, with a standard deviation of 0.84% for L and 0.89% for M, as opposed to the 1.8% value for luminance. An expected improvement of this sort formed the basis of the suggestion by von Campenhausen (1986) and Shepard (1992) that the requirements of lightness constancy have generated significant selective pressure for the evolution of trichromacy. But the already very close proportionality found for luminance undermines this argument: a system using the slightly broader L+M function as its sole spectral sensitivity could achieve very good constancy through normalization (with a root mean square error of 1.8% in the comparison of intensities between extreme daylight illuminants), so the improvement in constancy obtainable by dividing the red-green spectral range between the L and M cones is minimal.

The real world, then, is far from chaotic; and while it is not perfectly orderly like the three-band world, on this evidence it exhibits normalization-compatible mapping to a perceptually acceptable approximation. The Gaussian idealization appears to be a viable model in this respect, and the added intricacy of the linear model may not be necessary or even appropriate. We next ask whether the real world exhibits the deviations from normalization-compatible mapping that characterize the Gaussian World.

When the light gets red, the reds get lighter

Fig. 4 shows for the same scene the change in luminance of individual pixels in going from the bluish 20000K illuminant to the reddish 4000K illuminant, plotted versus the pixel's r

coordinate. The expected correlation is present: the reds (mostly) get lighter and the blues and greens dimmer. The magnitude of the effect is only a few percent, but as the Figure shows it is a major source of all departure from normalization-compatible mapping for the luminance axis (since the departures are themselves small). The effect is expected for the Gaussian World idealization but is also generated in the Linear one with appropriate basis functions.

Shift resistance and gamut compression

While it is inevitable on physical grounds that surfaces with sufficiently narrow-band reflectance will undergo a smaller chromaticity shift toward the illuminant, this fact cannot be exploited in the recovery of surface colour unless the observers can recognize the differences in surface bandwidth. Here we assume that the only information available for that purpose is the triple of cone excitations generated by the surface. This means that only those bandwidth effects that are correlated with surface colour are relevant. Two such surface bandwidth effects emerged in our simulations.

The first is illustrated for the Park4 scene in Fig. 5. In this and in the other scenes, surface-reflectance bandwidth (or more precisely, signed spectral curvature, as defined in the Appendix) tends to vary markedly along the b axis, with narrower bandwidth for the abundant yellowish colours than for whites, or for the relatively infrequent purplish colours, which tend to have concave-upward spectra (negative spectral curvature). Accordingly, Fig. 5 shows a result otherwise surprising: the illumination-induced shift in r is greater for surfaces with high b . While not large, this deviation from rigid logarithmic translation is enough, in many scenes, to skew the distribution of pixels in the $(\log(r), \log(b))$ plane in just the manner illustrated for the Gaussian World in Figure 1.

Second, the simulations revealed a weak tendency for the lighter surfaces to be more stable in physical chromaticity under change of illuminant. This was largely mediated by the correlation between luminance and b , the lighter surfaces being more yellowish in these scenes.

Gamut compression by narrow-band illumination was not clearly evident in our simulations. Variances among the scene elements in log cone excitation space did not differ systematically for the three illuminants. But the chosen daylight illuminants are not well suited to reveal such an effect. As Fig. 2 shows, they are roughly exponential in the visible range, and in the algebraic analysis of the Gaussian World (Appendix) they would be assigned spectral curvatures close to zero, and therefore would not be expected to restrict the range of $\log(r)$ or

$\log(b)$ much. Gamut compression is physically inevitable when the illumination is of sufficiently restricted bandwidth, but in the natural world it may become pronounced (if at all) only under restricted or indirect illumination (for instance in forest scenes).

3. Compensation by the Visual System: which world do we internalize?

Having established that the illumination-dependent mappings found both in a plausible ideal world (the Gaussian World) and in the real one are approximately normalization-compatible, we now ask whether the visual system has adapted to this environmental regularity by adopting something like a normalization algorithm for its recovery of surface colour. Evidence to be reviewed supports this, but leaves it uncertain whether the visual process is able to compensate for the relatively small natural deviations from normalization-compatible mapping.

Normalization for Colour

Brainard and Wandell (1992) made memory matches between computer-simulated surfaces viewed under two different (computer-simulated) conditions of illumination, and compared the errors in prediction of these matches for different candidate models of human colour constancy. The normalization model (with only 3 free parameters, one for scaling each cone sensitivity) predicted the human matches very nearly as well as Linear World models in which the human observer is supposed to internalize a linear description of the spectral distributions, using 9 or more parameters per illuminant, and do a matrix inversion to recover the surface colours. Under very different conditions involving binocular matching, Chichilnisky and Wandell (1995) obtained similar results. And under more natural conditions, Brainard, Brunt and Speigle (1997) found the same.

In all these studies, a simple cone-sensitivity scaling or normalization model mimics human visual judgments almost about as well as the Linear World based models can do with their extra parameters. Nor is there any suggestion that specific features of the linear model's predictions are reflected in the human matches.

Further evidence that our visual systems are designed for a normalization-compatible world comes from an interesting experiment by Nascimento and Foster (1997). They made careful simulations of illuminant changes applied simultaneously to several surfaces in a CRT display, and asked observers to compare these with precisely normalization-compatible

transformations. Although the latter are not as accurately representative of the physical effects of illumination change, they were *more* likely to be identified as illumination changes by the observers.

When the Lighting gets Red, the Reds get Lighter.

Lightness variations like this are easily demonstrated or observed in common experience with artificial light—when the light gets red, the reds do get lighter perceptually—but they seem to be taken for granted and have not been discussed in the context of colour constancy or even experimentally quantified. They represent a partial or complete failure of constancy. It is a failure that could in principle be eliminated relatively easily by a sophisticated visual system that has internalized the relevant environmental regularities. But its elimination would require interaction between the intensive and the purely chromatic components of the neural representation, and this might be problematic for a primitive compensation mechanism. Whatever the reason, this environmental deviation from normalization-compatible mapping may not, on the present limited evidence, have been internalized in the functional organization of our visual system.

Resistance of Narrow-band Surfaces to Chromaticity Shift

We have noted that to compensate for this deviation from normalization-compatible mapping, the visual system must exploit correlations between surface colour and surface bandwidth, which cause certain colours, in general, to be more resistant to chromaticity shifts. Two such effects were noted in §2.

The reduced illumination-induced (redward) chromaticity shift for yellows as compared with whites, in the metric of Fig. 5, is a violation of normalization-compatible mapping, that calls for the visual system to make a greater compensatory correction for whites (or desaturated purples) than for yellows. The question whether such differential compensation actually occurs has not been specifically addressed experimentally, and it is not clear, in the data of Ware and Cowan (1982) for example, whether the visual system merely adopts the simple but environmentally suboptimal normalization algorithm. Chichilnisky and Wandell (1995) and Whittle (this volume) note that equal logarithmic differences relative to differently coloured backgrounds in Whittle's haploscopic display are generally perceptually equal, as if only normalization were operative. But deviations from this principle appear with saturated

backgrounds. These differences are in the expected direction, viz. that equal logarithmic differences are perceptually larger in the case of the saturated background, but it is not yet clear whether they quantitatively or comprehensively support the idea that the visual system may have successfully internalized this environmental deviation from normalization-compatible mapping.

The tendency for brighter surfaces to change less in chromaticity, mentioned in Section 2, provides a possible (eco-)logical basis for the Helson-Judd effect (Helson, 1938), in which the lighter surfaces in an artificial physically achromatic scene (made up of spectrally non-selective surfaces) perceptually assume a faint tint of the colour of the illuminant, and darker surfaces assume the complementary colour. In Helson's situation the change in stimulus chromaticity was independent of luminance. If the visual system has a basis in environmental statistics for expecting the more luminous surfaces to change less than the darker ones, then the shifts toward the illuminant chromaticity for the more luminous surfaces in Helson's scenes would be unexpectedly large, and the resulting stimulus chromaticities might logically be taken as evidence that the lighter surfaces have an inherent colouration similar to the illuminant.

