

## **Colour discrimination, colour constancy and natural scene statistics**

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

### ***Abstract***

Does the visual system allocate discriminative ability to different regions in colour space in a way that optimizes discrimination among natural colours? If so, discrimination should satisfy a “cube root rule”: in an optimized system differential sensitivity will be greatest for the most commonly encountered conditions, dropping to half its maximum under conditions of relative frequency  $1/8$ . Quantitatively, this principle is only very roughly consistent with psychophysical data, but it does account for some of the salient findings, such as the relative sensitivity for different directions in colour space; the fit between theory and psychophysical observation is improved by considering the stimulus to be the local contrast between test field and background, rather than absolute luminance and chromatic values of individual pixels. Comparison with physiological data shows less satisfactory agreement: M cells appear to be too nonlinear, and P cells too linear, for optimal metric representations of luminance and colour respectively. The good colour discrimination of some strongly anomalous trichromats may result from an optimization, during development, of postreceptoral nonlinearity to match the limited range of inputs delivered by the anomalous photoreceptors.

For natural colours under natural illuminants, the cone excitations for all surfaces in an image are scaled by approximately the same factor with a change of illumination. This allows the effect of varying illumination to be simply corrected by reciprocal adjustments of sensitivity in the different cone types. The resulting representation is illumination-invariant, but also fails to preserve information about the overall chromatic

cast of a scene. Yet experimentally, colouration of the image is perceptually attributed in part to the illuminant and in part to the viewed surfaces, resulting in “underconstancy”.

When the statistical variation among natural illuminants and scenes is considered, underconstancy can be viewed not as a failure of constancy, but as a best guess about illuminant colour appropriately based on knowledge of relevant environmental statistics.

Natural images do generate small deviations from the scaling principle. These can provide useful cues to the illuminant: statistics (other than the mean) of the distribution of an image’s elements in cone excitation space can in principle resolve the ambiguity inherent in the mean alone. Experiment suggests that vision does exploit these cues and gives them statistically justifiable weight.

## **1. Introduction**

This year in Cambridge we remember Thomas Young (Mollon, this volume), but in this particular college, Peterhouse, we also remember James Clerk Maxwell, who so admirably filled out Thomas Young’s sketch of the trichromatic theory while still an undergraduate here and at Trinity (Sherman, 1981). I think of Maxwell as a kind of prototype for all of us, because he combined a nice appreciation for the phenomena of colour with a penchant for quantitative analysis and mechanistic modeling. Although, as Goethe complained, all theory is gray, treatments of colour vision ever since Maxwell’s time have tended to involve quantitative models, and the present one is no exception. My discussion also typifies a much more recent theoretical trend: I will attempt to relate the processes of colour vision to the characteristics of the natural environment. The reference to natural scene statistics in the main connection between the two parts of the chapter—the first (sections 2 to 6) concerned with colour discrimination, and the second (section 7)

with colour appearance and its transformation or constancy under changes of illumination. Also central to both discussions is the indispensable and now familiar intellectual device that we owe to Maxwell: cone excitation space, the alpha and omega of colour spaces, where 3 Cartesian co-ordinates specify the colour stimulus through the rates of isomerization of the three visual pigments of normal trichromatic colour vision.

## ***2. Discrimination and the distribution of natural colours***

Following von der Twer and MacLeod (2001), colour discrimination is here considered as a slicing of cone excitation space into distinguishable cells, and the questions considered are: What is the distribution of naturally occurring stimuli in cone excitation space? Is the slicing pattern that is revealed in psychophysical measures of colour discrimination well adapted for representing colours from that distribution? And how is the slicing pattern related to the neural representation of colour within the visual system?

The key idea to be developed is that it is advantageous to slice cone excitation space more finely in regions where natural colour stimuli are most densely concentrated. Imagine that successive, reliably different, levels of each neural signal define bins or slices into which our visual system can reliably divide colour stimuli (for instance in Figure 1 at left, separate signals originating from the L and M cones slice the shown plane of the space horizontally and vertically with 10 distinguishable levels each). The total number of these slices is fixed for a given signal by the range of output signal strength, going from zero to some maximum firing rate, and also varies reciprocally with the degree of intrinsic randomness or noise present in the neural signal. It has been clear

since the time of Fechner and Helmholtz that the visual system slices cone excitation space in a distinctly non-uniform manner. In some respects at least the non-uniformity takes the form of finer slicing where natural stimuli are most frequent. For example: the first characteristic of the slicing pattern to be recognized (Helmholtz, 1896) was that equal percentage changes in the linear cone excitations are roughly equally detectable, as if the excitation differences seen by each cone are divided by the mean excitation, a scenario supported by psychophysical (He & MacLeod, 1997, MacLeod, Williams & Makous, 1992) and physiological evidence (Dacey, Lee, Stafford, Pokorny & Smith, 1996). Since equal percentage changes correspond to equal intervals in the logarithm, a log-cone-excitation space in which each cone excitation is logarithmically compressed is more uniformly sliced than the linear one. The approximately logarithmic slicing pattern is efficient for representing natural surface colours, since natural stimuli are more symmetrically distributed in the logarithmic cone excitation space (Ruderman, Cronin & Chiao, 1998), without the concentration of dark colours near the origin seen in the linear space of Figure 1.

But two main features of the distribution of natural colours are not well dealt with by independent compression of the cone excitations. First, the stimulus distribution is centrally peaked: near-grays of moderate reflectance are relatively common, and the margins of cone excitation space, even in its logarithmic form, are only sparsely populated. Accordingly, fine discrimination is less important in the margins of colour space than in the centre, near white, where natural colour stimuli are most densely packed. Much psychophysical evidence on colour discrimination has indeed indicated a slicing pattern with best discrimination near white (or more precisely, near the stimulus to

which the visual system is adapted) and a progressive deterioration as the saturation of the reference stimulus increases (Friele, 1961, LeGrand, 1949, Mollon, 1982, Pugh & Mollon, 1979). Superior discrimination for the naturally abundant near-neutral colours cannot be explained by independent compression of the three cone excitations, but requires a comparable compressive nonlinearity in the neural code at the colour-opponent level, where it is the differences between cone excitations that are represented (Friele, 1961). The requirement for finer slicing near white thus yields a rationale for colour opponent encoding.

Second, the distribution of natural colours in cone excitation space is highly anisotropic, with much less dispersion in the red-green direction (the negative diagonal in Figure 1, right) than in the luminance direction (the positive diagonal) or in the blue-yellow direction. Correspondingly, colour discrimination thresholds in the red-green direction, expressed in terms of cone contrast, are an order of magnitude lower than they are in the achromatic or luminance direction (Chaparro, Stromeyer, Huang, Kronauer & Eskew, 1993), implying finer “slicing” of log-cone-excitation space in that direction than in the luminance direction. This anisotropic slicing pattern also calls for colour-opponent signals. In the nonlinear analysis to be developed here, it is the reduced dispersion of natural inputs in the red-green direction that makes such a slicing pattern feasible and advantageous. Indeed, according to von der Twer and MacLeod (2001), only a nonlinear analysis can properly justify this obviously appropriate choice of slicing directions, although the principle of “decorrelation” within a linear framework dictates a similar choice (Buchsbaum & Gottschalk, 1983, Fukurotani, 1982, Zaidi, 1997).

Von der Twer and I considered quantitatively how the visual system's distribution of discriminative capacity over cone excitation space ought, in principle, to depend on the distribution of natural colour stimuli as represented in two data sets (Brown, 1994, Ruderman et al., 1998). For 574 haphazard samples of natural colours measured in San Diego by Richard O. Brown, the correlation between L and M cone excitation is 0.985. These samples were not selected in such a way as to be representative in any sense, and they may have included a disproportionate number of highly chromatic surfaces since these would more readily have caught the investigator's eye. Ruderman, Cronin and Chiao (1998) obtained spectral reflectance estimates, pixel by 3 min arc pixel, for 12 entire views of natural environments. For this data set comprising nearly 200,000 pixels, the correlation between L and M cone excitation is even higher at 0.9983. In that sense, the L and M cones measure almost the same thing—luminance—when used in natural environments. Surface luminance is given simply by the summed excitations of L and M cones, which we denote here simply by L and M (Eisner & MacLeod, 1980, Lennie, Pokorny & Smith, 1993). The standard deviation of  $\log_{10}(L+M)$  in Brown's data set is 0.46, which corresponds to a factor of three in luminance; for Ruderman et al.'s data the value is 0.24, a little less than a factor of two. A convenient measure of the red/green dimension of colour is the L cone excitation per unit luminance,  $r = L/(L+M)$ . As the high *LM* correlation implies, the standard deviation of  $r$  is strikingly smaller than that for luminance: only 7.5%, or 0.03 in the decimal logarithm for Brown's data, and only 1% for the entire scenes of Ruderman et al. For the remaining chromatic axis, we adopt the luminance-normalized S cone excitation,  $b = S/(L+M)$ , which is low for yellows and very high for violets. The standard deviation of  $b$  is an order of magnitude greater than for  $r$ ,

though still lower than the one for luminance: 0.39 in  $\log_{10}(b)$  or a factor of about 2.5 in  $b$  for Brown's data, and 0.114 in  $\log_{10}(b)$  or 30% in  $b$  for the data of Ruderman et al (crosses in Figure 2 show the histogram for  $\log_{10}(b)$ ). In both these data sets, the average of the natural stimuli is near to, but slightly yellower than, the equal energy white, which plots at a horizontal coordinate of zero in the Figure.

We can now ask: what arrangement of slice thicknesses could provide the most precise representation, averaged over all cases, of natural colours conforming to the given distribution? This criterion allows determination of an optimum nonlinear code, subject to the constraint of a limited output range. As will be seen in the next section, the optimal code with a single neuron per stimulus dimension is a sigmoid. For two neurons per dimension, it is a split range code employing rectifying opponent cells (Section 4). In Section 5 this theoretically optimal behaviour is compared with psychophysical results on the one hand and with physiological data on nonlinearity of the postreceptoral neural code on the other.

### **3. Optimal nonlinearity: the pleistochrome**

The optic nerve constitutes an informational bottleneck for vision, where the number of nerve fibers is relatively limited, and the number of nerve impulses is much less than the number of absorbed photons at daylight light levels (Barlow, 1965). Hence relatively large errors are introduced by random fluctuations in the optic nerve impulse counts (Bialek & Rieke, 1992, Lee, Wehrhahn, Westheimer & Kremers, 1993). So initially we assume, for simplicity, that variability originating at the retinal output predominates over sources of error at earlier stages.

When discrimination is limited by noise introduced into an output signal, good discrimination around any input value can be achieved by making the gradient of the input-output function at the relevant point as steep as necessary. But the limited total available range in firing rate means that an increase in gradient at one point has to be paid for by a decrease at other points within the input range, and hence by reduced discrimination at those points. By suitable choice of a nonlinear response function, relative discriminative precision can be distributed in any desired way over the range of input values. But which choice is best? For example, what input-output function gives the smallest RMS error in the estimated input, averaged over all naturally occurring cases?