Gamut Compression with Narrow Band Illuminants

When illuminant bandwidth becomes narrow, the configuration of surfaces in log cone excitation space is compressed toward to the locus of physically non-selective reflectances (as discussed quantitatively in the Appendix). Cone signal normalization can effect compensatory translations in this space but cannot counteract the compression. Yet there is evidence that we have quite powerful perceptual compensatory mechanisms available to deal with such gamut compression. Brown and MacLeod (1992, 1997) found that when a test field is surrounded by elements of uniform or nearly uniform colour (as would occur with a narrow-band illuminant) the colour of the test field is perceptually enhanced when comparison is made with a test field in a chromatically varied surround. This perceptual gamut expansion, cued by the physical heterogeneity of other stimuli within the scene, would be useful in compensating for physical gamut compression when the illuminant bandwidth becomes narrow. This is one aspect of colour constancy that can owe nothing to normalization. It could, however, originate from something as primitive as sensitivity modifications at post-receptoral (spatially opponent and/or colour-opponent) stages of visual processing. Perceptual gamut expansion may not, however, have evolved for dealing with natural illuminants of restricted bandwidth, since as we have noted, bandwidth restrictions with unimpeded daylight illumination are not pronounced.

Perceptual gamut expansion is probably more valuable under hazy conditions (the situation investigated for the achromatic domain by Gilchrist and Jacobsen, 1983) than in the rare case where the illuminant bandwidth becomes narrow.

Our conclusion with regard to the compensation problem is that the visual system employs a colour-correcting mapping approximately equivalent to normalization, but may not compensate appropriately for the relatively small natural deviations from normalization-compatible mapping that characterize both the real world and its most useful idealizations. We now turn from the compensation problem to the complementary problem of colour constancy: the problem of estimating the illuminant.

4. The estimation problem: Inferring the Illuminant from Scene Statistics

Here we consider one source of information for estimating the illuminant: the illumination-dependent mapping itself. Illuminants may be recognizable by their effects on the scene statistics, or specifically on the distribution of scene elements in cone excitation space. This leaves aside many other potential cues to the illuminant, particularly ones that depend on the three-dimensional geometry of the scene (Hurlbert, 1998). Effectively we consider a world of co-planar diffusely reflecting surfaces, with illumination incident uniformly over each scene. The most obvious candidate scene statistics are the scene-averages or maximum values of relevant quantities, but here we will conclude that environmental deviations from normalization-compatible mapping allow other scene statistics to play an important independent role.

The limited value of the scene-averaged chromaticity for estimating the illuminant

When the illumination of a scene is changed (for example towards more energy at long wavelengths), all reflected lights will change correspondingly, and the chromaticity averaged over the entire image will become more reddish. A simple approach to colour constancy could therefore take the space average chromaticity as a cue for the chromaticity of the illuminant and use this to correct for shifts in chromaticity of objects due to non-neutral illuminants.

One model for colour constancy that uses a spatial average of the receptor responses to estimate the illumination of the image was proposed by Buchsbaum (1980). His estimation of the illuminant is based on the assumption that for all scenes the field average reflectance is equal to an internal reflectance standard $S_0(\lambda)$.⁶ The illuminant is estimated by determining which illuminant would have resulted in the actual obtained average receptor response, assuming that the scene is illuminated uniformly and the spatial average reflectance of the scene is $S_0(\lambda)$. This estimate is then used together with the responses for the subfields in the image in order to obtain illuminant independent reflectance descriptors for each subfield.

For scenes for which the mean reflectance function differs from $S_0(\lambda)$, the algorithm wrongly attributes the mean receptor response to a spectrally biased illuminant. Every change in the mean receptor response is interpreted as due to a change in the illumination. Thus this algorithm, like many others relying on what has been called the “Gray World assumption”, is not able to deal with scenes in which coloured surfaces different in a particular direction from the known standard are predominant. This weakness is inherent to all models that use a space averaged chromaticity of the scene to estimate the illuminant. The reason is that this measure is ambiguous: for a given scene the mean chromaticity could be reddish due to a predominance of reddish surfaces within this scene, or to a reddish illuminant (Fig. 6(a)).

One approach that takes into account more information about the scene than mean chromaticity is the framework of probabilistic colour constancy. D’Zmura et al. (1995) presented a stochastic linear model for estimating the illuminant of a scene. This scheme uses a priori knowledge about reflectance and illuminant probabilities in order to calculate how likely it is that the viewed scene is illuminated by a particular light given the chromaticities of the reflected lights. The most likely illuminant is then determined by a maximum likelihood estimation procedure.

Let $p(S(\lambda))$ describe the a priori probability that the surface reflectance function $S(\lambda)$ is encountered in the world. If the probability distribution $p(S(\lambda))$ is known, one can determine the probability distribution of the receptor responses $r = (r_1, r_2, r_3)$ for a trichromatic visual system. Let $I(\lambda)$ be the scene illuminant and $C_i(\lambda)$ the sensitivity function of the i -th receptor, then it follows that

⁶ Even though this internal reflectance standard does not have to be spectrally neutral, this approach is

$$r_i = \int C_i(\lambda) I(\lambda) S(\lambda) d\lambda \quad (i = 1, 2, 3).$$

The conditional probability distribution $p((r_1, r_2, r_3) | I(\lambda))$ for the receptor responses given a particular illuminant $I(\lambda)$ can thus be calculated.

Now consider a scene containing a set of N_s surfaces drawn independently and viewed under an unknown illuminant resulting in a set of N_s cone responses. For a particular illuminant $I(\lambda)$ the likelihood $L(I(\lambda))$ for this set of cone responses can be calculated as

$$L(I(\lambda)) = \prod_{n=1}^{N_s} p((r_1, r_2, r_3)_n | I(\lambda)).$$

A simple way to estimate the illuminant of an image is to take the illuminant that has the highest likelihood for the given set of responses. The result of this maximum likelihood estimation is also called maximum a posteriori (MAP) estimate.

If one also has prior knowledge about the probability of encountering various illuminants in the world, one can refine the estimation by finding the maximum for the likelihood that also takes this probability distribution into account:

$$L(I(\lambda)) = \prod_{n=1}^{N_s} p((r_1, r_2, r_3)_n | I(\lambda)) p(I(\lambda))$$

This estimate of the illuminant then can be used to recover the surface reflectance functions of the scene.

D'Zmura et al. (1995) implemented a maximum likelihood estimation scheme that uses only two chromaticity values for each reflected light and employs a linear model for the surface reflectances and illuminations. They present a Monte Carlo simulation which shows that the chromaticities of relatively few reflected lights are sufficient to recover an illuminant accurately.

Brainard & Freeman (1997) presented a similar Bayesian scheme to determine the illuminant, although they use a different optimality criterion for finding the best estimate. They argue that this so-called maximum local mass (MLM) estimator is more appropriate for perceptual tasks than other estimators. They compare the simulated performance of this scheme

with a Bayesian scheme using a MAP estimator, a Gray World algorithm and other colour constancy algorithms. The MLM method performs better than all other algorithms. But if the mean chromaticity of the sample of surfaces in a scene is biased, all algorithms perform poorly. Thus, these methods have only a limited capability to separate illuminant changes from changes in the surface collection under the simple viewing conditions considered here.

By using the probability of the observed chromaticities, both Bayesian schemes take into account information about the entire distribution of the reflected lights, not just the mean. If other statistics change in a characteristic way under changing illumination and the employed *a priori* knowledge about probabilities mirrors these regularities of the world, these schemes automatically gain improved performance.

Nevertheless, in adopting the assumption that all scenes draw their surfaces independently from a single $p(S(\lambda))$ distribution that is characteristic of the world in general, these schemes still embody the Gray World assumption (albeit in a statistical reformulation), and inherit its problems. In the real world, where different scenes have surfaces drawn from different populations, the expected values of all statistics differ from one type of scene to another, and these schemes deliver correspondingly incorrect estimates of the illuminant.