Clearly it would be inefficient to make the code linear (with constant gradient and constant discrimination) over an input range greater than what is naturally encountered, as in the straight line in Fig. 2, where the crosses represent the natural distribution of  $\log_{10}(b)$  from the data of Ruderman et al. This sacrifices discrimination among frequently occurring stimuli in order to preserve discrimination in ranges where it is never needed. Also inefficient is the opposite extreme, a response function that steps abruptly from minimum to maximum firing rate at the peak of the distribution of natural colours (step function, Fig. 2). This provides only categorical colour perception, with exquisite discrimination between bluish and yellowish colours, but no distinction among blues or among yellows. The best choice will be an intermediate one (circles, Fig. 2): a gently curving sigmoid, which retains some discrimination in the tails while slicing colour space most finely at the peak. More specifically, the optimal response function has a gradient matched to the cube root of the probability density function of the input



distribution, and generates discrimination thresholds inversely proportional to that gradient.

Denote the input by  $x$  (which might be, for instance, some weighted sum of cone excitations), and the noisy output (e.g. a retinal ganglion cell firing rate) by a random variable  $y^*$ , with mean  $y = g(x)$ , a generally nonlinear function of  $x$ , and with a standard deviation  $\mathbf{s}(y^*)$  due to the addition of a random noise term,  $y^* - y$ . With the simplifying assumption that  $g(x)$  can be treated as linear over the limited range of confusion, the input  $x^*$  that would elicit  $y^*$  in the absence of noise is

$$x^* = g^{-1}(y^*) = x + (y^* - y) / g'(y)$$

Since the mean of  $(y^* - y)$  is zero, the mean of  $x^*$  is simply the true input  $x$ , and the mean squared error for input  $x$  is the variance of  $x^*$ , which we denote  $\mathbf{s}^2(x^*)$ . This is proportional to the output variance  $\mathbf{s}^2(y^*)$ , which we initially assume to be fixed and independent of  $y$ :

$$MSE(x) = \mathbf{s}^2(x^*) = \frac{\mathbf{s}^2(y^*)}{(g'(x))^2} \quad (1)$$

The constant of proportionality is the inverse square of  $g'(x)$ , the gradient of the input-output function  $y = g(x)$  at  $x$ : where  $g(x)$  is steep, a smaller range of inputs suffices to span the range of the noise that is added at the output.

Denote the environmental probability distribution of  $x$  (for all stimuli encountered, or of interest) by  $p(x)$ . The mean squared error to be minimized is the average of  $\mathbf{s}^2(x^*)$  for all inputs, or its probability-weighted integral over  $x$ :

$$MSE = \int p(x) \mathbf{s}^2(x^*) dx$$

To see how this can be minimized, consider the effect of small variations in the response gradient or incremental gain  $g'(x)$  around its optimal value. An increase in  $g'(x)$  at one value of  $x$  has to be paid for with an equal decrease at other values of  $x$ , and in the optimal condition the effects of such complementary changes must cancel, that is  $p(x)d(\mathbf{s}^2(x^*))/dg'(x)$  must be independent of  $x$ . Thus

$$d(\mathbf{s}^2(x^*))/d(g'(x)) = \alpha/p(x) \quad (2)$$

with a constant  $\alpha$ .

Since from (1)

$$d(\mathbf{s}^2(x^*))/d(g'(x)) = -2(\mathbf{s}^2(x^*))/(g'(x))^3$$

the optimal condition occurs when the gradient is matched to the cube root of the pdf:

$$g'(x) = \mathbf{b} p(x)^{1/3} \quad (3)$$

The scaling factor  $\mathbf{b}$  absorbs the factor equal to the output variance. It serves only to define units of measurement for  $y$  and can be set to 1 if those units are not defined independently. Hence

$$g(x) = \int_{-\infty}^x \mathbf{b}(p(u))^{1/3} du \quad (4)$$

Here the input quantity can be any function of the cone excitations (e.g. the log, in this example.) The optimum response gradient just has to be matched to the cube root of the probability distribution of the input quantity for which error is to be minimized.

Von der Twer and I refer to this optimal response function, illustrated by the circles in Figure 2, as the *pleistochrome*, from the Greek *pleistos* meaning ‘most’. As Fig. 2 shows, the pleistochrome is roughly similar to the cumulative distribution of environmental stimulus values dictated by the principle of histogram equalization and suggested by Simon Laughlin, the first person to consider these issues quantitatively (Laughlin, 1983). With histogram equalization all of the “slices” are equally populated. But the pleistochrome is wider than the histogram-equalization sigmoid by a factor of about the square root of three, allowing much better discrimination in the margins of colour space.

#### **4. The benefit of split range coding**

Our discussion has not yet yielded any rationale for the familiar colour-opponent codes (Derrington, Krauskopf & Lennie, 1984, DeValois & DeValois, 1975). On the contrary, the sigmoidal nature of the pleistochrome is incompatible with a null response

to white. But when more than one neuron is available to represent a single stimulus dimension, new opportunities for coding are introduced. The non-opponent pleistochrome that optimizes encoding by a single neuron has an opponent counterpart when the encoding is done by a pair of neurons. Two rectifying neurons, one red-excitatory and the other green-excitatory, and each with a purely compressive nonlinearity, can represent opposite halves of the red-green stimulus continuum with positive firing rates, but with little or no response to greenish or to reddish stimuli respectively. Such a representation is almost equivalent to that produced by a single neuron with sigmoidal nonlinearity (Marr, 1974). The responses of two such neurons correspond to the two halves of the single-neuron pleistochrome sigmoid, but with the left half flipped up so that the response gradients for the neuron responding in the left half of the stimulus range are simply reversed. This, however, uses only the upper half the output range of each neuron. If the appropriate segment of the input range elicits the maximum possible response range from each neuron, the gradients of the response functions are doubled everywhere.

By using two neurons in this way the visual system can therefore double the precision in its representation of the input in the presence of output noise. If, alternatively, the two neurons had each been endowed with the same sigmoidal nonlinearity that is optimal for single neurons, then averaging of their signals (on the generous assumption of independent noise) would have reduced average error by only the square root of two. Thus the net benefit of adopting the 'split range' code (as opposed to the alternative of similar neurons operating in parallel with optimal nonlinearity) is a square root of two reduction of average error.

*Location of white.* If the split-range colour opponent code is designed for optimal characterization of natural colours in the least-average-error sense, and if the null stimulus for the colour-opponent neurons is the subjectively achromatic white, typical natural colours should be nearly white. This is of course roughly correct, but the prediction is not fulfilled exactly. Typical natural colours, at least in the chosen environments, where vegetation tends to be predominant, are greenish and yellowish. There are a number of more or less plausible post hoc rationalizations for the somewhat unexpected placement of the white point (and of the optimum point for colour discrimination). First, although most natural surfaces are yellowish, the sky is bluish. If equilibrium hue loci are adaptively fixed by the average input, the blue of the sky might act as a massive low- $r$  and high- $b$  counterweight to shift the environmental mean substantially from the mean of surface colours. Second, the location of the white point could be a compromise between optimizing discrimination for the most frequent surface colours (which, if the images of Ruderman et al. are typical, would put it in the part of colour space we actually identify as yellow-green) and preserving some discrimination for saturated blue and red surfaces. Even if this choice is not optimal by the unweighted least-squared-error criterion, it could be appropriate if saturated colours (other than the greens) tend to have greater than average biological importance. Osorio and Bossomaier (1992) suggest that discrimination among the greens of vegetation is not important or even desirable, whereas discrimination of reddish fruits from vegetation is important. A null point, with optimal discrimination, near white might usefully promote those discriminations at the expense of the less important ones. Such considerations are the

focus of much recent discussion (Regan, Julliot, Simmen, Viénot, Charles-Dominique & Mollon, 2001)

### **5. Comparison of the pleistochrome with psychophysical and electrophysiological data**

*Testing the cube root rule.* The proposal that visual nonlinearity is optimized for discrimination among natural colours in the presence of output noise leads to a simple prediction: *the average error in visual discrimination or matching should be inversely proportional to  $g'(x)$  in equation (4), and hence to the cube root of the natural probability density function  $p(x)$ .* This prediction holds even for non-uniform output noise, provided that the nonlinearity  $g(x)$  is appropriately optimized in each case (von der Twer and MacLeod, 2001).

Psychophysical data for evaluating this prediction are available. For achromatic intensity differences between comparison and standard stimulus patches, set in a common grey background, comparison error is doubled for standards that create a contrast of about 20% with the background. For isoluminant yellow-blue differences the cone contrast (for S cones in this case) at which error is doubled is again about 20%. But in the case of red-green isoluminant stimuli a standard L cone contrast of only 2% is enough to double the mean comparison error (Figure 3; Leonova and MacLeod, in preparation).<sup>1</sup>

---

<sup>1</sup> The implicit assumption here that contrasts along the three axes of color space are perceptually related by simple scaling factors is consistent with contrast matching results of Switkes and Crognale(1999). Switkes and Crognale, however, found greater relative sensitivity to achromatic contrast. Their parafoveal grating stimulus may have been high enough in spatial frequency to place the chromatic mechanisms at a

If the visual system adopts the encoding principle of the pleistochrome, we would therefore expect to find the probability density function  $p(r)$  for natural colours dropping to 1/8 of its peak value at  $r$  values that give a contrast of 2% with white, with half-heights at around 20% contrast for the other axes. This prediction is in one respect roughly upheld, since as we saw in Section 2, the distribution of environmental colours the ranges for luminance and for  $b$  are indeed wider than for  $r$ , by an order of magnitude or so. In that sense the operating ranges of the various relevant neurons are fairly well matched to the very diverse distributions of environmental inputs that they have to represent. Webster and Mollon (1997) likewise noted that the mean contrast of their images along each of these axes was very approximately a constant multiple of the visual contrast threshold. In fact, however, the environmental distributions of  $b$  and (especially) of luminous reflectance are somewhat broader than would be expected for strict consistency with the pleistochrome principle (Table 1). Why should the operating range of the visual system be narrower than ‘optimal’ in this way?

*Importance of contrast.* One answer appeals to visual adaptation. Since retinally stabilized images fade, vision is evidently sustained by the temporal transients that are generated when small eye movements scan spatial gradients in the image. It is therefore relevant to consider the distribution of the spatial *differences* in cone excitation in natural images. In the images of Ruderman et al, the differences in luminance, in  $b$ , and in  $r$

---

disadvantage. The spatial frequency dependence of the relative sensitivities for difference directions in color space (Parraga et al., 1998) is indeed not fully consistent with the statistics of natural scenes, a fact that complicates the simple correspondence noted in this section and limits the generality of the cube root rule.

between adjacent 3 min arc pixels, have standard deviations of only 30%, 13% and 0.6% respectively. The pdfs drop to 1/8 (where optimal differential sensitivity is halved) at about twice these values. These distributions are naturally somewhat tighter than those for the absolute values, owing to correlated variation in the values across the scene. The visual system can therefore advantageously employ an adaptively roving null point for the colour-opponent code (Krauskopf & Gegenfurtner, 1992, Thornton & Pugh, 1983), if its objective is the precise representation of local contrast (from which a metric representation of local brightness and colour can then be resurrected, perhaps in the way discussed by Land and Marr (Land, 1959, Marr, 1974). With a roving null point, the range of input values spanned by the neural response functions need only be wide enough to capture the relatively small *deviations* in the stimulus values from their time and space varying adapting levels—and the precision with which those values can be represented then becomes correspondingly greater. Local-contrast pleistochromes—contrast-response functions that lead to least error in the representation of pixel-wise spatial differences in the images of Ruderman et al.—are narrow enough to be fairly consistent with the cited psychophysical results in the case of the chromatic variables. (Table 1)

This analysis implies, in agreement with observation (Leonova and MacLeod, in preparation; Smith, this volume) that the very sharp optimum in colour discrimination is observed only when test and reference field are embedded in a common uniform surround that sets the adaptation state.

When light adaptation provides accurately reciprocal sensitivity adjustments in each cone (Dacey et al., 1996, He & MacLeod, 1997, MacLeod et al., 1992), postreceptoral signals become functions of contrast. In the classical bipartite field for



which the MacAdam ellipses (MacAdam, 1942) were measured, the contrast between the two halves of the field is nearly zero. Postreceptoral neurons sensitive to contrast then operate at a fixed point on their response function for all colours to which the observer is well adapted; the margins of colour space no longer need be disadvantaged. On this view, all the MacAdam ellipses are really the same ellipse, but measured with different cone sensitivities. As noted, LeGrand (1949) showed that this view is too simple: MacAdam's discrimination thresholds do increase somewhat with increasing saturation of the reference colour. But that increase is far less precipitous than what is found with a common uniform surround for reference and test (Figure 3; Smith, this volume). In view of the effort that has been lavished on the refinement of colour difference formulas to account well for MacAdam's data, it is important to remember that when test and reference fields are non-adjacent, the pattern of results is very different. Colour difference formulas for that situation have yet to be worked out.

*Discrepancies for luminance.* For luminance, the contrast operating range implicit in the discrimination results is narrower by a factor of about three than the theoretically optimal pleistochrome. A second limitation in our initial framework, also connected with the role of adaptation, may underlie this remaining discrepancy. We have taken for granted that the purpose of colour and lightness vision is to represent colours and lightnesses with the least possible error and allow these attributes of a surface to be estimated as precisely as possible. But of course, differences in lightness and colour are also indispensable for the detection of spatial features (Boynton, 1980, Morgan, Adam & Mollon, 1992). For spatial vision, local contrasts should be detected with the greatest

possible sensitivity wherever they are present in the image. For this purpose, an all-or-none or categorical encoding scheme, with a step function nonlinearity at a small threshold offset from the adapting background stimulus is ideal (since the large, all-or-none spatial contrast signal resists obliteration by fluctuations in the output), and the graded response of the pleistochrome is not needed. Visual nonlinearity more step-like than the pleistochrome could therefore reflect a compromise in design between the conflicting requirements of surface identification and characterization on the one hand, and of detection of spatial features on the other. This is supported by the electrophysiological data considered below.

*Comparison with physiological results.* We next apply a Fechnerian construction (Fechner, 1860) to derive nonlinear opponent codes for lightness and colour from discrimination data such as those of Figure 3 and to compare the result with physiological data on the response functions of single neurons in the optic nerve.

In Fig. 3, predictions for the two extreme cases that were introduced in Fig. 2 are illustrated for comparison with the data. A linear code predicts uniform precision of discrimination (horizontal dashed line). An all-or none response, that distinguishes sharply between reddish and greenish colours but makes no distinction among the colours of each category, permits standards of any redness to be distinguished only from greenish tests, and vice versa; hence the threshold  $\mathbf{Dr} = |r - r^*|$ , where  $r^*$  is the colour category boundary. In Fig. 3 the steep dashed V illustrates this prediction, assuming a category boundary at  $r^* = 0.7$  (the value for white). Neither of these extreme models describes the data well; the condition for discrimination is neither constant nor as abruptly standard-

dependent as the step nonlinearity would require. Instead, the linear increase in threshold on each side of the white point suggests, by a straightforward extension of Fechner's argument to the colour domain, a logarithmic compression of each of the two colour-opponent neural signals that form the split range code. The linear variation of the discrimination threshold with  $r$  on each side of the null point  $r = .7$  in Fig. 3, with an abscissa intercept at  $r_0$ , leads to a response-intensity function of the form

$$N = \ln|(r - r_0)| \quad (5)$$

where  $r_0$  has a value of about .714 for the 'green' response (applicable for  $r < .7$ ) and .68 for the 'red' response (applicable for  $r > .7$ ). The value obtained by reflecting  $r_0$  around the null point,  $r = 0.7 + (0.7 - r_0)$ , is the value of  $r$  associated with a doubling of threshold, or a halving of differential sensitivity. This condition occurs at an L cone contrast of about 2% with respect to the null white stimulus in each case.

By reversing the argument that led to Equ. (4), one can then ask: for what distribution of environmental inputs is the Weber Law discrimination function—and the logarithmic response nonlinearity of Equ. (5)—optimal? The answer is  $p(x) = p_{max} / (1 + (x/x_0))^3$ . This function does fit tolerably well the central core of the distribution of local contrast in the images of Ruderman et al. Whether we accept Fechner's integration or not, the need to perceptually reconstruct values distributed in this way adds a new functional rationale for Weber's Law for contrast.

The nonlinearity implied by the reviewed psychophysical data is quite severe. The gradient of the red-green response function, assumed in Equation (4) to be directly proportional to differential sensitivity, is halved at an L cone contrast of roughly 2%. No

physiological data suggest so severely compressed a response function for responses to chromatic stimuli: half-saturation L cone chromatic contrasts of around 10% appear to be more typical, for the red-green sensitive P cells of the parvo-cellular stream (Lee, Pokorny, Smith, Martin & Valberg, 1990). Thus although the psychophysically estimated visual operating range along the red-green axis is efficiently matched to the range of environmental inputs, the physiological one apparently is not. Elsewhere (MacLeod and von der Twer, in preparation) we consider reasons for this. These include the possibility that output noise increases with mean firing rate (as in a Poisson process), which makes the optimum physiological nonlinearity a more gentle one.

Along the achromatic axis of colour space, both the psychophysical operating range in cone contrast and the dispersion of the environmental inputs are (as noted above) at least tenfold greater than along the red/green one. The M cells of the magnocellular pathway are a plausible substrate for achromatic discrimination (Livingstone & Hubel, 1987). These cells, however saturate at very *low* achromatic cone contrasts, with half-saturation values of around 5% (Kaplan, 1986, Lee et al., 1990, Wachtler, Wehrhahn & Lee, 1996). The M cells thus deviate by an order of magnitude from the optimal behaviour embodied in the pleistochrome. They could not support the observed keen discrimination between test patches with relatively high achromatic contrast relative to their surrounds.

It is therefore likely that the M cells are not responsible for representing the achromatic attributes of surfaces in a continuous fashion, but serve instead as all-or-none detectors of spatial contrast and of change over time. Thus, although responses of P cells to red/green contrast incorporate a degree of nonlinearity only slightly less pronounced

that that of the M cells for achromatic contrast, comparison with natural scene statistics actually suggests different functions for these two systems. The red/green nonlinearity of the P cells is roughly appropriate for a metric representation of their natural inputs, but the nonlinearity of the M cells is not. This is consistent with the common view that the luminance system (sometimes tentatively, though questionably, identified with the magnocellular pathway) is more concerned with form and with detection of spatial structure than are the chromatic ones (Boynton, Hayhoe & MacLeod, 1977, Gregory, 1979, Livingstone & Hubel, 1987). The metric representation of the achromatic axis as well as the chromatic axes could be the job of the P cells (Allman & Zucker, 1990, Pokorny & Smith, 1997), which have an almost linear response to achromatic contrast.

### ***6. Is the nonlinearity optimized developmentally? Anomalous observers as a test case***

Leaving aside the discrepancies noted above, we have seen that the operating ranges of postreceptoral neurons, or the ones implicit in psychophysical discrimination performance, are very roughly appropriate to the range of inputs they receive from natural environments. But is the appropriate form of the neural nonlinearity genetically determined (and specified via natural selection during evolution), or is it shaped during development? In anomalous trichromats, nature has provided a test case that may decide this point. Here our theoretical concerns finally make contact with the work of Guy Verriest, by providing a new perspective on deficiencies of trichromatic colour vision.

Previous treatments of colour discrimination by anomalous observers have implicitly neglected output noise, supposing that colour discrimination thresholds will,

like the cone contrasts, vary inversely with the separation between the L and M pigment absorption curves. This expectation is not generally upheld: although there is some disagreement about the interpretation of the finding (Pokorny & Smith, 1977), discrimination and pigment spectra (assessed from colour matching measures) vary almost independently among anomalous observers (Hurvich, 1972, Pickford, 1958, Pokorny, Smith & Verriest, 1979, Verriest, 1960, Willis & Farnsworth, 1952), .

Consideration of output noise can resolve this paradox. If discrimination is limited by noise added to the colour opponent signal, discrimination around any given point in colour space can be improved, in principle without limit, by increasing the gain of the stage relating the colour opponent signal to the receptor signals. The total number of just-noticeable differences, or slices, along any chromatic continuum is then fixed by the output noise amplitude and output signal range, not by the photoreceptor inputs. With a plastic nonlinearity optimized to the range of photoreceptor excitations generated by the natural environment, the entire available range of output signals can be mapped onto the range of inputs made available by any particular trichromat's photoreceptors. The range of signals might in this way "expand to fill the neural space available" (Regan & Mollon, 1997). Discrimination thresholds for trichromats can therefore in principle be completely independent of the (non-zero) separation between the pigment sensitivities!<sup>2</sup>

---

<sup>2</sup> Alternatively, if the nonlinearity is genetically fixed, or for any reason fails to be shaped appropriately by input during development, the impoverished cone contrasts provided by the anomalous observer's photoreceptors will fail to span the full range of the colour opponent signal, and discrimination will then be impaired just as if noise were added at the input (e.g. to the photoreceptor excitations themselves). In this case, the colour opponent signals of anomalous trichromats, being smaller than those of normal observers,

Thus the existence of strongly anomalous observers who violate the naively expected relation between pigment spectra and discrimination by showing near-normal colour discrimination at all chromatic contrasts can be explained by—and taken as circumstantial evidence for—both the importance of output noise, and plasticity of colour-opponent nonlinearity. I myself may be such a case. I am deuteranomalous, with a very deviant match point (Nagel  $\log_{10}(AQ) = 0.56$ ). The implication that my anomalous M photoreceptor spectral sensitivity is close to my normal L sensitivity is supported by dichromat-like behaviour under selective adaptation or selective bleaching: these manipulations alter my relative sensitivity for red and green by less than 10 per cent, as opposed to the tenfold change possible in normals (Eisner & Macleod, 1981).

Yet in the experiment of Figure 3, my thresholds were only about twice the shown values (which are typical of practiced normal observers). More significantly, this held approximately for all tested chromatic contrasts of test and reference relative to the white surround, with no clear trend toward improved performance at high contrast. This preliminary evidence suggests that my colour-opponent neurons were overloaded to about the same degree as those of normal observers by these coloured stimuli, just as expected if colour-opponent nonlinearity is shaped during development.

Similar theoretical possibilities exist for acquired deficiencies, if these have the effect of reducing chromatic signal strength in the retina before noise is added to the

---

will undergo less nonlinear compression, with the interesting consequence that the anomalous observer will have relatively good (conceivably even super-normal) discrimination in situations where the compared stimuli each have high chromatic contrast relative to their common background. Inferior sensitivity at low chromatic contrast will be accompanied by relatively good sensitivity at high chromatic contrast.

signal at the retinal output. But since plasticity is generally thought to decrease with age it is less likely that neural nonlinearity can be optimized in response to an acquired deficit.

Instead of long-term plasticity, one could alternatively invoke short-term adaptation (Webster and Mollon, 1997) and gamut compression/expansion by context (Brown and MacLeod, 1997) as the means for “tuning” the nonlinearity of colour-deficient observers to suit their diet of impoverished chromatic contrast. On this view, acquired colour deficiencies and congenital ones should behave alike.