Consideration of scene statistics does, however, allow a more radical departure from the Gray World assumption. The assumption of a single $p(S(\lambda))$ can be abandoned, and replaced by other, more general constraints on the distribution of cone excitations. An interesting example is the proposal of Forsyth (1990), which considers the gamut of cone excitations that are physically realizable under a particular illuminant. This gamut is fixed, for any given illuminant, by the constraint that whatever collection of surfaces is present, none can reflect more than 100% of the incident light at any wavelength (Koenderink and van Doorn, this volume). Forsyth (1990) modelled the colour of an object as the receptor responses of the object under a fixed canonical light. In order to achieve colour constancy one must then estimate which illuminant is present and try to estimate what the receptor responses for the objects would have been if the scene had been illuminated by the canonical light. Every illuminant is associated with a linear mapping that maps the responses under the canonical light to a different set of responses under this particular illuminant. Once one has estimated the illuminant in a scene, one can then apply the inverse mapping to determine the responses under the canonical light and thus obtain the constant colour descriptors. Forsyth (1990) described the circumstances under

which these mappings are invertible and the conditions and assumptions necessary for finding the illuminant of an image (and therefore the mapping).

The estimation of the illuminant is based mainly on the fact that a surface cannot reflect more light than is cast on it at any wavelength. For a particular illuminant, therefore, many receptor responses cannot be achieved; the set of possible receptor responses is bounded. The illuminant is constrained by the observed image: some illuminants are not compatible with the set of responses given in the image. If, for example, a patch strongly excites the long wavelength receptor, it cannot be illuminated by a blue light. Using the gamut of the receptor responses of an image one can therefore exclude some illuminants leaving only the set of possible illuminants (the 'feasible set'). Other visual information may then provide cues that reduce the feasible set further or an estimator is needed to choose the most likely illuminant from this set.

Forsyth (1990) specified an algorithm, Crule, using the convex hull of the receptor responses of an image in order to find the feasible set. He compared the performance of this algorithm with that of the Retinex algorithm of Land et al. (1971). Unlike the Retinex algorithm, Crule is only slightly disturbed if the mean chromaticity of an image is biased by a predominance of one colour in the scene. Use of the hull of the receptor responses instead of an average allows the algorithm to better separate changes in illumination from changes in surface ensemble.

The power of the Crule algorithm derives from the fact that the gamut is both characteristic of the illuminant and invariant (as a limit) with changing scene content. Given uncertainty about scene composition, the gamut constraint can be helpful for illuminant estimation even in an environment where illumination-dependent mapping is normalization-compatible. But deviations from normalization-compatible mapping provide further cues for illuminant estimation. One example is the case of almost-monochromatic illuminants, which each produce a highly characteristic distribution in cone excitation space, whatever surfaces may populate the scene. Unfortunately, the gamut is seldom approached by natural surfaces. But the real-world constraints that determine the gamut may also influence appropriate statistics of the distribution of cone excitations, making these statistics usefully diagnostic of the illuminant. We next discuss the usefulness of various scene statistics for this purpose.

Higher-order chromatic scene statistics as cues for the illuminant: their distribution in natural scenes

Correlations. As mentioned above, the visual system has to deal with an ambiguity if it uses only the mean chromaticity of an image to estimate the chromaticity of its illumination. This statistic alone does not allow us to distinguish a scene under a chromatic illumination from a scene with a chromatically biased surface ensemble (Fig. 6(a)). But since a red illumination makes red surfaces (relatively) lighter as shown in the analysis of the natural scenes (§2) and as predicted by the Gaussian World (§1), the correlation between redness and luminance may be diagnostic for the illumination. A high luminance-redness correlation among the elements of a scene might thus by itself suggest a reddish illuminant, no matter what the scene-averaged chromaticity (Fig. 6(b)), and could thus provide an estimate of illuminant colour balance even in a world where different scenes differ in predominant colour and hence violate the Gray World assumption.

To express the point more generally: by evaluating both mean and correlation, an observer can estimate two unknowns - the predominant colour inherent in the objects making up the scene, and the chromaticity of the light source that illuminates the scene. In this way higher-order statistics of the distribution of surface luminance and chromaticity within a scene can resolve the ambiguity encountered in considering scene-averaged chromaticity alone.

We performed a theoretical analysis of higher-order chromatic scene statistics in natural scenes in order to investigate whether these measures can usefully support inferences about the illuminant.

Fig. 7 shows the correlation between redness and luminance within a scene versus the mean scene chromaticity for the 12 scenes under each illuminant. The mean redness of all scenes is of course highest under 4000K illumination. Contrary to our initial expectation, the correlations are almost independent of illumination.

But note that in Fig. 7 for each illuminant the distribution of the two statistics across scenes is negatively sloped as indicated by the regression lines. Thus, despite this invariance with illuminant, the correlation measure can resolve an ambiguity with which a visual system would have to deal, if it took only the mean chromaticity of a scene as a cue for the illuminant. This is possible because scene redness and illuminant redness affect the correlation differently:

For a reddish scene under neutral lighting, the visual system's diminishing sensitivity for long wavelengths introduces a negative correlation between pixel redness and pixel luminance.

(This effect is reduced or absent for predominantly greenish scenes, since the redder pixels within such scenes are likely to have spectral distributions better placed in relation to the luminosity function. This accounts for the sloped regression lines in Fig. 7.)

For a neutral scene under reddish light, the low luminosity of reds is counteracted by the illuminant's greater energy at long wavelengths, making the correlation between pixel redness and pixel luminance greater in such a scene than for the reddish scene under neutral light.

Consider as an example a scene for which the mean $\log_{10}(r)$ is -0.155 . If only this mean chromaticity is known, this could be a reddish scene under neutral light or a neutral scene under reddish light. But if the visual system takes into account the correlation between redness and luminance, it can distinguish between illumination redness and scene redness. If the correlation is low, then a reddish scene under a neutral illuminant is more likely. Thus the use of the correlation measure can improve the estimation of the illuminant, even though this statistic is almost unaffected by changes in illumination.

The correlation between luminance and blueness within a scene was almost independent of illumination for the used set of scenes. Unlike the correlation between luminance and redness, the correlation between luminance and blueness varied too widely in the set of scenes we used, to play a useful disambiguating role.

Variances. We also investigated the variance of the chromaticities within scenes. As Fig. 8 shows, the standard deviation of $\log(r)$ increases for only some scenes when the colour temperature of the illumination changes from 20000K to 4000K, and there is no indication that this statistic could play the disambiguating role suggested for the luminance-redness correlation in Figure 7. The lack of a consistent effect of illuminant on chromatic variance within scenes is not surprising since, as noted above, the illuminants we used did not vary much in bandwidth. The variance statistic could be diagnostic of a chromatically biased illuminant in more extreme cases than the ones considered here.

The standard deviation of $\log(b)$ as well as the standard deviation of $\log(\text{luminance})$ were likewise almost independent of illumination, and were similarly distributed for scenes of different mean chromaticity. This means they cannot be very useful for diagnosing the illuminant.

Skewness. For the reddest surfaces in any scene, the redward shift in stimulus chromaticity under a reddish illuminant is ultimately limited by the spectrum locus. One might therefore suppose that a shift to reddish illumination would shift the already reddish pixels less toward red

than greenish pixels. The distribution of chromaticities would in this way get skewed away from the chromaticity of the illuminant.

The results of the simulation actually revealed an opposite trend. For each of the 12 natural scenes the shift in $\log(r)$ in going from 20000K illumination to 4000K illumination was higher for pixels with high redness (as measured by $\log(r)$ values under 7000K illumination) - if we consider separately sets of pixels of similar blueness (similar values of the chromaticity coordinate b). This may be due in part to the abundance in the natural scenes of greenish vegetation, which tends to have a relatively narrow reflectance band, and consequently to be resistant to chromaticity shift. In any case, this relation does not emerge clearly if one considers the whole set of pixels for a scene, since there is also a dependence of the shift in redness on the blueness of the pixels (§2) and this latter effect counterbalances the former if one pools all slices of constant blueness values. Thus the calculated Fisher skewness in $\log(r)$ or in $\log(b)$ within the scenes is almost independent of the illuminant.

The use of higher-order scene statistics by the visual system

One experiment that investigated the influence of higher-order scene statistics on colour constancy was performed by Mausfeld & Andres (Mausfeld, 1998). They created a new type of stimulus that makes it possible to vary certain statistics of the chromaticity distribution of the image independently. These computer-generated displays consist of a random structure of overlapping circles around a centre test spot.⁷ By varying the modulation of the colour of the circles along the luminance axis and along the red-green axis Mausfeld & Andres obtained a set of such “Seurat” stimuli, each with a different combination of variances for chromaticity and luminance. The spatial mean chromaticity and luminance of the surround were equal for all stimuli.