*Plasticity of color appearance.* There is reason to believe that the null points of the color opponent code, as well as the operating range, may be plastic. First, normal observers and anomalous observers show close agreement when asked to identify the unique phenomenally pure spectral yellow or blue that they consider to be neither reddish nor greenish, an agreement possible only if individual variation in the cone spectral sensitivities is in some way compensated postreceptorally; moreover, the slight differences between normal and anomalous unique yellow settings are roughly consistent with the idea that some post-receptor adaptation to the history of photoreceptor signals sets the opponent null appropriately in relation to the near-white average of environmental stimuli. (Pokorny & Smith, 1977). Second, among normal observers the variation in unique yellow has been found to be relatively small despite large variation in the relative contribution of L and M cones to spectral sensitivity (Brainard, Roorda, Yamauchi, Calderone, Metha, Neitz, Neitz, Williams & Jacobs, 2000, Pokorny & Smith, 1977). Paul Wise and I have obtained further evidence on this point. Twenty-four normal male observers made flicker photometric matches and unique yellow settings using CRT stimuli. The mean decimal logarithm of the relative M cone weight derived from the



flicker photometric settings was 0.045 (where agreement with the standard luminosity curve would imply a zero value), but the values for different observers differed widely, with a true standard deviation of 0.35 (after deduction of experimental variance, assessed from variability in the settings of each observer across sessions). For unique yellow settings, however, the standard deviation among observers was only 0.023, sixteen-fold less (around a mean of 0.429). The yellow setting showed no clear correlation with the observer's spectral sensitivity. The post-receptoral compensation hypothesis is consistent with such lack of correlation, but substantial correlations that it predicts between spectral yellow and other measures are not always found (Mollon & Jordan, 1997, Webster, Miyahara, Malkoc & Raker, 2000).

Any such compensation could not only promote uniformity in perception across observers (perhaps a questionable benefit), but could also improve stability in perception over the life span (Werner, 1998). It could also minimize differences between sensations from the two eyes, or different regions in the same eye. Donders (Donders, 1884, Mollon & Jordan, 1997) reports the case of a certain Dr. Sulzer, whose two eyes gave different impressions for spectral lights but were in agreement for more natural, broadband stimuli. In a more extreme instance, Peter Lennie and I (MacLeod & Lennie, 1976) studied a unilateral deuteranope, RH. Using his dichromatic eye RH saw colours ranging from orange at long wavelengths, through desaturated green, to a slightly reddish blue at short wavelengths. This is qualitatively consistent with expectation for a system designed to minimize, for each deuteranopic chromaticity, the average disagreement with the trichromatic eye. A similar suggestion was made for unilateral tritanopia by Alpern, Kitahara and Krantz (Alpern, Kitahara & Krantz, 1983). I have tried (MacLeod &

Lennie, 1976) to demonstrate a monocular compensation process experimentally by wearing goggles with a red filter over one eye and a green filter over the other, but this failed to produce an obvious perceptual asymmetry. Recent binocular experiments are more encouraging (Yamauchi, Williams, Carroll, Neitz & Neitz, 2001). Perhaps the adaptation process has difficulty treating the two eyes independently. Monet's reaction to unilateral cataract surgery (Werner, 1998) is relevant: with the operated eye things appeared too blue, with the other eye too yellow.

### **7. Colour constancy: the “anchoring problem” for natural scenes**

*Ambiguity of mean chromaticity: the anchoring problem.* The cone excitations associated with a reflecting surface depend both on the surface spectral reflectance function and on the spectral power distribution of the illuminant. The effects of the illuminant are determined by the interplay between lights and surfaces in the natural environment. If vision is to achieve constancy of apparent surface colour under changing illumination, those effects must be allowed for, by a kind of reciprocal mapping from cone excitations to colour appearance.

Fechner and Helmholtz recognized that logarithmic compression of the cone excitations—or, equivalently, sensitivity changes associated with adaptation—could play a key role in making surface colours independent of the illumination despite changes in the retinal stimulus. Fechner (1860) noted that when the sun comes out from behind a cloud, all the logs of the cone excitations go up by the same amount and the differences of the logs are invariant change. A reciprocal adjustment of sensitivity in each cone type, a kind of von Kries normalization, will effectively subtract the same number from all the

logs. This yields a representation of the scene that's invariant with illuminant intensity (Cornsweet, 1970). But in principle, the effects of changes of illuminant colour need not be removable in this simple way—a point made in many modern discussions (e.g. Brill, 1978, Worthey & Brill, 1986). The difficulty is that in general, a change of illumination scales the cone excitations by different factors for different surfaces. Conveniently, though, as noted by Foster and Nascimento (1994) the cone excitations for natural surfaces are scaled by approximately the same factor with a change of illumination, and this would allow the effect of varying illumination to be simply corrected by reciprocal adjustments of sensitivity in the different cone types. Jürgen Golz and I (in preparation) have found that in the natural scenes of Ruderman et al (1992) the log shift principle does apply quite well in the 3 dimensions of cone excitation space. So von Kries normalization for each cone type separately could in principle achieve good constancy.

But this leaves us with a problem, one that Alan Gilchrist (e.g. Gilchrist, Kossyfidis, Bonato, Agostini, Cataliotti, Li, Spehar, Annan & Economou, 1999) has called the anchoring problem: Although the resulting representation is invariant with illumination as desired, it also shows a less welcome invariance: it fails to reflect the overall chromatic cast of a scene, rendering each scene as (on average) indeterminate in colour. How, then, do we know what to normalize by? One idea, sometimes termed the “gray world assumption” is that the average chromaticity of the retinal image might be attributed entirely to the illuminant, and not to a predominant colouration of the viewed surfaces themselves (Buchsbaum, 1980). If this assumption were operative, the space-average colour of any scene would be perceptually normalized to gray. But constancy is generally not complete: neutral surfaces appear to take on a faint tint of the illuminant colour. Such

“under-constancy” is an aspect of the visual system’s response to the anchoring problem. Consideration of the statistical variation among natural illuminants and scenes will show that under-constancy is actually not a failure of constancy, but an appropriate best guess about illuminant colour, appropriately based on knowledge of relevant environmental statistics.

*An ecological rationale for under-constancy.* Under-constancy would not be expected if the gray world assumption were strictly applied. But it may instead reflect commitment to a weaker but more justifiable probabilistic version of that assumption—an assumption that the overall chromatic cast of an individual scene is most likely to be neutral, but that exceptions are frequent enough that deviations from overall neutrality of the retinal image can most plausibly be ascribed to a combination of a chromatic illuminant with a similar chromatic bias inherent in the scene itself. The statistics of natural scenes and illuminants determine what particular apportionment of image colour to the illuminant is most statistically justifiable.

Consider a set of images generated by randomly sampled natural scenes and illuminants. For each image with mean chromaticity  $r_{image}$ , let  $r_{scene}$  be the mean surface chromaticity (the mean of  $r$ ) for that external scene when observed under neutral illumination, and  $r_{source}$  be the illuminant chromaticity. Let  $dr_{scene}$ ,  $dr_{source}$  and  $dr_{image}$  represent the deviations of  $r_{scene}$ ,  $r_{source}$  and  $r_{image}$  from the mean chromaticity for all naturally occurring scenes, illuminants and images respectively. The distributions of  $dr_{scene}$  and  $dr_{source}$  will be independent; if we assume they are approximately Gaussian, then the joint distribution is bivariate Gaussian, and contours of constant likelihood in

$(dr_{scene}, dr_{source})$  are origin-centred ellipses oriented along either the  $dr_{scene}$  or the  $dr_{source}$  axis (Figure 4). The loci of constant mean image chromaticity ( $dr_{image}$ ) in  $(dr_{scene}, dr_{source})$  will be approximately the negative diagonals, reflecting the equal contributions of illuminant and mean scene chromaticity to the image mean chromaticity:  $dr_{image} = dr_{scene} + dr_{source}$ . (There is no exact and general relation between image, source and overall scene chromaticities, but the simulations of Figure 6 below support this intuitively plausible and simple relationship as an approximation for natural conditions). Consequently, for a given value of  $dr_{image}$ , the corresponding negative diagonal defines the various possible origins (or interpretations) of the image in terms of average surface chromaticity  $dr_{scene}$  and illuminant chromaticity  $dr_{source}$ .

If the image chromaticity is attributed entirely to the illuminant and the scene is assumed neutral—the gray world assumption—the inferred values of  $dr_{scene}$  and  $dr_{source}$  become simply 0 and  $dr_{image}$  respectively, and this is what is regarded as “full constancy”. But other combinations of illuminant and scene characteristics  $dr_{scene}$  and  $dr_{source}$  are actually more likely. The most likely combination is specified by the point  $(dr_{scene}^*, dr_{source}^*)$  where the negative diagonal

$$dr_{scene} + dr_{source} = dr_{image}$$

is tangent to some constant-likelihood ellipse (Figure 4). At other points on the diagonal, a lower likelihood will prevail, since likelihood decreases monotonically with distance from the origin along any radius. To determine  $dr_{scene}^*$  and  $dr_{source}^*$ , note that any ellipse of constant likelihood has an equation of the form

$$(dr_{scene}/\mathbf{S}_{scene})^2 + (dr_{source}/\mathbf{S}_{source})^2 = \text{constant}$$

where  $\mathbf{s}_{scene}$  and  $\mathbf{s}_{source}$  are the standard deviations of  $r_{scene}$  across scenes and of  $r_{source}$  across illuminants. The slope of such an ellipse is  $-(dr_{scene}/dr_{source})(\mathbf{s}_{source}/\mathbf{s}_{scene})^2$ .

Hence the tangent point satisfies the equation

$$dr_{scene} / dr_{source} = (\mathbf{s}_{scene}/\mathbf{s}_{source})^2$$

Combining the last two equations, the relation between  $r_{image}$  and the parameters  $(dr_{scene}^*, dr_{source}^*)$  that give the most likely interpretation of the image is

$$dr_{scene}^* = (dr_{image})(\mathbf{s}_{scene})^2 / (\mathbf{s}_{scene}^2 + \mathbf{s}_{source}^2)$$

$$dr_{source}^* = (dr_{image})(\mathbf{s}_{source})^2 / (\mathbf{s}_{scene}^2 + \mathbf{s}_{source}^2)$$

Equivalently,

$$dr_{source}^*/dr_{image} = (\mathbf{s}_{source})^2 / (\mathbf{s}_{scene}^2 + \mathbf{s}_{source}^2)$$

This last equation gives the ecologically optimal degree of “underconstancy”. It expresses the maximum likelihood estimate of the illuminant chromaticity as a fraction of the “gray world” estimate that attributes the image chromaticity entirely to the illuminant. That fraction is about 0.5 for the case of equality between  $\mathbf{s}_{scene}$  and  $\mathbf{s}_{source}$ ,

This last equation gives the ecologically optimal degree of “under-constancy”. It expresses the maximum likelihood estimate of the illuminant chromaticity as a fraction of the “gray world” estimate that attributes the image chromaticity entirely to the illuminant. That fraction is about 0.5 for the case of equality between  $\mathbf{s}_{scene}$  and  $\mathbf{s}_{source}$ ,

For the Ruderman et al. scenes and the Judd-MacAdam-Wyszecki daylight distribution (Wyszecki & Stiles, 1982),  $\mathbf{s}_{scene}$  tends to be less than  $\mathbf{s}_{source}$  but comparable in magnitude; for Krinov’s spectral data for the space-averages of natural terrains, scene variation is greater and may exceed illuminant variation (Brown, in press).