Subjects made red - green equilibrium settings ('unique yellow' settings) by a double random staircase procedure. This subjective yellow point toward the (mean) surround chromaticity, as expected, but the shift was greater for surrounds of reduced chromatic variance than for very heterogeneous surrounds.

Mausfeld (1998) discusses these results within an ethologically inspired perspective: certain stimulus characteristics trigger the visual system to interpret an image in terms of certain elementary perceptual categories. He argues that a reduced chromatic variance increases the

tendency of the visual system to interpret a scene with a biased average chromaticity as chromatically illuminated.⁸ This results in a stronger correction for the appearance of all patches including the test spot. Thus the effect of a chromatically non-neutral surround on unique yellow settings is larger if the surround has low chromatic variance than if it has high chromatic variance. That is, these surrounds are not 'functionally equivalent' even though the space-averaged chromaticity of both surrounds is equal. This contradicts algorithms for colour constancy in which only the space-average of the surround elements is important (for example, those that rely on a Gray World assumption).

Another experiment that tackled the question whether surrounds with same space-averaged chromaticity always causes the same changes in appearance of a test region was reported by Jenness & Shevell (1995). They used red backgrounds with randomly scattered sparse white or green dots and compared them to the red background without these dots and to uniform backgrounds with the same space-averaged chromaticity as the inhomogeneous backgrounds. Subjects had to set the test field so that it appeared neither reddish nor greenish. The influence of the inhomogeneous background on these unique yellow settings differed from that of the uniform background with the same space average.

To find out whether human vision employs higher-order scene statistics to estimate the illuminant, we performed experiments using stimuli similar to the one used by Mausfeld & Andres (Mausfeld, 1998), but varying other statistics. We asked subjects to adjust the colour of a circular test field embedded in such a computer-generated display so that it appeared neutral gray.

In our main experiment we varied the correlation between $\log(r)$ and $\log(\text{luminance})$ independently of other statistics (means and variances). For a given condition, the chromaticity and luminance values for the circles in the surround were chosen to achieve a certain correlation value (-1, -0.8, 0.0, 0.8 or 1). If the perceived colour of the centre test spot was not influenced by the varied correlation, then the settings to make the test spot neutral gray should be the same for all 5 conditions. This is because the space-averaged chromaticities of the backgrounds were the

⁷ If the diameter of the circles is chosen to be very small, these images resemble Neo-Impressionistic paintings. Therefore Mausfeld (1998) refers to these stimuli as Seurat-type configurations.

⁸ This assumption is based on the fact that a chromatically biased illuminant that is narrow in bandwidth will lead to a reduced variance of chromaticity values for the reflected lights of a scene. For a discussion whether the chromatic variance of natural scenes changes under a range of typical daylights see section 4.2.

same. The backgrounds would then be functionally equivalent with respect to the perceived colour of the centre test spot.

Results for subject JG are shown in Fig. 9. For conditions with higher correlation between redness and luminance, a more reddish chromaticity was required to make the test field subjectively achromatic. The data for eight of ten subjects tested were quantitatively similar and individually statistically significant. When the correlation between redness and luminance was positive, all subjects selected a physically more reddish (higher r) test field as neutral gray. Since higher r -values are associated with redder illumination of a physical neutral surface, this is the result expected if the observer infers a more reddish illumination in the case of positive luminance-redness correlation, and perceives neutral gray when a correspondingly reddish light stimulus is received. These results are thus consistent with the possible use of the luminance-redness correlation as a cue for the chromaticity of the illuminant, as suggested by the analysis of natural scenes earlier in this section.

How much weight should a smart visual system give to the correlation between redness and luminance in estimating the illumination? To answer this question for our simulated world, we calculated a maximum likelihood estimate of the chromaticity of the illumination based on the mean and correlation values of the scenes. Fig. 10 shows the effect of the correlation between redness and luminance on the test spot settings for all subjects. The dashed lines through the data are parameter-free theoretical predictions, on the hypothesis that optimal weight is given to the correlation measure in estimating the illuminant. Two cases are illustrated. The steeper line is obtained for an optimal visual system that uses (in addition to the luminance-redness correlation) the luminance-weighted mean chromaticity (or the mean cone excitations) of the surround; the more shallow sloped line is for a system that applies no luminance weighting in evaluating the mean chromaticity. The small effect observed in our experiment is thus roughly consistent with optimal computation.

Our experiments revealed no comparable effect for the correlation between blueness and luminance, or for the skewness of the blueness and redness distributions. This is also consistent with the theoretical analysis of natural scenes, in which the latter statistics did not appear helpful for estimating the illuminant.

To summarize our discussion of the estimation problem: the effects of changing illumination on natural scenes indicate that the statistics of the cone excitations associated with individual scene elements are potentially helpful in resolving the ambiguity inherent in use of

scene-averages alone. And recent evidence from perceptual experiments suggests that certain of these statistics are indeed exploited, possibly in a statistically nearly optimal manner.

5. Concluding Summary

In §1, we introduce an idealization of the world of colour in which the interplay of illuminants, surfaces and photoreceptors becomes mathematically tractable. Normalization is a fairly effective algorithm for colour correction in this Gaussian World, and also in the real world as viewed by human retinas (§2). The recent evidence reviewed in §3 indicates that normalization may be a fairly good model for the human visual system's compensation for changing illumination. In §4 we turn to the problem of illuminant estimation. There we show, theoretically and experimentally, how considering higher-order statistics of the cone excitation distribution make it possible to dissociate scene-average colouration from illuminant colouration.

ACKNOWLEDGEMENTS

We thank D. L. Ruderman, T. W. Cronin, and C. C. Chiao for supplying us with their spectral data of natural scenes. This work was supported by NIH grant EY01711. J. Golz was supported by the German-American Fulbright Commission.

References

- Brainard, D. H., Brunt, W. A., & Speigle, J. M. (1997). Color constancy in the nearly natural image. I. Asymmetric matches. *Journal of the Optical Society of America a. Optics and Image Science*, 14(9), 2091-110.
- Brainard, D. H., & Freeman, W. T. (1997). Bayesian color constancy. *Journal of the Optical Society of America a. Optics and Image Science*, 14(7), 1393-411.
- Brainard, D. H., & Wandell, B. A. (1992). Asymmetric color matching: how color appearance depends on the illuminant. *Journal of the Optical Society of America a. Optics and Imagescience*, 9(9), 1433-48.
- Brill, M. H. (1978). A device performing illuminant-invariant assessment of chromatic relations. *Journal of Theoretical Biology*, 71(3), 473-8.
- Brown, R. O., & MacLeod, D. I. A. (1992). Saturation and color constancy. *Advances in Color Vision Technical Digest (Optical Society of America)*, 4, 110-111.
- Brown, R. O., & MacLeod, D. I. A. (1997). Color appearance depends on the variance of surround colors. *Current Biology*, 7(11), 844-9.
- Buchsbaum, G. (1980). A Spatial Processor Model for Object Colour Perception. *Journal of The Franklin Institute*, 310(1), 1-26.
- Campenhauen, C. v. (1986). Photoreceptors, lightness constancy and color vision. *Naturwissenschaften*, 73(11), 674-5.
- Chichilnisky, E. J., & Wandell, B. A. (1995). Photoreceptor sensitivity changes explain color appearance shifts induced by large uniform backgrounds in dichoptic matching [see comments]. *Vision Research*, 35(2), 239-54.
- Dannemiller, J. L. (1993). Rank orderings of photoreceptor photon catches from natural objects are nearly illuminant-invariant. *Vision Research*, 33(1), 131-40.
- D'Zmura, M., Iverson, G., Singer, B. (1995). Probabilistic Color Constancy. In D. Luce, D'Zmura, M., Hoffman, D., Iverson, G., Romney, A. (Ed.), *Geometric Representations of Perceptual Phenomena*. Mahwah: Lawrence Erlbaum Associates.
- Forsyth, D. A. (1990). Colour Constancy. In A. a. T. Blake, T. (Ed.), *AI and the Eye*. Chichester: John Wiley & Sons.
- Foster, D. H., & Nascimento, S. M. (1994). Relational colour constancy from invariant cone-excitation ratios. *Proceedings of the Royal Society of London. Series B: Biological Sciences*, 257(1349), 115-21.
- Gilchrist, A. L., & Jacobsen, A. (1983). Lightness constancy through a veiling luminance. *J. Exp. Psychol.*, 9, 936-944.
- Helmholtz, H. v. (1896). *Handbuch der Physiologischen Optik*. (2nd ed.). Hamburg: Voss.
- Helson, H. (1938). Fundamental problems in color vision. I. The principle governing changes in hue, saturation, and lightness of nonselective samples in chromatic illumination. *J. Exp. Psychol.*, 23, 439-476.
- Hurlbert, A. C. (1998). Computational models of color constancy. In V. Walsh & J. Kulikowski (Eds.), *Perceptual constancy: Why things look as they do* (pp. 283-322). Cambridge: Cambridge University Press.
- Jenness, J. W., & Shevell, S. K. (1995). Color appearance with sparse chromatic context. *Vision Research*, 35(6), 797-805.