Environmental statistics therefore call for a considerable degree of “under-constancy”. Observers who show moderate under-constancy are closer to a statistically optimal partitioning of scene color into illuminant and surface color components than those who show full constancy. In this sense the nearly complete constancy typically found in natural conditions (Kraft & Brainard, 1999), or suggested by color appearance judgements even in severely reduced situations (Thornton & Pugh, 1983), is excessive. This conclusion has relevance for a fundamental and long-debated question about the mechanistic basis of constancy. Full or nearly full constancy is to be expected if constancy corrections are shaped by reciprocal sensitivity regulation in the retina. A sophisticated and statistically informed interpretation of the available cues would lead to less complete constancy than is observed. It is possible, however, that nearly complete constancy under natural conditions is supported in part by cues other than space-average colour of the image (Kraft & Brainard, 1999). One source of such cues, considered next, is the distribution (in colour space) of the component elements of the scene.

*Possible disambiguating cues from other image statistics.* We have seen that the anchoring problem originates from an ambiguity: the overall chromatic cast of the retinal image could reflect the predominant colour of the surfaces, or the colour of the illuminant. This ambiguity can be expressed as a problem of discrimination: a reddish scene under white light can produce the same mean stimulation as a neutral scene in red light. How can we tell those two apart? Various potential cues that have been identified and studied, e.g. mutual illumination (e.g. Bloj, 1999) and specular reflections (Lee, 1986) turn out to be at best weakly effective.

Jürgen Golz and I have found that higher order statistics of the cone excitation distribution for the surfaces within a natural scene provide a possible basis for solving the anchoring problem—even if the mean redness of a neutral scene under red light and a red room under neutral light are the same, the higher order statistics will generally be different—and we found evidence that the visual system is sensitive to that information and may use it appropriately to achieve colour constancy. In particular, the correlation between surface chromaticity and intensity within an image appears to be a useful and effective cue for constancy.

Reddish surfaces have the colour that they do because they reflect a larger proportion of long-wavelength than of short-wavelength light. This suggests that in the room lit by reddish (predominantly long-wavelength) light, the reddish surfaces will be more luminous relative to other colours within the scene owing to the greater overlap of their spectral reflectances with the spectral range in which the illuminant has its highest power. As a result, the reddish light cast on the predominantly neutral scene will not only shift the mean image chromaticity toward red but, will create a correlation between redness and luminance among the elements of the retinal image - a correlation which will be absent (or less positive) in the case of the reddish scene under neutral illumination. This intuition is confirmed by an analysis of natural images. Thus, a high luminance-redness correlation within the image is evidence that the illuminant is reddish, no matter what the mean chromaticity. By evaluating both mean and correlation, two independent quantities, an observer can estimate two unknowns – the predominant colour (in this case, the degree of redness) inherent in the objects making up the scene, and the redness of the light source that illuminates the scene. In this way, statistics of the distribution of



luminance and chromaticity within the retinal image might resolve the ambiguity encountered in considering mean chromaticity alone (Fig. 5).

*Simulation using natural scene data.* To check how well chromatic statistics in images of natural scenes can support inferences about illuminant colour, we used the hyperspectral data for the scenes of Ruderman et al. (1998). To these 12 scenes four CIE daylight illuminations (Wyszecki & Stiles, 1982, pp142-146) were applied (correlated colour temperatures 4,000K, 5,500K, 8,500K, 20,000K; with increasing correlated colour temperature the illuminants become less reddish (lower  $r$ -values) and more bluish (higher  $b$ -values)). For each illuminated scene several statistics (mean, variance, correlation, skewness) of the luminance and chromaticity values were calculated. Somewhat discouragingly at first, all statistics we evaluated except the means were almost independent of illumination. Although within each of the 12 scenes individually, the reddish surfaces became relatively more luminous when the illumination became more reddish, this increase in luminance for reddish surfaces was too weak to introduce consistently higher luminance-redness correlations for images under more reddish illumination. The correlations were indeed almost independent of illumination.

It turns out, however, that the correlation measure can nevertheless resolve the ambiguity inherent in the mean image chromaticity, because when used together with the mean image chromaticity it separates the clusters of images from different illuminations (the four diagonally oriented clusters in Fig. 6) and permits the estimation of illumination colour. Thus an image with a mean  $r$  value of 0.7 could be a slightly reddish scene under neutral light or a neutral scene under reddish light (Fig. 3). But if the visual system takes

into account the correlation between redness and luminance, it can distinguish between illumination redness and scene redness. The more positive the correlation, the greater the evidence for redness of the illumination. In the above example image A belongs with high probability to the cluster of the reddish 4,000K illuminant whereas image B is a reddish scene under more neutral 5,500K illumination.

The reason why the luminance-redness correlation can resolve the ambiguity inherent in the mean is that scene redness and illuminant redness affect the correlation differently. Under “neutral” lighting, the eye’s spectral sensitivity may favour the neutral or comparatively reddish pixels within a neutral or greenish scene, but the eye’s diminishing sensitivity at long wavelengths will discriminate against the most reddish pixels of a reddish scene, and these will therefore tend to be of low luminance. The luminance-redness correlation thus becomes more negative the more reddish the scene, as the sloped regression lines in Fig. 6 show. No comparable effect occurs if it is not the scene, but the light that is reddish. In this case, the low luminosity of reds is counterbalanced by the illuminant's greater energy at long wavelengths. Reddish scenes, but not reddish illuminants, generate images with negative luminance-redness correlations. The ambiguity inherent in mean image chromaticity can thus be resolved.

*Is this cue exploited? Experimental tests.* To find out whether human vision exploits scene statistics in this way, we used stimuli of a type introduced by R.Mausfeld and J. Andres (Mausfeld, 1998, Mausfeld & Andres, 1999)) which make it possible to vary independently various statistics (means, variances, correlations) of the distribution of colour and lightness within the display. A circular test field was surrounded by random

patterns of overlapping circles; these were of a fixed diameter but varied in colour and luminance to a degree typical of the natural scenes of Ruderman et al. We asked subjects to adjust the colour of the test field so that it appeared neutral grey.

In our main experiment we varied the luminance-redness correlation for the elements surrounding the test field (by introducing a linear dependence 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, -0.8, 0.0, 0.8 or 1.0). 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 grey should be the same for all 5 conditions, since the space-averaged chromaticities of the surrounds were the same. The surrounds would then be functionally equivalent with respect to the perceived colour of the centre test spot.

For conditions with higher correlation between redness and luminance, a more reddish chromaticity was required to make the test field subjectively achromatic (Figure 7). The data for eight of ten subjects tested were quantitatively similar and individually statistically significant ( $p < 0.001$ , linear trend test). When the correlation between redness and luminance was positive, subjects selected a physically more reddish (higher  $r$ ) test field as neutral grey. Since higher  $r$  values are associated with redder illumination of a physically 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 grey when a correspondingly reddish light stimulus is received from the test field.

Though small, the effect of the luminance-redness correlation is not inconsequential; for comparison, a just noticeable difference in  $r$  is only about 0.001 (Figure 3).

These results are consistent with other recent work showing that the correction subserving colour constancy is not governed merely by the space-averaged chromaticity (Bäumli, 1994, Brown & MacLeod, 1997, Jenness & Shevell, 1995, Kraft & Brainard, 1999, Mausfeld & Andres, 1999, Webster & Mollon, 1997); some current colour constancy algorithms provide general frameworks for effects of this sort (Brainard & Freeman, 1997, D'Zmura, Iverson & Singer, 1995, Forsyth, 1990).

*Is the luminance-chromaticity correlation cue given appropriate weight?* 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 of natural scenes under different illuminations, we calculated a maximum likelihood estimate for the chromaticity of the illumination given the mean and correlation value of an image. An optimal observer, adopting this maximum likelihood estimate, would have reacted to the stimuli of our experiment much as our observers did, giving only about 30% more weight to the correlation cue than is implied by the straight line fit to the data of Figure 7. The size of the observed effect of this statistic on colour perception is therefore roughly consistent with optimal computation.

### **Acknowledgements**

Much of the material of sections 2-5 and section 7 was developed during the author's participation in the project, "Perception and the Role of Evolutionary Internalized Regularities of the Physical World." organized by R. Mausfeld, D. Heyer and H. Hecht at the Zentrum für interdisziplinäre Forschung, Bielefeld, 1995-96. Thanks are due to all participants in that project, including collaborators Tassilo von der Twer (Physics, Wuppertal) and Richard O. Brown (Exploratorium, San Francisco), and also to my other

collaborators in this work, Anya Leonova (now at Chicago) and Jürgen Golz (Kiel). Dan Ruderman kindly made available the hyperspectral data of Ruderman, Cronin, and Chiao on natural scenes. Barry Lee, John Mollon and Mike Webster provided helpful comments. Jürgen Golz's work at UCSD was supported by the German-American Fulbright Commission. All the work was supported by NIH Grant EY01711.

## **References**

- Allman, J., & Zucker, S. (1990). Cytochrome oxidase and functional coding in primate striate cortex: a hypothesis. *Cold Spring Harb Symp Quant Biol*, *55*, 979-982.
- Alpern, M., Kitahara, K., & Krantz, D.H. (1983). Perception of colour in unilateral tritanopia. *J Physiol*, *335*, 683-697.
- Barlow, H.B. (1965). Optic nerve impulses and Weber's law. *Cold Spring Harbor Symposium on Quantitative Biology*, *30*, 539-546.
- Bäumli, K.-H. (1994). Color appearance: effects of illuminant changes under different surface collections. *J. Opt. Soc. Am. A*, *11* (2), 531-542.
- Bialek, W., & Rieke, F. (1992). Reliability and information transmission in spiking neurons. *Trends Neurosci*, *15* (11), 428-434.
- Bloj, M., Kersten, D. & Hurlbert, A. C (1999). Perception of three-dimensional shape influences colour perception through mutual illumination. *Nature*, *402*, 877-879.
- Boynton, R.M. (1980). Design for an Eye. In: D. McFadden (Ed.) *Neural mechanisms in behavior* (Berlin: Springer-Verlag).
- Boynton, R.M., Hayhoe, M.M., & MacLeod, D.I.A. (1977). The gap effect: chromatic and achromatic visual discrimination as affected by field separation. *Optica Acta*, *24*, 159-177.

- Brainard, D.H., & Freeman, W.T. (1997). Bayesian color constancy. *Journal of the Optical Society of America a. Optics and Image Science*, 14 (7), 1393-1411.
- Brainard, D.H., Roorda, A., Yamauchi, Y., Calderone, J.B., Metha, A., Neitz, M., Neitz, J., Williams, D.R., & Jacobs, G.H. (2000). Functional consequences of the relative numbers of L and M cones. *J Opt Soc Am A Opt Image Sci Vis*, 17 (3), 607-614.
- Brill, M.H. (1978). A device performing illuminant-invariant assessment of chromatic relations. *Journal of Theoretical Biology*, 71 (3), 473-478.
- Brown, R.O. (1994). The world is not grey. *Invest. Ophthalmol. Vis. Sci. (Suppl.)*, 35/4, 2165.
- Brown, R.O. (in press). Backgrounds and illuminants: the yin and yang of color constancy. In: R. Mausfeld (Ed.) *Colour Perception: From Light to Object* Oxford: Oxford University Press.
- Brown, R.O., & MacLeod, D.I.A. (1997). Color appearance depends on the variance of surround colors. *Current Biology*, 7 (11), 844-849.
- Buchsbaum, G. (1980). A Spatial Processor Model for Object Colour Perception. *Journal of The Franklin Institute*, 310 (1), 1-26.
- Buchsbaum, G., & Gottschalk, A. (1983). Trichromacy, opponent colours coding and optimum colour information transmission in the retina. *Proc. Roy. Soc. Lond. B*, 220, 89-113.
- Chaparro, A., Stromeyer, C.F., Huang, E.P., Kronauer, R.E., & Eskew, R.T.J. (1993). Colour is what the eye sees best. *Nature*, 361, 348-350.
- Cornsweet, T.N. (1970). Visual perception. (pp. xiii, 475). New York,: Academic Press.

- Dacey, D.M., Lee, B.B., Stafford, D.K., Pokorny, J., & Smith, V.C. (1996). Horizontal cells of the primate retina: cone specificity without spectral opponency. *Science*, *271*, 656-659.
- Derrington, A.M., Krauskopf, J., & Lennie, P. (1984). Chromatic mechanisms in lateral geniculate nucleus of macaque. *Journal of Physiology (London)*, *357*, 241-265.
- DeValois, R.L., & DeValois, K.K. (1975). Neural coding of color. In: E.D. Carterette, & M.P. Friedman (Eds.), *Handbook of Perception*, 5 (pp. 117-166). New York: Academic Press.
- Donders, F.C. (1884). Farbgleichungen. *Archiv fur Anatomie und Physiologie. Physiologische Abt. Leipzig*, *8*, 518-552.
- 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.
- Eisner, A., & Macleod, D.I. (1981). Flicker photometric study of chromatic adaption: selective suppression of cone inputs by colored backgrounds. *J Opt Soc Am*, *71* (6), 705-717.
- Eisner, A., & MacLeod, D.I.A. (1980). Blue-sensitive cones do not contribute to luminance. *J. Opt. Soc. Am.*, *70*, 121-122.
- Fechner, G.T. (1860). *Elemente der Psychophysik*. (p. 2 v.). Leipzig,: Breitkopf und Härtel.
- Forsyth, D.A. (1990). A Novel Algorithm for Color Constancy. *International Journal of Computer Vision*, *5* (1), 5-36.

- 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), 115-121.
- Friele, L.F.C. (1961). Analysis of the Brown and Brown-MacAdam colour discrimination data. *Farbe*, 10, 193-224.
- Fukurotani, K. (1982). Color information coding of horizontal cell responses in fish retina. *Col. Res. Appl.*, 7, 146-148.
- Gilchrist, A.L., Kossyfidis, C., Bonato, F., Agostini, T., Cataliotti, J., Li, X., Spehar, B., Annan, V., & Economou, E. (1999). An anchoring theory of lightness perception. *Psychol Rev.*, 106, 795-834.
- Gregory, R.L., and P. Heard (1979). Border locking and the Cafe Wall illusion. *Perception*, 8, 365-380.
- He, S., & MacLeod, D.I.A. (1997). Contrast-modulation flicker: dynamics and spatial resolution of the light adaptation process. *Vision Res.*, 38, 985-1000.
- Helmholtz, H.v. (1896). *Handbuch der Physiologischen Optik*. (Hamburg: Voss.
- Hurvich, L.M. (1972). Color vision deficiencies. In: *Handbook of Sensory Physiology*, VII/4 (pp. 582-624). Berlin: Spriger-Verlag.
- Jenness, J.W., & Shevell, S.K. (1995). Color appearance with sparse chromatic context. *Vision Research*, 35 (6), 797-805.
- Kaplan, E., and R. M. Shapley (1986). The primate retina contains two types of ganglion cells, with high and low contrast sensitivity. *Proceedings of the National Academy of Sciences (U.S.A.)*, 83, 2755-2757.



- Kraft, J.M., & Brainard, D.H. (1999). Mechanisms of color constancy under nearly natural viewing. *Proceedings of the National Academy of Sciences of the United States of America*, 96 (1), 307-312.
- Krauskopf, J., & Gegenfurtner, K. (1992). Color discrimination and adaptation. *Vision Res*, 32 (11), 2165-2175.
- Land, E.H. (1959). Experiments in color vision. *Scientific American*, 200, 84.
- Laughlin, S.B. (1983). Matching coding to scenes to enhance efficiency. In: O.J. Braddick, & A.C. Sleigh (Eds.), *Biological processing of images* (pp. 42-52). Berlin: Springer Verlag.
- Le Grand, Y. (1949). Les seuils différentiels de couleur dans la théorie de Young. *Revue d'Optique*, 28, 261-278.
- Lee, B.B., Pokorny, J., Smith, V.C., Martin, P.R., & Valberg, A. (1990). Luminance and chromatic modulation sensitivity of macaque ganglion cells and human observers. *J Opt Soc Am [A]*, 7 (12), 2223-2236.
- Lee, B.B., Wehrhahn, C., Westheimer, G., & Kremers, J. (1993). Macaque ganglion cell responses to stimuli that elicit hyperacuity in man: detection of small displacements. *J Neurosci*, 13 (3), 1001-1009.
- Lee, H.C. (1986). Method for computing the scene-illuminant chromaticity from specular highlights. *Journal of the Optical Society of America a. Optics and Image Science*, 3 (10), 1694-1699.
- LeGrand, Y. (1949). Les seuils différentiels de couleur dans la théorie de Young. *Revue d'Optique*, 28, 261-278.

Lennie, P., Pokorny, J., & Smith, V.C. (1993). Luminance. *J. Opt. Soc. Am. A*, *10*, 1283-1293.

Livingstone, M.S., & Hubel, D.H. (1987). Psychophysical evidence for separate channels for the perception of form, color, movement, and depth. *Journal of Neuroscience*, *7* (11), 3416-3468.

MacAdam, D.L. (1942). Visual sensitivities to color differences in daylight. *J. Opt. Soc. Am.*, *32*, 247-273.

MacLeod, D.I., & Lennie, P. (1976). Red-green blindness confined to one eye. *Vision Res*, *16* (7), 691-702.

MacLeod, D.I.A., Williams, D.R., & Makous, W. (1992). A visual nonlinearity fed by single cones. *Vision Research*, *32*, 347-363.

Marr, D. (1974). The computation of lightness by the primate retina. *Vision Research*, *14* (12), 1377-1388.

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.

Mausfeld, R., & Andres, J. (1999). Detecting the presence of a 'non-normal' illumination: cues based on second-order statistics of colour codes. *Perception*, *27* (Supplement), 31-32.

Mollon, J.D. (1982). Color vision. *Annual Review of Psychology*, *33* (2), 41-85.

Mollon, J.D., & Jordan, G. (1997). On the nature of unique hues. In: D. Carden (Ed.) *John Dalton's Colour Vision Legacy* (pp. 381-392). London: Taylor and Francis.

- Morgan, M.J., Adam, A., & Mollon, J.D. (1992). Dichromats detect colour-camouflaged objects that are not detected by trichromats. *Proc. R. Soc. Lond. B.*, *248*, 291-295.
- Mollon, J.D., & Jordan, G. (1997). On the nature of unique hues. In: C. Dickinson, I. Murray, & D. Carden (Eds.), *John Dalton's Colour Vision Legacy* (pp. 381-392). London: Taylor and Francis.
- Osorio, D., & Bossomaier, T.R.J. (1992). Human Cone-Pigment Spectral Sensitivities and the Reflectances of Natural Surfaces. *Biological Cybernetics*, *67* (3), 217-222.
- Párraga, C.A., Brelstaff, G., Troscianko, T., & Moorehead, I.R. (1998). Color and luminance information in natural scenes. *J Opt Soc Am a Opt Image Sci Vis*, *15* (3), 563-569.
- Pickford, R.W. (1958). A review of some problems of colour vision and colour blindness. *Advancement of Science*, *15*, 104-117.
- Pokorny, J., & Smith, V.C. (1977). Evaluation of single-pigment shift model of anomalous trichromacy. *J Opt Soc Am*, *67* (9), 1196-1209.
- Pokorny, J., & Smith, V.C. (1997). Psychophysical signatures associated with magnocellular and parvocellular pathway contrast gain. *J Opt Soc Am A*, *14* (9), 2477-2486.
- Pokorny, J., Smith, V.C., & Verriest, G. (1979). Congenital Color Defects. In: J. Pokorny, V.C. Smith, G. Verriest, & A.J. Pinckers (Eds.), *Congenital and Acquired Color Vision Defects* (pp. 183-241). New York: Grune and Stratton.
- Pugh, E.N., Jr., & Mollon, J.D. (1979). A theory of the pi1 and pi3 color mechanisms of Stiles. *Vision Res*, *19* (3), 293-312.

Regan, B.C., Julliot, C., Simmen, B., Viénot, F., Charles-Dominique, P., & Mollon, J.D. (2001). Fruits, foliage and the evolution of primate colour vision. *Philos Trans R Soc Lond B Biol Sci*, 356, 229-283.

Regan, B.C., & Mollon, J.D. (1997). The relative salience of the cardinal axes of colour space in normal and anomalous trichromats. In: C.R. Cavonius (Ed.) *Colour Vision Deficiencies XIII* (pp. 261-270). Dordrecht: Kluwer.

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.

Sherman, P.D. (1981). *Colour Vision in the Nineteenth Century*. (Bristol: Adam Hilger.

Switkes, E., & Crognale, M.A. (1999). Comparison of color and luminance contrast: apples versus oranges? *Vision Res*, 39, 1823-1831.

Thornton, J.E., & Pugh, E.N., Jr. (1983). Red/Green color opponency at detection threshold. *Science*, 219 (4581), 191-193.

Verriest, G. (1960). Contribution a l'étude de la corrélation entre la quotient d'anomalie et l'étendue des confusions colorimétriques dan les systèmes trichromatiques anormaux. *Revue d'Optique*, 39, 467-471.

von der Twert, T., & MacLeod, D.I.A. (2001). Optimal nonlinear codes for the perception of natural colours. *Network: Computation in Neural Systems*, 12, 395-407.

Wachtler, T., Wehrhahn, C., & Lee, B.B. (1996). A simple model of human foveal ganglion cell responses to hyperacuity stimuli. *J Comput Neurosci*, 3 (1), 73-82.

Webster, M.A., Miyahara, E., Malkoc, G., & Raker, V.E. (2000). Variations in normal color vision. I. Cone-opponent axes. *J Opt Soc Am A Opt Image Sci Vis*, 17 (9), 1535-1544.

- Webster, M.A., & Mollon, J.D. (1997). Adaptation and the color statistics of natural images. *Vision Research*, 37 (23), 3283-3298.
- Werner, J.S. (1998). Aging through the Eyes of Monet. In: W. Backhaus, R. Kliegel, & J. Werner (Eds.), *Color Vision* (pp. 1-41). Berlin/New York: De Gruyter.
- Willis, M.P., & Farnsworth, D. (1952). Comparative Evaluation of Anomaloscopes. (Bureau of Medicine and Surgery, U.S. Navy.
- 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), 1708-1712.
- Wyszecki, G., & Stiles, W.S. (1982). Color science: concepts and methods, quantitative data and formulae. (New York: John Wiley & Sons.
- Yamauchi, Y., Williams, D.R., Carroll, J., Neitz, J., & Neitz, M. (2001). Chromatic adaptation can cause long term shifts in color appearance that arise in binocular visual pathways. *Invest. Ophthalmol. Vis. Sci.*, 42 (Suppl.), #3873.
- Zaidi, Q. (1997). Decorrelation of L- and M-cone signals. *J Opt Soc Am A*, 3430-3431.

## **LEGENDS**

Figure 1. Distribution of Brown's haphazard sample of natural colours in the (L,M) plane of cone excitation space. Linear co-ordinates show a clustering near the origin. Slicing pattern at left achievable by independent L and M cone signals (for which the constant response are vertical and horizontal lines respectively). Slicing pattern at

right uses a luminance signal (lines of negative slope) and a red-green opponent (L-M) signals (lines of positive slope).

Figure 2. Crosses: frequency distribution of  $\log_{10}(b)$  for Ruderman et al.'s set of natural colours;  $b$  specifies S cone excitation per unit luminance, i.e.  $b = S/(L+M)$ . Whites and greens are near the middle of the distribution, with equal energy white at  $\log_{10}(b)=0$ . To the right of that lie bluish colours; to the left, generally yellowish or reddish ones. Candidate input-outputs functions: pleistochrome (circles), compared with linear(squares) and stepwise alternatives.

Figure 3. Circles, with straight lines fit, show  $\Delta r$ , the difference in  $r = L/(L+M)$  just sufficient for 84% correct discrimination between isoluminant, non-adjacent test and standard stimuli, embedded in a common white surround, as a function of value of  $r$  for the standard stimulus. The surround was an equal energy white, for which  $r = 0.70$ ; hence abscissa values of .707 and .693 correspond to a 1% L cone contrast between standard and background.

Dashed lines show predictions for the extreme cases of linear encoding (horizontal dashed line: threshold is independent of reference colour) and step function encoding (steep dashed line: discrimination is only categorical). The inset illustrates these encoding schemes.

Figure 4. Ecological basis for underconstancy. The horizontal axis represents illuminant redness and the vertical axis represents the scene-averaged inherent redness of

illuminated surfaces, each expressed as a deviation from its average for natural environments ( $dr_{source}$ , horizontal axis and  $dr_{scene}$ , vertical axis). If these have approximately Gaussian and independent distributions, the contours of constant likelihood in their joint distribution will be approximately ellipses, as shown. A given reddish image is consistent with the various pairings of scene and illuminant redness traced out by the indicated line of slope  $-1$ . The most likely interpretation of the image is that both scene and illuminant are slightly reddish rather than that the redness originates in either one alone.

Figure 5. (a) The ambiguity of the mean image chromaticity for estimating the illuminant and how it could be resolved by higher order statistics. A given mean image chromaticity could result either from a reddish scene under white illumination or from a white scene under reddish illumination. (b) Taking higher order statistics into account can resolve this ambiguity.

Figure 6. Luminance-redness correlation and mean redness for images formed by natural scenes viewed under 4 different illuminations. The vertical coordinate of each data point represents the correlation between pixel redness ( $r$ ) and pixel luminance within an image of a particular scene under a particular illuminant. This correlation is plotted here against space average image redness (the mean of  $r$  averaged over the entire image). The four clusters each show data for 12 natural scenes under one of four different CIE daylight illuminants (■: colour temperature 4,000K, ○: 5,500K, ▼: 8,500K, ◇:

20,000K). For a given scene, the luminance-redness correlation is almost independent of illumination, but within each illuminant cluster the correlation is 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 in a way that is not possible using mean image redness alone.


Figure 7 Dependence of centre test spot settings on the luminance-redness correlation within the surround. Average results for all subjects (circles), together with the best fitting straight line. Error bars for the experimental results are  $\pm$  one standard error for variability between subjects (subject sampling error); variability of individual subjects settings is much smaller.

Table 1. L cone contrasts (for luminance and for isoluminant red-green direction) and S cone cone contrast (for isoluminant blue-yellow direction) for which colour discrimination thresholds are doubled (“Psychophysical”) column. “Neural” column shows rough estimates from physiological data of the contrast at which the gradient contrast-response function relevant neurons falls to half maximum.

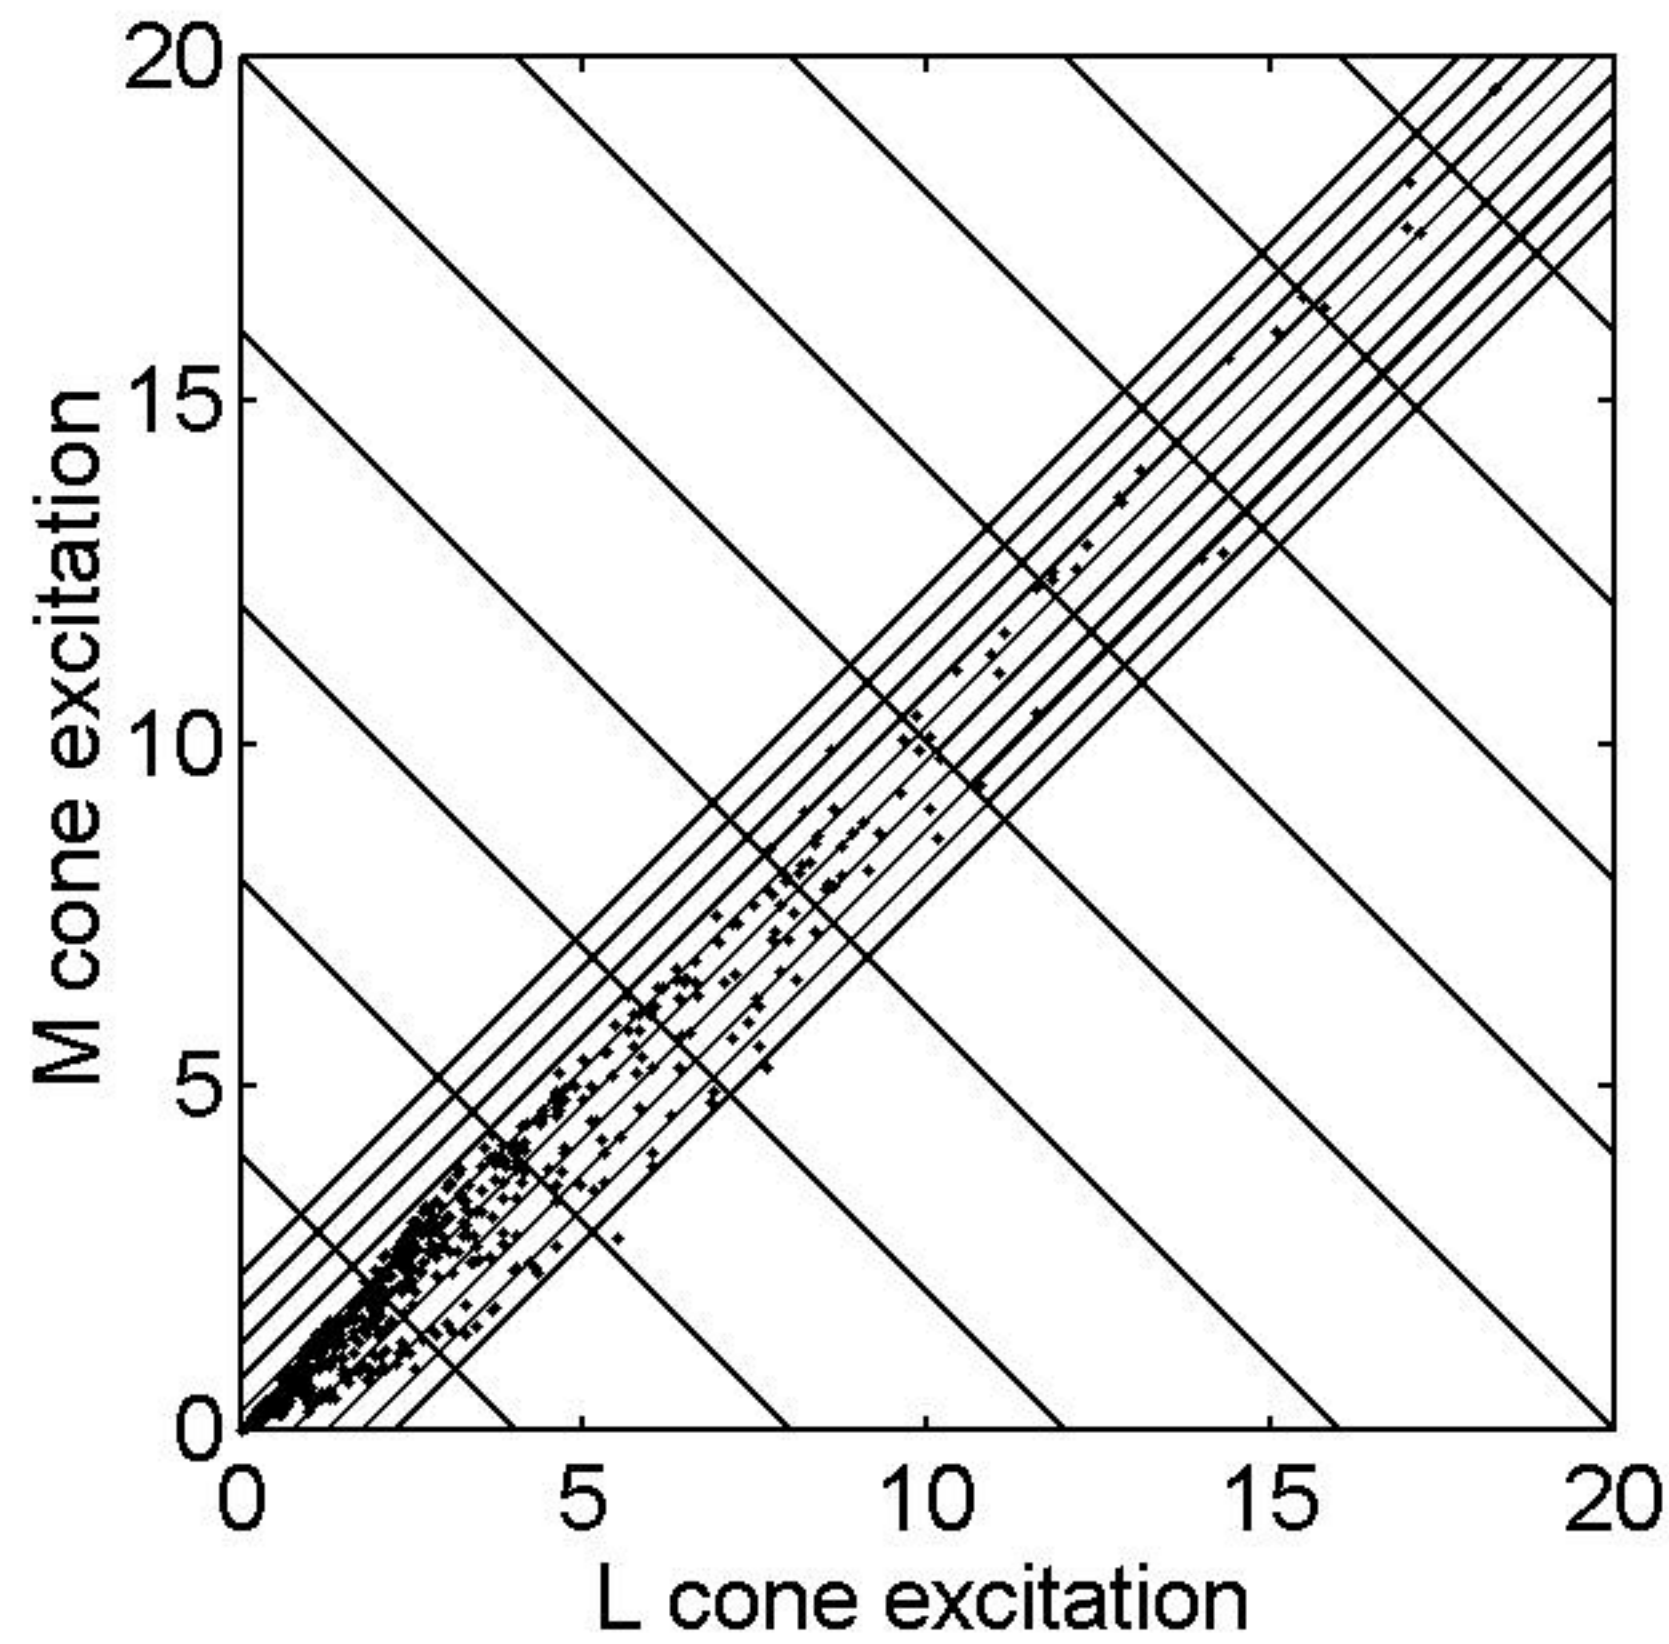


&lt;TABLE 1 of MacLeod.doc &gt;&gt;


<b>“Half-gradient” Contrasts</b>	Psychophysical (Leonova and MacLeod)	Neural (Lee, and others)
Luminance	20%	5% for contrast (M)
Blue/yellow	20%	??
Red/green	2%	10% to 20% for contrast (P)




(a)

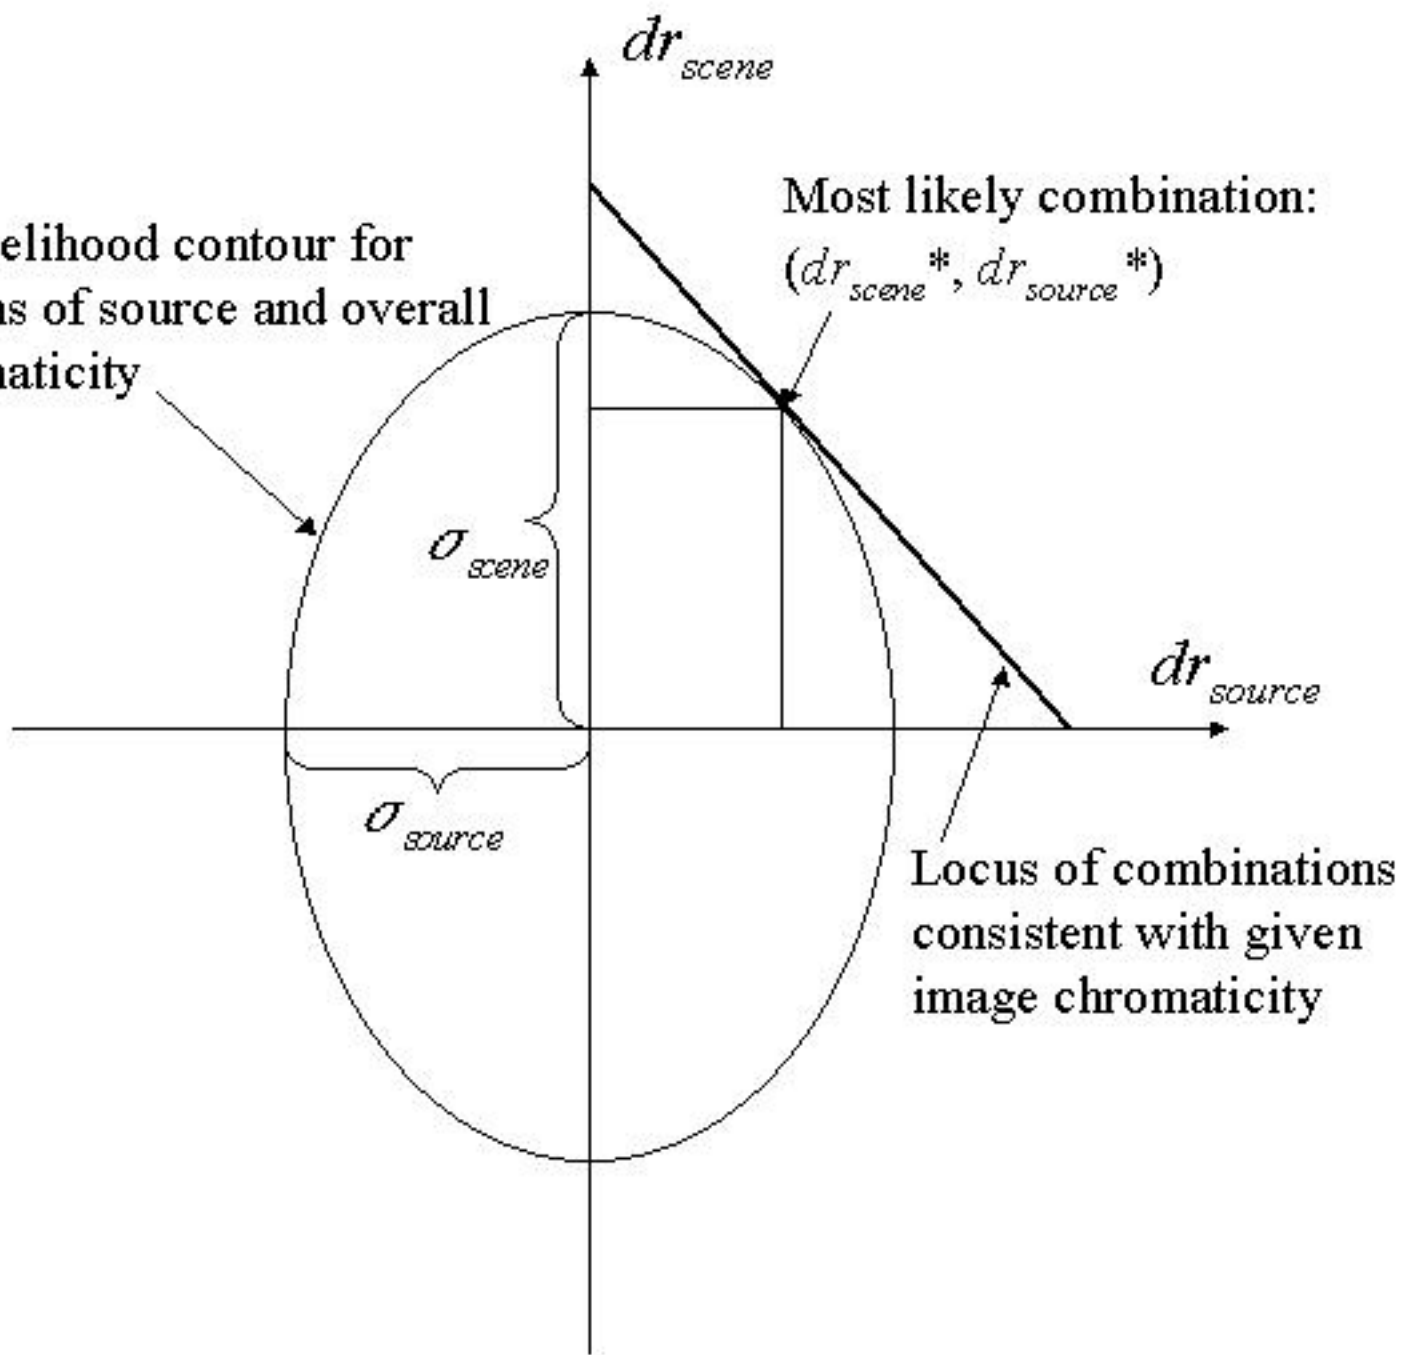


(b)

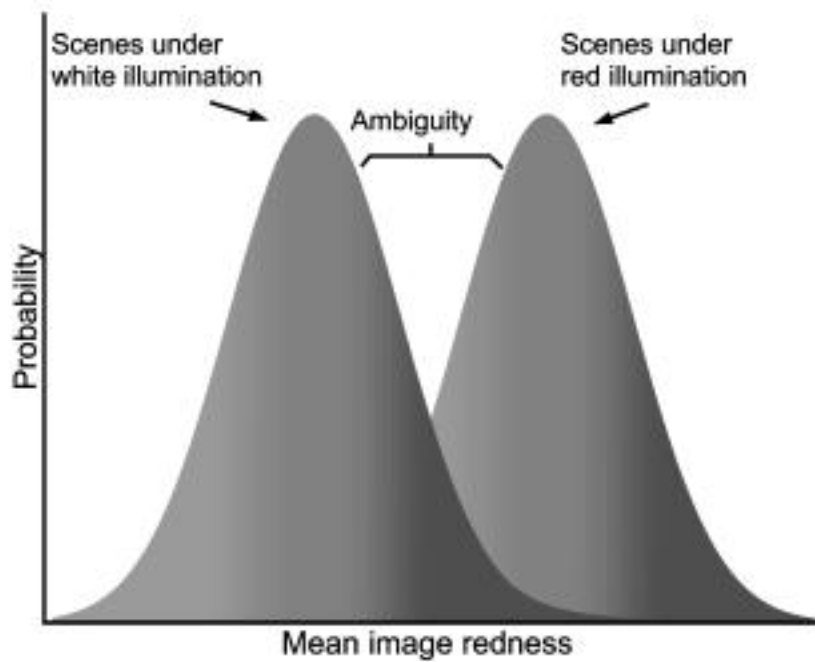




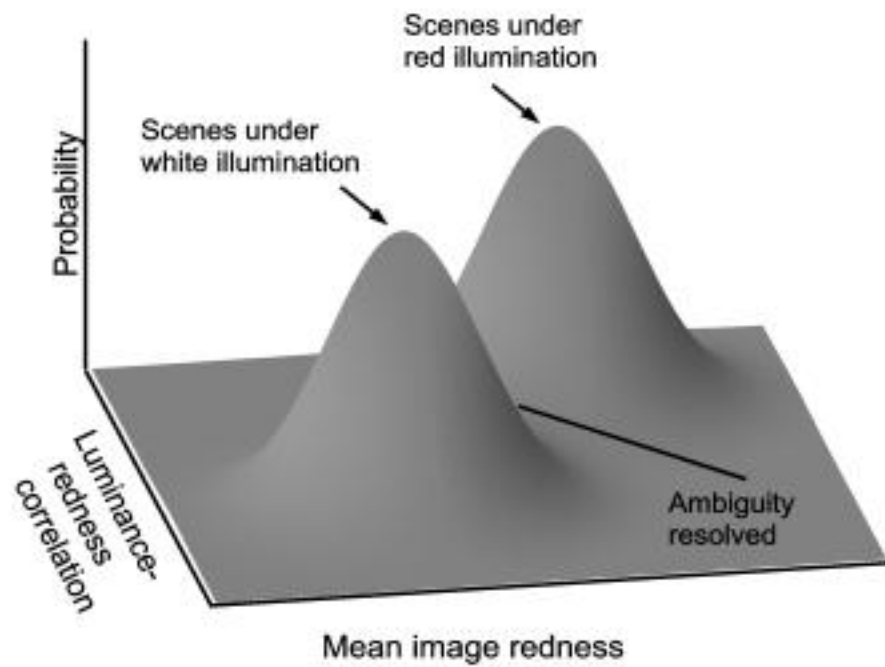
Constant likelihood contour for combinations of source and overall scene chromaticity




a



b



Luminance-redness correlation



Mean image redness