- Land, E. H., & McCann, J. J. (1971). Lightness and retinex theory. *Journal of the Optical Society of America*, 61(1), 1-11.
- MacLeod, D. I. A. (1985). Receptor constraints on color appearance. In O. D. & Z. S. (Eds.), *Central and Peripheral Mechanisms of Color Vision*. London: MacMillan.
- MacLeod, D. I. A., & Boynton, R. M. (1979). Chromaticity diagram showing cone excitation by stimuli of equal luminance. *Journal of the Optical Society of America*, 69(8), 1183-6.
- MacLeod, D. I. A., Williams, D. R., & Makous, W. (1992). A visual nonlinearity fed by single cones. *Vision Research*, 32(2), 347-363.
- Marr, D. (1982). *Vision : a computational investigation into the human representation and processing of visual information*. San Francisco: W.H. Freeman.
- Mausfeld, R. (1998). Color Perception: From Grassmann Codes to a Dual Code for Object and Illumination Colors. In W. Backhaus, R. Kliegel, & J. Werner (Eds.), *Color Vision* (pp. 219 - 250). Berlin/New York: De Gruyter.
- Moon, P., & Spencer, D. E. (1945). Polynomial representation of spectral curves. *Journal of the Optical Society of America*, 35, 597 - 600.
- Nascimento, S. M., & Foster, D. H. (1997). Detecting natural changes of cone-excitation ratios in simple and complex coloured images. *Proceedings of the Royal Society of London. Series B: Biological Sciences*, 264(1386), 1395-402.
- Ruderman, D. L., Cronin, T. W., Chiao, C. C. (1998). Statistics of cone responses to natural images: implications for visual coding. *J. Opt. Soc. Am. A*, 15, 2036-2045.
- Sällström, P. (1973). *Colour and physics: Some remarks concerning the physical aspects of human colour vision* (Report No. 73-09). Stockholm: Institute of Physics, University of Stockholm.
- Shepard, R. N. (1992). The Perceptual Organization of Colors: An Adaptation to Regularities of the Terrestrial World? In J. H. Barkow, L. Cosmides, & J. Tooby (Eds.), *The Adapted Mind*. New York: Oxford University Press.
- Stockman, A., MacLeod, D. I. A., & Johnson, N. E. (1993). Spectral sensitivities of the human cones. *Journal of the Optical Society of America a. Optics and Imagescience*, 10(12), 2491-521.
- Ware, C., & Cowan, W. B. (1982). Changes in perceived color due to chromatic interactions. *Vision Research*, 22(11), 1353-62.
- Worthey, J. A., & Brill, M. H. (1986). Heuristic analysis of von Kries color constancy. *Journal of the Optical Society of America a. Optics and Image Science*, 3(10), 1708-12.
- Wyszecki, G., & Stiles, W. S. (1982). *Color science: concepts and methods, quantitative data and formulae*. (2nd ed.). New York: John Wiley & Sons.

LEGENDS

Figure 1: (a) Illumination-dependent chromatic values in the Gaussian World. The kite represents stimuli generated by a set of natural surfaces centred on white and representative of natural scenes; it is formed from lines of constant spectral centroid (the radii) or of constant spectral curvature (the crossbars). With illuminants roughly approximating 4000K (left), equal-energy (centre) or 20000K (right), the kite undergoes a translation, coupled with slight compression and skewing. Units for the cone excitations are chosen to place the equal-energy white at the origin.

(b) The same plot, with the addition of the spectrum locus as the illumination-invariant limit approached for monochromatic surface reflectances. This illustrates how the chromatic shifts for colours typical of natural scenes are theoretically related to the spectrum locus (although the model does not provide a usefully accurate construction of the spectrum locus itself).

Figure 2: Spectra of the three CIE daylights applied to the natural images.

Figure 3: Chromaticity and luminance values for individual pixels of the PARK4 scene of Ruderman et al. (1998), under 4000K illumination (horizontal axis) and under 20000K illumination (vertical axis). Normalization-compatible mapping requires proportionality between the values, or clustering of the points around lines of slope 1 in these double-logarithmic plots. Note the very different scales for the different plots.

Figure 4: Change in $\log_{10}(\text{luminance})$ for individual pixels of PARK4 when the 4000K illuminant replaces the 20000K illuminant, plotted versus redness of the pixel (as measured by $\log_{10}(r)$ under 7000K).

Figure 5: Change in redness ($\log_{10}(r)$) for individual pixels of PARK4 when the 4000K illuminant replaces the 20000K illuminant, plotted versus blueness of the pixel (as measured by $\log_{10}(b)$ under 7000K). The yellowish surfaces, at left, undergo smaller shifts.

Figure 6: (a) The mean chromaticity of a scene is an ambiguous measure for estimating the illuminant. (b) How this ambiguity could be resolved by high order scene statistics.

Figure 7: Luminance-redness correlation and mean redness of natural scenes under different illuminations. The vertical coordinate of each data point represents the correlation between pixel redness ($\log(r)$) and pixel luminance ($\log(\text{luminance})$) within an image of a scene under a particular illuminant. This correlation is plotted here against space average image redness (mean of $\log(r)$). The four clusters show data for 12 natural scenes under four different illuminants (\blacksquare : colour temperature 4,000K, \circ : 5,500K, \blacktriangledown : 8,500K, \blacklozenge : 20,000K). For a given scene, the luminance-redness correlations are almost independent of illumination, but within each illuminant cluster they are more negative for the redder scenes, as shown by the negatively sloped regression lines. Thus correlation and mean together separate the distributions of images resulting from different illuminants and make it possible to distinguish scene redness from illuminant redness.

Figure 8: Standard deviation of pixel redness within a scene, plotted versus scene-average redness for 12 natural scenes under four different illuminants (\blacksquare : colour temperature 4,000K, \circ : 5,500K, \blacktriangledown : 8,500K, \blacklozenge : 20,000K).

Figure 9: Dependence of centre test spot settings on correlation between redness and luminance in background for subject JG. Closed circles are the means of $\log(r)$ for perceptually achromatic test fields, error bars are \pm one standard error of the mean. If the settings were not dependent on the correlation but only on the mean of the chromaticity of the background, the settings should be the same for all conditions (horizontal dashed line). For the case that the settings are not dependent on the correlation but only on the luminance-weighted mean of the chromaticity of the background the results are predicted by the oblique dashed line. The measured settings are significantly different from both models ($p < 0.001$).

Figure 10: Experimental results for all subjects (circles) compared with predictions for an estimation of the illuminant making optimal use of the correlation statistic (dashed lines). The steeper sloped line is for the model that the visual system uses the luminance-weighted mean

chromaticity of the background, the more shallow sloped line for the not luminance-weighted mean. Error bars for the experimental results are \pm one standard error for subject variability.

APPENDIX: CONE EXCITATIONS IN THE GAUSSIAN WORLD

In the Gaussian World, three parameters characterize illuminants, surfaces or cone sensitivities. These are the spectral centroid; the spectral curvature, which is related to spectral dispersion or bandwidth; and the value at the centroid, which is a maximum or a minimum depending on the sign of the spectral curvature. We denote the spectral centroids for illuminant power distribution, surface reflectance and cone sensitivity by λ_I , λ_S and λ_C respectively, where λ_C is of course a different number for each cone type. Spectral curvatures are similarly k_I , k_S and k_C ; centroid values (maximum or minimum values, in either case scaling the entire function) are I_{\max} , S_{\max} and C_{\max} respectively. The following three equations then apply, with the parameter values appropriate to any given surface, illuminant and cone type.

$$\text{Illuminant spectral power density for illuminant I: } I(\lambda) = I_{\max} \exp[-k_I(\lambda - \lambda_I)^2]$$

$$\text{Surface spectral reflectance function for surface S: } S(\lambda) = S_{\max} \exp[-k_S(\lambda - \lambda_S)^2]$$

$$\text{Cone sensitivity function for cone C: } C(\lambda) = C_{\max} \exp[-k_C(\lambda - \lambda_C)^2]$$

If $k_I + k_S + k_C > 0$, as is generally the case since the cone sensitivity function is typically narrower than $I(\lambda)$ or $S(\lambda)$, the cone excitation elicited in cone C by surface S under illuminant I is given by:

$$\begin{aligned} E(I, S, C) &= \int_{-\infty}^{\infty} I(\lambda) S(\lambda) C(\lambda) d\lambda \\ &= I_{\max} S_{\max} C_{\max} \int_{-\infty}^{\infty} \exp\left[-k_I(\lambda - \lambda_I)^2 - k_S(\lambda - \lambda_S)^2 - k_C(\lambda - \lambda_C)^2\right] d\lambda \end{aligned}$$

The integral can be evaluated by grouping the terms involving λ and then completing the square in the exponent, transferring the term independent of λ outside of the integrand:

$$E(I, S, C) = I_{\max} S_{\max} C_{\max} \exp\left[-(k_I \lambda_I^2 + k_S \lambda_S^2 + k_C \lambda_C^2)\right]$$

$$\begin{aligned}
& \int_{-\infty}^{\infty} \exp\left[-(k_I + k_S + k_C)\lambda^2 + 2(k_I\lambda_I + k_S\lambda_S + k_C\lambda_C)\lambda\right] d\lambda \\
&= I_{\max} S_{\max} C_{\max} \exp\left[-(k_I\lambda_I^2 + k_S\lambda_S^2 + k_C\lambda_C^2) + \frac{(k_I\lambda_I + k_S\lambda_S + k_C\lambda_C)^2}{k_I + k_S + k_C}\right] \\
& \int_{-\infty}^{\infty} \exp\left[-\left(\sqrt{k_I + k_S + k_C}\lambda - \frac{k_I\lambda_I + k_S\lambda_S + k_C\lambda_C}{\sqrt{k_I + k_S + k_C}}\right)^2\right] d\lambda \\
&= k_0 \exp\left[-(k_I\lambda_I^2 + k_S\lambda_S^2 + k_C\lambda_C^2) + \frac{(k_I\lambda_I + k_S\lambda_S + k_C\lambda_C)^2}{k_I + k_S + k_C}\right]
\end{aligned}$$

where $k_0 = I_{\max} S_{\max} C_{\max} \sqrt{\frac{\pi}{k_I + k_S + k_C}}$. The square root factor here is the value of the

Gaussian integral in the previous expression.

The terms in the exponent can now be rearranged as a sum of squares, thereby representing the cone excitation as a product of three Gaussian factors (together with the intensity factor k_0 , which in many contexts may be neglected, since it is independent of the spectral variables but incorporates a relatively slight dependence on bandwidth).

$$\begin{aligned}
E(I, S, C) &= k_0 \exp\left[-\frac{1}{k_I + k_S + k_C} (k_I k_S (\lambda_I - \lambda_S)^2 + k_I k_C (\lambda_I - \lambda_C)^2 + k_S k_C (\lambda_S - \lambda_C)^2)\right] \\
&= k_0 \exp\left[-\frac{k_I k_S}{k_I + k_S + k_C} (\lambda_I - \lambda_S)^2\right] \exp\left[-\frac{k_I k_C}{k_I + k_S + k_C} (\lambda_I - \lambda_C)^2\right] \exp\left[-\frac{k_S k_C}{k_I + k_S + k_C} (\lambda_S - \lambda_C)^2\right]
\end{aligned}$$

(Equation 1)

The first Gaussian factor depends both on the illuminant and the surface parameters, and therefore might appear to jeopardize colour constancy based on normalization, since it can't be normalized out for all surfaces at the same time. But this factor is independent of the cone spectral centroid, which means that it generates variations in effective intensity only. What this factor means is simply that (for example) when the lighting gets red, the reds get lighter relative to the shorter wavelength surfaces. Unlike other deviations from strict normalization-compatibility mentioned below, this factor is potentially significant even in the important case of broadband illuminants and surfaces, that is if $k_S \ll k_C$ and $k_I \ll k_C$, provided neither spectral curvature is zero.

The second Gaussian factor similarly determines the illuminant colour. A key point to note is that this factor is approximately the same for all broadband surfaces, (surfaces for which $k_s \ll k_l + k_c$), as it then depends only on $(\lambda_l - \lambda_c)$, the difference between the illuminant spectral centroid and that of the cone spectral sensitivity. This means that this illuminant factor can be removed, for all such surfaces, by a simple normalization, in just the way that was possible only for effectively monochromatic radiations in the three-band world of section 1.

The final Gaussian is a factor representing surface colour, independent of the illuminant. It is a Gaussian function (or if $k_s k_c$ is negative, the reciprocal of one) of the spectral distance of the surface's spectral centroid from that of the cone. This is what normalization preserves, and it is what a colour constant visual system needs to preserve.

Colour differences approximately invariant with illumination. Chromatic values in the Gaussian World are most easily analyzed by considering intensity-invariant chromatic co-ordinates in log cone excitation space. One such coordinate is the difference in the logs of the L and M cone excitations. To evaluate this we replace λ_c by λ_M and λ_L respectively in Equation 1, assuming the same cone spectral bandwidths k_c in each case (both for simplicity, and because the bandwidths can be made equal by adopting a more suitable function of wavelength - approximately its logarithm - as the spectral variable at the outset). We also assume the same value peak cone sensitivity C_m ; this can always be arranged by appropriate choice of the units for the cone excitations. By substituting into Equation 1 we obtain the difference between the natural logs of the L and M cone excitations:

$$\log(L) - \log(M) = \log(L/M) = \frac{2(\lambda_L - \lambda_M)k_c}{k_l + k_s + k_c} (k_l(\lambda_l - (\lambda_L + \lambda_M)/2) + k_s(\lambda_s - (\lambda_L + \lambda_M)/2))$$

Equation 2)

This expression gives a particularly simple mathematical embodiment to the tradeoff between illuminant colour and surface colour in the determination of the retinal stimulus, since illumination and surface colour here contribute two symmetrical terms, derived respectively from the second and the third Gaussian factors in Equation 1. (The k_0 factor and the first of the three Gaussian factors in Equation 1 have been cancelled, being equal for the two cone types compared). The sum of the terms in

parentheses involving k_I and k_S is the effective spectral centroid of the stimulus, measured from the mean peak wavelength of the L and M cones. The difference in log excitation is proportional to that spectral distance and also to two other factors: the spectral separation of the L and M cone sensitivities, and a squared-bandwidth factor $\frac{k_C}{k_I + k_S + k_C}$. The bandwidth factor is an inverse measure of the spectral curvature of a cone's effective stimulus, relative to that of the cone sensitivity itself.

Similarly, if λ_B is the spectral centroid of the spectral sensitivity of the blue or short-wavelength cones (referred to elsewhere as the S cones, but here, to avoid confusion with the surface identifier, as the B cones, with excitation B),

$$\log(B) - \log(L) = \frac{2(\lambda_B - \lambda_L)k_C}{k_I + k_S + k_C} (k_I(\lambda_I - (\lambda_B + \lambda_L)/2) + k_S(\lambda_S - (\lambda_B + \lambda_L)/2)).$$

(Equation 3)

Equations 2 and 3 allow us to locate surface-derived stimuli in the chromatic plane $(\log(L) - \log(M), \log(B) - \log(L))$, shown in Figure 1.1. (Note that the quantities r and $\log(r)$ employed in much of our analysis of natural images are both practically linear with $\log(L) - \log(M)$, making the three quantities practically equivalent for most theoretical purposes. Likewise, $\log(b)$ is almost linear with $\log(B) - \log(L)$ for natural colours).

Under spectrally flat illumination ($k_I = 0$), spectrally non-selective surfaces ($k_S = 0$) are at the origin in Figure 1.1 since the illuminant term and the surface colour term are both zero. For selective surfaces, the nonzero surface colour term generates the stimulus coordinates

$$\frac{2(\lambda_L - \lambda_M)k_S k_C}{k_S + k_C} (\lambda_S - (\lambda_L + \lambda_M)/2) \quad , \quad \frac{2(\lambda_B - \lambda_L)k_S k_C}{k_S + k_C} (\lambda_S - (\lambda_B + \lambda_L)/2).$$

Loci of constant λ_S , traced out by varying k_S , radiate in straight lines from the origin, each coordinate being proportional to a purity measure, $\frac{k_S}{k_S + k_C}$, that approaches unity for monochromatic reflectances. Loci of constant surface bandwidth k_S lie on parallel straight lines, such as the crossbars of the "kites" in Figure 1.1, that have a slope equal to the ratio of the separations between the cone

sensitivity peaks. As a special case, the spectrum locus is entirely straight in these logarithmic coordinates, though not of course in a chromaticity diagram where the coordinates are linearly related to the cone excitations.

Under chromatic illumination, the illuminant term becomes nonzero and the spectrally non-selective surface shifts from the origin to the point

$$\frac{2(\lambda_L - \lambda_M)k_I k_C}{k_I + k_C} (\lambda_I - (\lambda_L + \lambda_M)/2) \quad , \quad \frac{2(\lambda_B - \lambda_L)k_I k_C}{k_I + k_C} (\lambda_I - (\lambda_B + \lambda_L)/2).$$

Selective surfaces undergo a comparable displacement. In the important range of conditions where the illuminant bandwidth and surface bandwidth are both large enough, relative to the cone bandwidth, to keep the factor $\frac{k_C}{k_I + k_S + k_C}$ close to 1, the displacement is approximately the same for all surfaces. Under such conditions, normalization achieves approximate colour constancy.

Deviations from this simple pattern of rigid translation take two forms: gamut compression due to illuminant bandwidth restriction and shift resistance due to surface-reflectance bandwidth restriction. If illuminant bandwidths differ, the narrower band illuminants (with high k_I) reduce saturation, compressing the constellation of stimulus colours by the factor $\frac{k_C}{k_I + k_S + k_C}$ toward the locus of the neutral (non-selective) surface under the prevailing illuminant; we refer to this as gamut compression. Likewise if surface bandwidths differ, variation in k_S reduces the shift of the narrower band surfaces toward the illuminant colour, by this same factor; in Figure 1.1 this is seen as a skewing of the kites as their radii each pivot around a fixed point on the spectrum locus. In the limit of narrow illuminant bandwidth, surface stimuli all assume the chromaticity of the monochromatic illuminant. In the limit of narrow surface bandwidth (monochromatic reflectance), changing illumination makes no difference to the chromaticity of the stimulus. Under narrow-band illuminants, shift resistance and gamut compression work together to compress colour differences orthogonal to the spectrum locus (“saturation” values). As a result these differences are compressed more than the “hue” differences that are parallel to the spectrum locus.

The illumination-invariant plane. The reason why changing illumination does not lead to strictly rigid translation in the coordinate system of Figure 1.1 is that the coordinates (from Equations 2 and 3) are

not additive combinations of an illuminant component and a surface component. Indeed, the bandwidth factor $\frac{k_C}{k_I + k_S + k_C}$ by itself is not additive in that sense. But its reciprocal is, and so too is the effective spectral centroid $k_I \lambda_I + k_S \lambda_S$. By adopting these two quantities as coordinates, we can create a colour plane that is illumination-invariant, in the sense that differences between any pair of surface colours are represented by illumination-invariant vectors. In this plane, the effect of changing illumination is always a precisely rigid translation of the constellation of points representing the difference surfaces within a scene. Geometrically, the plane of Figure 1.1 is a perspective view of this illumination-invariant plane. Instead of regarding the “kites” as vertically oriented in the plane of the paper and tethered to the spectrum locus below, they should be imagined as oriented in an orthogonal plane that recedes in depth toward the spectrum locus at infinite distance. The triples of apparently converging “tether” lines are parallel within the illumination-invariant plane, where they recede toward their vanishing points on the spectrum locus. The depth coordinate is linear with $k_I + k_S$, a spectral curvature value that approaches infinity for monochromatic stimuli. The horizontal coordinate is linear with the effective spectral centroid $k_I \lambda_I + k_S \lambda_S$. These coordinates of the illumination-invariant plane are each expressible as a ratio of linear combinations of the logs of the cone excitations. But the visual system has little to gain by computing them, since for natural colours, illumination-dependent mapping is generally well approximated by translation in Figure 1.1, and is therefore normalization-compatible.

Figures

Figure 1:





Figure 2:

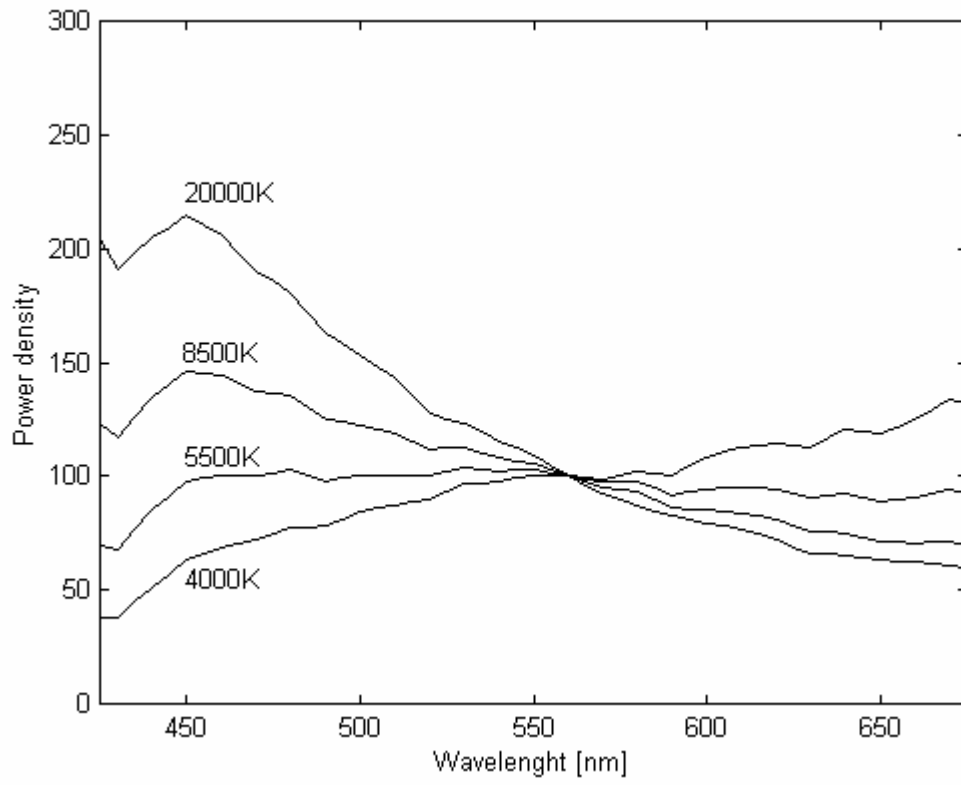


Figure 3:

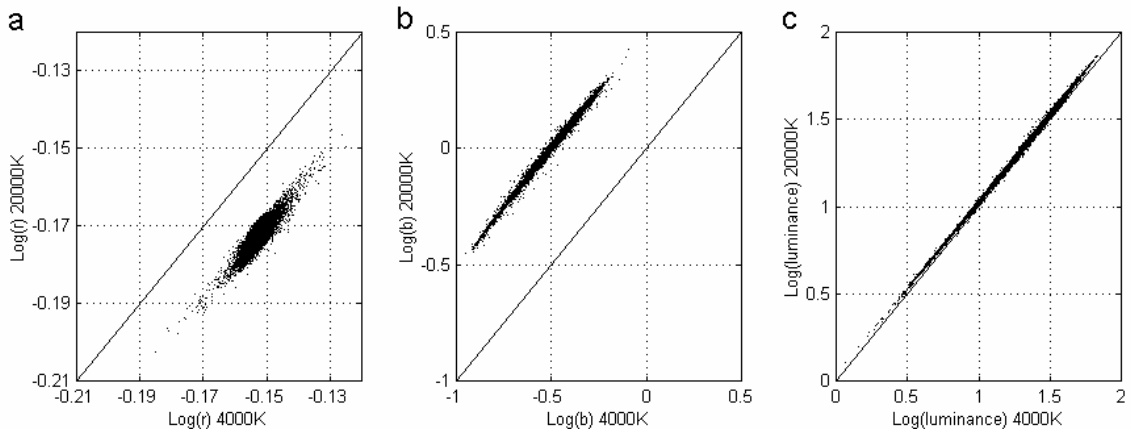


Figure 4:




Figure 5:


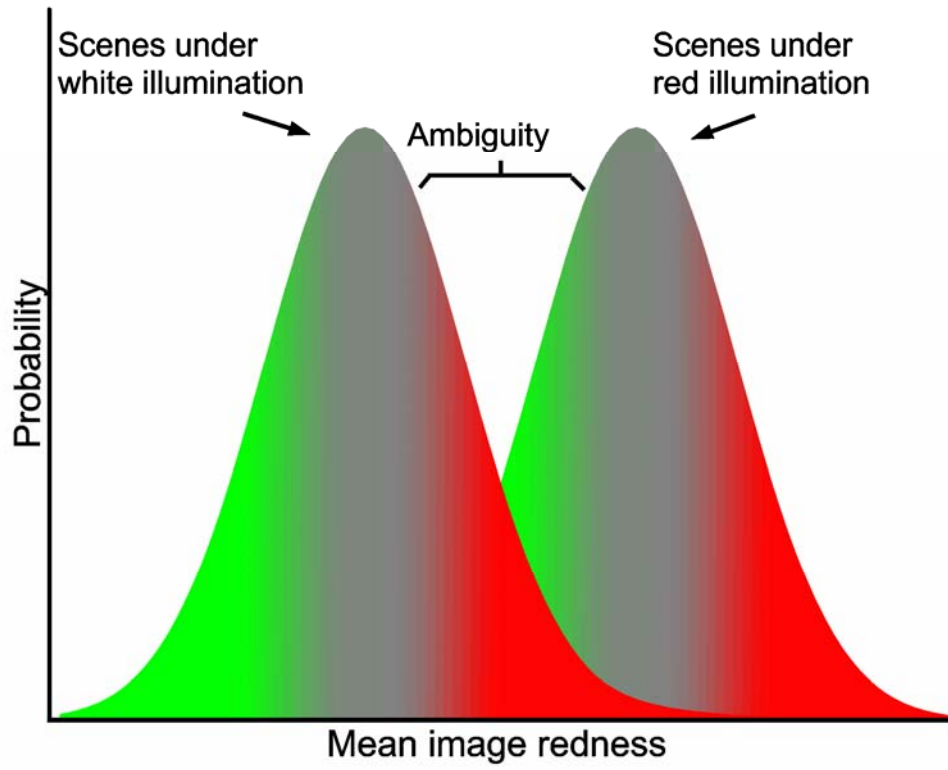


Figure 6:
a

b

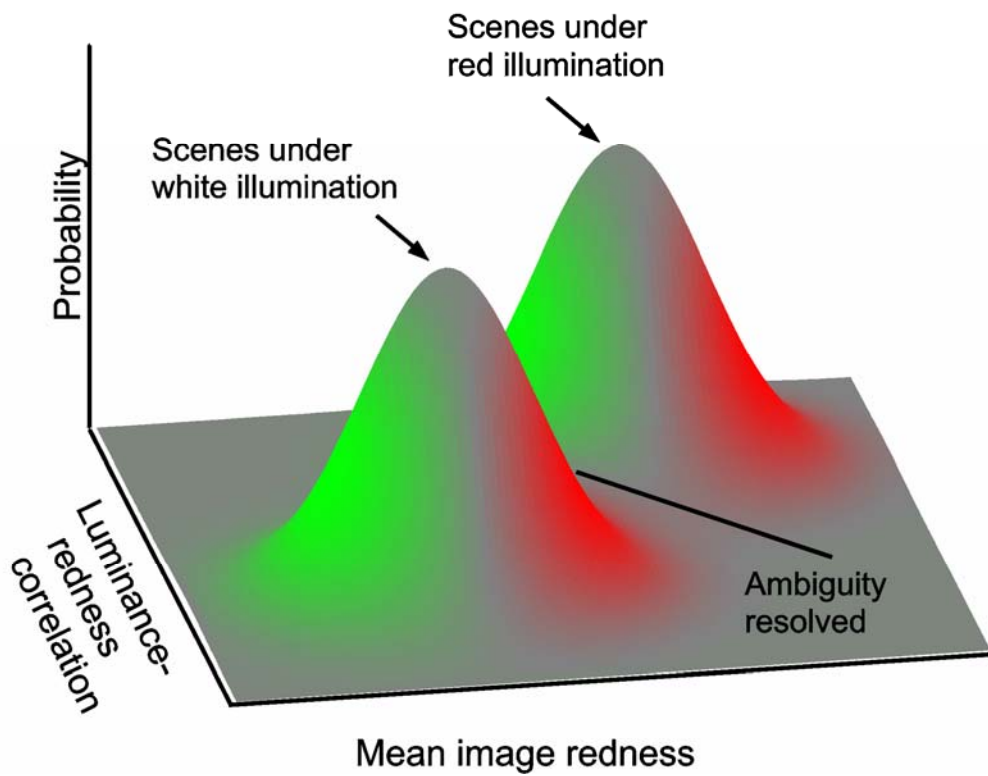


Figure 7:




Figure 8:




Figure 9:

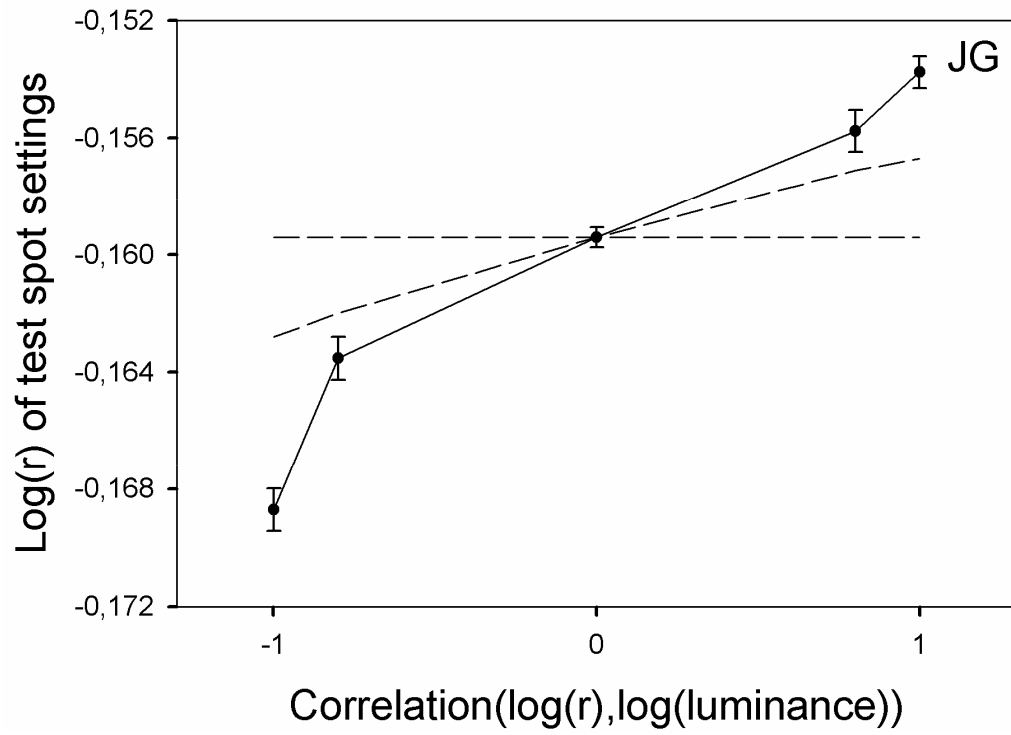


Figure 10:

