THE PLEISTOCHROME: OPTIMAL OPPONENT CODES FOR NATURAL COLOURS

Donald I. A. MacLeod

University of California at San Diego

La Jolla, CA 92093-0109

USA

Tassilo von der Twer

Bergische Universitat Wuppertal

D-42097 Wuppertal

Germany

Draft submitted for publication in: "Color Perception: from Light to Object", edited by R. Mausfeld and D. Heyer (OUP, forthcoming)

1. Introduction

Rich though it is, the perceptual world of conscious experience is far more impoverished, in terms of sheer information content, than either the external reality of which it is a representation, or the proximal stimulus from which it is constructed. Many distinctions that are present in the stimulus fail to register in perception. The domain of color vision provides two clear examples of this. First, the initial encoding of color by the human retina is only three-dimensional, with the result that very different spectral energy distributions may be absolutely equivalent visually. Second, because discrimination is limited, even a stimulus difference of a kind that could lead to a perceptible difference in color or brightness will escape notice if it is sufficiently small in magnitude. In this chapter we address this second limitation on color vision, by analysing the discrimination of color and brightness within a framework that is both mechanistic and ecological.

Discrimination is a primitive perceptual accomplishment that lends itself to a mechanistic analysis informed by neurophysiology. Two stimuli that are physically different will be perceptually discriminable only if their neural representations are reliably different. The earliest such representation in vision is the set of excitations of the photoreceptor cells in the retina. But a mechanistic analysis cannot be based on that alone: as we will see, postreceptoral recoding radically alters the neural representation of color, even within the retina, and information may be lost in this recoding. No matter which brain loci form the immediate substrates of visual experience, any distinctions that are lost in the retinal output can hardly be restored in the brain or in conscious experience [Brindley, 1970 #60]. Which stimulus distinctions are lost in this way, and which are retained, will depend on the nature of the neural code at the retinal output, as we will see below. This is a purely mechanistic, or neurophysiological issue.

The visual system's neural code for color can be regarded as a choice made either during evolution, or during individual development under the guidance of genetically allowed plasticity. In either case, its selection poses a problem of design. Here we encounter the ecological aspect of the problem. The code should be one that reliably represents important stimulus differences, obliterating only the less important ones. Thus there is no need for precision in the neural representation of stimuli that never occur in the natural

environment. Our argument in this chapter is that the postreceptoral neural code for color is nicely adapted to the representation of naturally common stimuli.

One of the simplest statistical characterizations of a sensory environment is the probability distribution of input values. In this chapter, we consider the implications of that distribution for perceptual coding. The distribution is commonly a peaked one: in the case of colour for example, natural colours are usually nearly neutral (delivering roughly comparable intensities of stimulation to each of the three types of cone photoreceptor) rather than vividly saturated (delivering very different intensities of stimulation). In the case of visual motion, the angular velocities of viewed objects relative to the observer are typically small or zero, and for any given direction of motion the larger velocities, positive or negative, are progressively less frequent.

Turning from the stimulus to its neural representation, we commonly find an opponent code. The clearest instance of this is in the representation of colour, where individual neurons at stages following the photoreceptor stage (Derrington, Krauskopf, and Lennie 1984; DeValois and DeValois 1975) are excited by certain parts of the spectrum and inhibited by others. This happens because these neurons are excited by one or two types of spectrally selective photoreceptor and inhibited by others. A somewhat paradoxical consequence of this encoding scheme is that the postreceptoral neurons are poorly responsive to the physiologically and phenomenally neutral stimuli that are most abundant in the environment. Likewise in the case of motion detection, directionally selective neurons respond poorly to static or nearly static stimuli, with inhibitory or zero response for motion in directions opposed to the preferred direction. Thus the cases of motion and colour both exemplify what we will call a '*split range*' code, where an input continuum such as colour, or (signed) input velocity, is divided at a physiological null point (near-white, or zero velocity), and where separate neurons respond to inputs on opposite sides of that null point. In this chapter we ask: Why is this nonlinear encoding scheme a good one? We first demonstrate theoretically a quantitative connection between the statistics of environmental inputs and the split range code evolved by the visual system. Namely, by adopting a split range code, and representing opposite segments of the input range by neurons that each have rectifying and compressive response nonlinearity, the visual system maximizes the average precision in its representation of natural inputs in the presence of neural noise introduced at the output. We then compare the optimal form of the split range nonlinearity with experimental estimates

from psychophysics and from neurophysiology, finding fair agreement. Our discussion is focused on colour vision, although the theoretical arguments are quite general. We therefore begin with a rough characterization of the distribution of natural colours.

2. The distribution of surface colours

Because the cone spectral sensitivities are broad, with substantial overlap, the ratio of the sensitivity of the long wavelength-sensitive cone photoreceptors (L cones) to midspectral (M) cone sensitivity varies only by 20:1 across the spectrum (Stockman, MacLeod, and Johnson 1993), and L and M cone excitations are strongly correlated. For equal energy spectral lights, both L and M cones are most strongly excited in the yellow-green part of the spectrum, and for spectral lights equally spaced in wavelength from 400 to 700 nm their excitations show a correlation of 0.84. As Fukurotani (Fukurotani 1982) and Buchsbaum and Gottschalk (Buchsbaum and Gottschalk 1983) have noted, this rather high correlation for spectral colours means that the L and M cones measure almost the same thing. The difference between their excitations, on which perception of colour depends, is very small. Moreover, this problem is enormously exacerbated by the characteristically broad spectral energy distributions of natural stimuli. For natural stimuli, the ratio of L to M cone sensitivity seldom even approaches the limiting values that can be attained by spectral lights, but both L and M cone excitations vary together with varying surface lightness. The correlation between the cone excitations is correspondingly higher for natural colours than for spectral lights. This is clearly apparent in each of two sets of measurements on natural colours (Brown 1994; Ruderman, Cronin, and Chiao 1998), on which we have mainly relied in our analysis. Fig. 1 shows results for a set of 574 natural spectral reflectance functions measured in San Diego by Richard O. Brown (Brown 1994; Brown, this volume). The sample included flowers, fruits, leaves, barks, and earths, with a few samples of water and sky. For these 574 samples, the correlation between L and M cone excitation is 0.985 (Fig. 1). Since it seems impossible to define a representative sample, Brown made no attempt to select in any systematic way but measured various things that happened to catch his eye; vivid colours are accordingly over-represented. Ruderman, Cronin and Chiao (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 than in Fig. 1 (0.9983).

Correspondingly, the variation among the natural colors of Figure 1 is far less in the red/green direction than in the luminance direction, even when the effects of variations in the daylight illuminants (which Brown has noted (Brown 1994; this volume) are particularly large in the luminance direction) are normalized out. To normalize for illumination variation, each surface was characterized by its spectral reflectance relative to a full reflectance white standard measured under the same illumination (Brown) or in the same scene (Ruderman et al.). We derived cone excitations from these reflectances by integrating their crossproducts with the cone sensitivities of Stockman et al (Stockman, MacLeod, and Johnson 1993), either without or with a weighting factor for the spectral energy distribution of typical (D65) daylight. The resulting numbers are the cone excitations produced by each surface viewed under, respectively, a standard equal energy white illuminant or a D65 illuminant. Surface luminance is given by the summed excitations of L and M cones, which we denote here simply by L and M (Eisner and MacLeod 1980; Lennie, Pokorny, and Smith 1993). The standard deviation of log10(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. The purely chromatic variations are conveniently indexed by two axes: r = L/(L+M) forms the (roughly speaking) 'red/green' axis of a photoreceptor-based chromaticity diagram (Luther 1927; MacLeod and Boynton 1979) and is proportional to L cone excitation per unit luminance. In Fig. 1 it is nearly proportional to the angle, from vertical, of the line connecting a colour point to the origin. As the high correlation between L and M implies, the standard deviation of r is far smaller than that for luminance (only 7.5%, or 0.03 in the decimal logarithm in Fig. 1, 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 more than 10 times that for r, and about as high as the one for luminance: 0.39 in log10(b) or a factor of about 2.5 for Brown's data, or 0.24 in log10(b) for the data of Ruderman et al. This greater variability for b than for r arises partly because natural surface reflectances have more variation at short wavelengths than at long, but partly it arises because the S cone spectral sensitivity curve is spectrally remote: it is displaced some 6 times further from the L and M sensitivity curves than those are from each other, with the result that the ratio of S sensitivity to L or M sensitivity varies more than a million-fold across the spectrum (Boynton 1980; McMahon and MacLeod 1998; Stockman, MacLeod, and Johnson 1993)

3. Colour discrimination as a slicing of colour space

We assume initially that the goal of the encoding of colour and lightness is to characterize or specify surfaces in terms of colour and lightness. While this may appear tautologous, there are other possibilities (Boynton 1980; Morgan, Adam, and Mollon 1992)—for instance, the goal of detecting all object boundaries has rather different requirements, which we briefly consider in §9 below. Roughly speaking, characterization of any given surface is made most precise by making the number of distinguishable colour-lightness combinations as large as possible. The number of distinctions made by a single neuron is limited by restrictions on its firing rate (from zero to some maximum) and by the random fluctuations in the firing rate. As an aid to intuition, we find it useful to imagine a discrete set of some number N of progressively increasing firing rates, ranging from zero to a maximum rate with each one just reliably different from the last, as defining the number N of distinct signals that the neuron can generate: somewhat more precisely, such a neuron could be regarded as classifying any stimulus into one of N classes, on the basis of the firing rate that the stimulus evokes. A neuron fed by L cones alone will distinguish colour stimuli on the basis of L cone excitation alone, and planes of constant L cone excitation in colour space will bound the N colour classes that it can distinguish. On this view the neuron performs a slicing of colour space, separating it into N distinguishable colours; in Fig. 1 an L cone-driven neuron would make vertical slices, while a purely M cone-driven neuron would make horizontal ones. The combination of an L-driven neuron with an M-driven one will slice both horizontally and vertically, creating a grid of N^2 distinguishable colours that appear as squares in the (L,M) plane of Fig. 1 (and also in the 3-dimensional (L,M,S) colour space, the S coordinate being irrelevant). Yet as Fig. 1 shows, only a small number of these potentially distinguishable squares are actually occupied by natural colours. In the case illustrated, N=10 and only 23 of a possible 100 squares are occupied.

Although the retinal ganglion cells of the optic nerve, as well as visual neurons in the brain, are fed by multiple cone types, the planar slicing analogy is straightforwardly applicable to them also, since the signals of these neurons depend (albeit through a nonlinear response function) on a weighted but (approximately, and with the proviso that the state of adaptation does not vary) linear combination of cone excitations (Derrington, Krauskopf, and Lennie 1984). The weighted linear combination of cone inputs has some parallels in psychophysically investigated opponent codes (Larimer, Krantz, and Cicerone 1974). The

slices associated with any such neuron remain planar and parallel, but are made at some characteristic angle to the axes of cone excitation space. Importantly, the planar slicings also have in general an unequal spacing, expressing a combined influence of response nonlinearity and noise factors (see §5 below). In §7, on multidimensional codes, we note the adequacy of such planar slicings for representing multidimensional input distributions that satisfy an independence criterion. When independence is violated, curved slicings are indeed superior; these, however, require a deeper nonlinearity than the simple response nonlinearity that we initially assume and that has empirical support—at least as an approximation—in the context of colour coding.

4. Best directions for slicing colour space

The visual system can associate any chosen slicing direction with an individual neuron or class of neurons, by appropriately setting the weights of the cone inputs to that neuron. In Fig. 1, an attractive choice would be to confer on one type of neuron a luminance (=L+M) sensitivity, allowing such a neuron to make, say, 10 slices spanning the diagonal that forms the major axis of the distribution of natural colours. Then if a second neuron generates a signal that depends on an opponent combination of L and M excitations, its 10 slices could be spaced much more finely so as to just span the more limited range of natural colours in the red/green direction, as shown in Fig. 1(b); many more of the squares representing potentially distinguishable colours would then be occupied. The domain of natural colours can in this way be divided into a much greater number of distinguishable colours than was possible using purely L- or M-cone driven neurons.

Note however that this benefit of slicing along diagonals occurs only if the number of slices made by each neuron is fixed, in keeping with the assumed restriction on response range. If there were no such restriction, the grid of Fig. 1(b) could be rotated arbitrarily without reducing appreciably the number of distinct colours. For fine slicings, that number is simply equal to the volume of the colour 'sausage' being sliced, divided by volume of the pieces (distinguishable colour domains) into which it is sliced. It is therefore invariant with the orientation of the slicing grid. The essential role of neural nonlinearities, such as range restriction, in determing the optimal slicing pattern is not recognized or allowed for in treatments of colour coding which suggest that the optimal directions for slicing are those of the principal components of the distribution of photoreceptor excitations (Atick, Li, and

Redlich 1992; Buchsbaum and Gottschalk 1983; Fukurotani 1982). Principal components analysis (PCA), appropriate as it is for a linear system, provides no mathematical or biological rationale for preferring one slicing grid orientation over another (even for a nonlinear system, whose nonlinearity is not allowed for in the PCA analysis)[1]. The benefits of encoding the diagonal variables in Fig. 1, or uncorrelated variables in general, are contingent on the limited range of neural response, a nonlinearity which limits the number (and thickness) of the slices that are created by each type of neuron. Moreover, while principal components analysis can specify a set of *directions* for slicing the input space, it cannot indicate a preferred origin of the coordinate system, and hence can give no rationale for opponent codes. We will present a rationale for opponent neural codes—and for split range codes in particular—which takes as its starting point the idea that the *thickness*, and not the direction, of the slices should satisfy a principled requirement: specifically, the arrangement of slice thicknesses should provide the most precise representation of natural colours. This criterion allows determination of an optimum nonlinear code, subject to the constraint of a limited output range. The optimal code is a split range one, with rectifying opponent cells. And as explained in §7, the axis directions for the optimum nonlinear code, if the distribution of colours satisfies an independence condition, are simply the principal component directions.

5. Best choice of slice thickness: the pleistochrome

We will assume for simplicity that the relevant neural signals depend upon a weighted combination of cone excitations (Derrington, Krauskopf, and Lennie 1984; Larimer, Krantz, and Cicerone 1974). In terms of our crude slicing analogy, this means that each neuron performs a slicing of cone excitation space into slabs bounded by parallel planes. If the neural signal were linear in its dependence on its net input, the thickness of these slices might be uniform. But owing to limitations on firing rate at both ends of the scale, the dependence of output on input is necessarily quite nonlinear, and the slice thickness will be correspondingly non-uniform. The pattern of slice thickness in the input

[1] The principal component axes are optimal when there is a requirement for dimensionality reduction in a linear system, but no functionally important dimensionality reduction is typically involved in the postreceptoral

space will depend not only on the form of the nonlinear neural input-output function, but also on the degree of random variability in the output, in a manner that we will consider shortly. A second assumption adopted for simplicity at the outset—we will dispense with it later—is that variability originating at the retinal output predominates over sources of error at earlier stages. The rationale for this is that the optic nerve constitutes an informational bottleneck for vision, where the number of nerve fibers and of nerve impulses is relatively limited: much less, for instance, 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 fiber impulse counts (Bialek and Rieke 1992; Lee and others 1993).

Since the retinal output represents the input stimulus for the perceiving organism, any random error originating at the output stage implies a corresponding error in the perceptual estimate of the input value. This error in the estimated input depends both on the output error and on the gradient of the input-output response function, as illustrated in Fig. 2. There the output noise, with root-mean-square (RMS) value σ, may be regarded as defining the (vertical) slice thickness at the output. In Fig. 2 the noise, and the vertical thickness of the illustrated horizontal slices, is constant with variation in the mean input or output. The associated variation in the represented input value (RMS equivalent input noise) is proportional to σ, and inversely to $f'(x)$, the derivative of the function relating input to output. This defines reliably distinguishable slices at the input; in Fig. 2 these are non-uniform in width owing to the compressively nonlinear input-output function.[2]

Note that because of the reciprocal relation between the response function gradient and the associated error in the estimated stimulus value, discrimination around any input value can be made as good as desired simply by making the response function gradient at the relevant point as steep as necessary. But the constraint imposed by the limited total available range in firing rate means that (for monotonic response functions) an increase in gradient at one point has to accompanied by a decrease at other points within the input range, and hence by reduced discrimination at those points. Thus a problem confronts the visual

---

coding of color.

[2] For now, we make the traditional simplifying assumption that the relevant visual nonlinearity can be treated as a static one. In reality, the nonlinearity of postreceptoral visual neurons is preceded by a sensitivity-regulating mechanism. In section 9 we consider how recogntion of this affects the analysis.

system. By suitable choice of a nonlinear response function, relative discriminative precision can be distributed in any desired way over the range of input values. So which choice is best? With a mathematical definition of 'best', this problem has, as we will show, a definite solution. We may ask, for example, what input-output function is best in the sense of giving the smallest error (in for instance, the least squares sense) in the estimated input, averaged over all naturally occurring cases. The optimal code in this sense must depend on the distribution of environmental inputs. 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 (dotted curve, Fig. 3). By doing this the system would sacrifice useful discrimination among naturally occurring stimuli in order to preserve discrimination in ranges where it is never needed. Other things being equal, it is advantageous to allocate discriminative power preferentially to the part of the input range where discrimination is most often needed—that is, the part where natural colours most frequently occur—by giving the response function a steep gradient at that point in the stimulus range. Also inefficient, however, is the opposite extreme from the linear code: a response function that steps abruptly from minimum to maximum firing rate at the peak of the distribution of natural colours (dashed curve, Fig.3). This choice would lead to categorical perception, in which blues and yellows, for instance, might be unmistakably distinguished at a precisely defined bounding chromaticity, but with no discrimination within each class. The best choice will be an intermediate one: a gently curving sigmoid with steepest gradient at the peak of the distribution, but with a non-zero gradient also in the tails of the distribution. This retains some discrimination in the tails while slicing colour space most finely at the peak.

The optimal response function can be found once the cost of a perceptual error has been defined. The optimal condition can then be determined by noting that the benefit of increasing the gradient very slightly from its optimum value at any chosen point must be exactly cancelled by the cost of the necessary equal reduction in the gradient at any other point. That is, the derivative of cost with respect to gradient must be the same at all points in the optimal condition. One definition of cost is implicit in the adoption of a mean squared error criterion: To minimize the sum of the squares of all errors made in the perception of the stimulus set is to minimize a cost proportional to the square of each error. We begin by considering the simplest case—where the cost of a given error is the same at all points in the stimulus range, and proportional to the squared error—and show that in that case, the

optimal condition is achieved by a response function with a gradient matched to the cube root of the probability density function of the input distribution.

Let the output signal of interest (one element of the vector making up the postreceptoral colour code) be y^*, with a mean $y = g(x)$ and an associated random variation (standard deviation) $\sigma(y)$, for a net input (a weighted sum of cone excitations) x. Denote the environmental probability distribution of x (for all stimuli encountered, or of interest) by $p(x)$.

Let x^* be the perceptual estimate of the input value x based on the output, y^*. We want to minimize the mean squared random error (MSE) in x^* originating from random variation in y^*. For given input x, this mean squared error in x^* is

$$\sigma^2(x^*) = \left(\frac{\sigma(y)}{g'(x)}\right)^2 \qquad (1)$$

where $g'(x)$ is the response gradient, or the derivative of the response function $y=g(x)$ at x. The average of $\sigma^2(x^*)$ for all inputs is its probability-weighted integral over x, which converges if $p(x)$ decreases more rapidly than x^3 for large x:

$$MSE = \int p(x)\, \sigma^2(x^*)\, dx$$

Consider the effect of small variations in the response gradient or incremental gain $g'(x)$ around its optimal value. As explained above in the optimal condition, $p(x)d(\sigma^2(x^*))/dg'(x)$ must be independent of x. Thus

$$d(\sigma^2(x^*))/d(g'(x)) = k/p(x) \qquad (2)$$

where k is a constant of proportionality.

Since from (1)

$$d(\sigma^2(x^*)/d(g'(x)) = -2(\sigma(y))^2/(g'(x))^3$$

the optimal condition occurs when

$$g'(x) = (1/k)(\sigma(y))^2 p(x))^{1/3} \quad (3)$$

The scaling factor k serves only to define units of measurement for y and can be set to 1 if those units are not defined independently.

If $\sigma(y)$ is independent of y, and more generally as an approximation if the factor $p(x)$ predominates, the above reduces to

$$g'(x) = p(x)^{1/3}$$

or

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

We refer to this optimal response function as the *pleistochrome*, from the Greek *pleistos* meaning 'most', since it may be roughly described as the function that makes available the greatest number of distinguishable colours. More strictly, it is the function that maximizes the average precision with which input colours are represented. For single-peaked stimulus distributions the pleistochrome is a sigmoidal curve centered near the peak of the distribution of x (Fig. 4). It is roughly similar to the cumulative distribution of x, but wider than that function by a factor of about the square root of three. A similarly motivated proposal of Laughlin (Laughlin 1983) aims to maximize the information about the input, given the output, through histogram equalization. The infomax criterion does not derive from a noise-based theoretical framework, however, and lacks a rationale in terms of minimization of random error or the associated cost. It leads to a steeper nonlinearity (by about a factor of the square root of three) than the least-error one embodied in the pleistochrome (Fig. 4)..

6. The pleistochrome under less restrictive assumptions: non-uniform noise and other complications

More general cases than the simple one that led to equation (4) turn out to be mathematically tractable. Here we list some ways to elaborate or extend the initial scenario. More formal and rigorous treatments of these cases, and of the pleistochrome in general, can be found in von der Twer and Macleod (in preparation). Readers with limited enthusiasm for quantitative theory may prefer to skip the next section or two.

Output noise can be non-uniform. Equation (3) determines the pleistochrome in the general case where the random variation introduced into the output has a standard deviation that varies with the mean output. Equation (4), on the other hand, treats the output noise as independent of the the mean output. But a very simple intuitive connection exists between that case and the general one. When output noise is output-dependent (but monotonic with mean output) there exists some nonlinear *monotonic transform* of the output that has a standard deviation independent of its mean: the transforming function need only have a derivative inversely proportional to the output noise standard deviation at every point. The optimal condition is when this function of the output should conform to Equ. (4).

For instance, in the case of the Poisson process that provides the simplest idealization of a spiking neuron and is often approximately descriptive of real ones (Levine, Cleland, and Zimmerman 1992; Levine, Zimmerman, and Carrion-Carire 1988; Tolhurst, Movshon, and Dean 1983; but see also (Croner, Purpura, and Kaplan 1993), the standard deviation goes up as the square root of the mean firing rate, and so the square root of the output has a constant standard deviation. For a neuron with a mean output firing rate $n(x)$ for input x, with a lowest firing rate n_{min} (perhaps representing spontaneous activity, or perhaps reflecting maximal stimulus-derived inhibition) at $x = -\infty$, a maximum rate of n_{max} and a standard deviation $an^{1/2}$, the standard deviation of $n(x)^{1/2}$ will be $a/2$ for all x, and the optimal dependence of n on x is such that $n^{1/2}$ satisfies Equ. (4), i.e.

$$n(x)^{1/2} = (n_{\min}(x))^{1/2} + ((n_{\max}(x))^{1/2} - (n_{\min}(x))^{1/2}) \int_{-\infty}^{x} p(u)^{1/3} du \qquad (5)$$

so that $n(x)$ is given by the square of this expression. A second example: with noise fluctuations whose range spans a constant fraction of the mean rate, the *log* of the rate has a constant standard deviation. Then it is the log of the output that should satisfy equation (4)

in order to minimize the average error in the estimate of the input, so the optimal nonlinear response function is an exponential function of $g(x)$ in (4).

A key point to note is that the variation of equivalent input noise with input x in such cases is completely independent of the function $\sigma(y)$ that specifies the variation of output noise with mean output y, since the effects of the latter are cancelled (except for an overall sensitivity factor) if the nonlinearity $y = g(x)$ is in each case the one appropriate for $\sigma(y)$.

Multiple nonlinearities are treatable stage by stage. In a system where there are multiple stages that impose nonlinear transformations on their inputs, equation (3) provides a prescription for the optimization of each such transform separately, proceeding downstream from the initial stimulus. Errors arising in segments of the system where processing is linear (including on the one hand those preceding, and on the other hand those following, the nonlinear stage of interest) can be lumped together. Optimizing response nonlinearity for each individual nonlinear stage by this procedure is not, however, equivalent to optimization of the whole system. For that, iteration would be necessary.

The input can itself be contaminated by error. Sources of error may exist prior to the nonlinear stage of interest, and this input noise may be stimulus-dependent. Conveniently enough, this doesn't affect the pleistochrome at all: Equation (3) still guarantees least mean squared error. This is because the input noise (if uncorrelated with the output noise) simply adds an optimization-irrelevant constant to the mean squared error that must be minimized. When input noise is stimulus-dependent, however, it may be appropriate to weight differently the errors of estimation for different stimuli, as discussed below.

Input variability can be non-uniform (stimulus-dependent). To the extent that Weber's Law applies, the cost of an error in estimating the stimulus value may be related less to its absolute magnitude than to its magnitude as a proportion of the value being estimated. More generally, the error cost will depend on the stimulus value in a way that reflects any prior uncertainty about the stimulus value--uncertainty due to sources of error like photon fluctuations and photoreceptor noise that contaminate the signal before it arrives at the nonlinear stage. Such non-uniformity in input variability can be dealt with by weighting the errors for different values of x by an additional suitably chosen factor besides the probability $p(x)$. In just the same way as was described for nonuniform output noise, this is equivalent to first determining a transformation of the input x for which estimation errors of equal

absolute magnitude are equally undesirable (e.g. a logarithmic one, if Weber's Law describes the error cost), and then applying equation (3) to derive the optimum dependence of the output firing rate on this function *f(x)*, and thence on *x*.

Figure 3 illustrates this, since there the probability distribution is for the logarithm of the *b* chromaticity coordinate. As a result the pleistochrome shown by the continuous curve in Fig. 3 is the response function that minimizes error in the estimate of log(*b)* rather than in the absolute value of *b*. The rationale for preferring this to the response function that minimizes error in the recovery of linear *b* values is that equal differences in log(*b)* are about equally noticeable (Boynton and Kambe 1980; Le Grand 1949).

The cost of a given error of estimation may vary depending on the estimated stimulus value. Some discriminations are biologically important, others less so. If important discriminations tend to be concentrated at certain points in the input domain (as suggested by Osorio and Bossomaier, in a paper to which we will return) some allowance should be made for this, and the establishment of a mathematical criterion for optimization might then seem hopeless. To this objection we have two responses. First, it is never appropriate to insist that every relevant factor be incorporated into a mathematical idealization of a problem. As Körner (Körner 1996) observes, mathematical analysis always require us to 'look at the rich complexity of the real world and replace it with a simple system which, at best, palely reflects one or two aspects of it.' Second, non-uniformity in the cost of error over the stimulus domain can in fact be handled formally in exactly the same way as non-uniform noise contamination of the input. When the cost of errors of estimation is different for different stimulus values (whether due to non-uniformity in input noise, or to other considerations), this can be dealt with simply by appropriately weighting those errors, or—equivalently—by applying Equ. (4) not to the initial input value *x* but to a transformation of it, *f(x)*, chosen such that estimation errors of equal absolute magnitude in *f(x)* are equally undesirable.

The relative cost of small and large errors can be chosen freely. As noted, the use of least-squared error in the estimated input as a criterion for optimization is equivalent to assuming that errors entail a cost simply proportional to their squares. If we choose to minimize the absolute error rather than its square, small errors are not as well tolerated, and the incentive to make very thin slices in high density regions is increased. By developments similar to those that led to Equation (3), this leads to an optimal incremental gain $g'(x) = (\sigma(y)p(x))^{1/2}$, and thus to an

input-output function spanning an input range a little less wide than in the least square case (Fig. 4).

An interesting situation arises if noise sources inherent in the input predominate over errors originating at the output. In the visual system, this situation arises at low light levels (Barlow, Levick, and Yoon 1971; Baylor 1987; Donner 1992). In this case, a relatively small output-derived MSE is added to the relatively large MSE inherent in the input itself to give the total MSE in the perceptually estimated input. Provided that the output-derived increments in total MSE are sufficiently small, any useful measure of average error or of the associated cost—mean square, mean absolute error or anything else—will increase approximately linearly both with the total MSE and with the *mean square* error contributed from sources of variation at the output. Here the optimum input-output nonlinearity is determined for *any* cost function by first transforming the input so that input-derived error (or more precisely, the cost per unit of added error in the estimated input) is independent of mean input, and then applying the MSE pleistochrome (the cube root construction of Equation (3)) to derive the best dependence of output on that input. In this case the cube root construction, with its relatively shallow derived pleistochrome, remains optimal even if the cost function for errors is not quadratic.

7. Multidimensional stimulus domains: slicing colour space

In extending these ideas to a multidimensional stimulus domain like colour we encounter interesting problems that we discuss more fully elsewhere (von der Twer and MacLeod, in preparation). A particularly simple extension is possible if the stimulus probability density function satisfies independence for some input quantities u and v (which need not be linearly related to the initial stimulus coordinates, for instance the cone excitations), that is if $p(u,v) = p_1(u)p_2(v)$. Then the pleistochromes for the marginal distributions $p_1(u)$ and $p_2(v)$ specify the optimal spacing of linear cuts parallel to the u and v axes—constant-response contours of signals encoding u and v –for minimizing perceptual error in (u,v). Moreover these linear cuts (or for higher dimensionality, planar cuts) in (u,v). are more efficient in that sense than any curved cuts, since the u-pleistochromes are the same for all v and vice versa, and since as we show elsewhere the choice of other axes than those satisfying independence introduces added reconstruction error. The optimal neural responses $f_u(u)$ and $f_v(v)$ are obtainable in

terms of the original stimulus coordinates x and y by inverting the transformation that generated u and v from those stimulus values. The constant-response contours in (u,v) will be curved unless the transformation between (u,v) and (x,y) is linear. But two *caveats* apply. First, the error is minimized in the coordinate system (u,v) where independence holds, not in the original stimulus coordinate frame. The optimization is valid, therefore, only where independence holds in some coordinate frame in which the cost associated with errors is uniform. This happy situation may exist for colour space, with the adoption of appropriate simple linear combinations of the logs of the cone excitations—quantities rather close to log(r), log(b) and log(luminance)—as coordinates (Ruderman, Cronin, and Chiao 1998).

Second, there is some freedom with respect to the orientation of the coordinate frame in (u,v). A suitable choice of u and v will make the marginal distributions Gaussian. If independence holds, i.e. $p(u,v) = p_1(u)p_2(v)$, the distribution $p(u,v)$ then becomes a bivariate Gaussian with radial symmetry. Then, if equal costs are associated with equal errors in different directions in (u,v), all orthogonal coordinate frames in (u,v) are equally efficient. In Fig. 1 for example, where the distribution is elongated obliquely in terms of the original stimulus coordinates, the independence condition requires adoption of other coordinates. These are necessarily orthogonal in (u,v), but need not be orthogonal in the original input space of Fig. 1. The strategy of encoding the principal axes of the distribution—roughly, luminance and colour as in Fig. 1(b)—is therefore only one possible choice; the other choices slice the plane of Fig. 1 into narrow diamonds rather than narrow rectangles.

More complicated (but still continuous) distributions, that violate independence overall, will still approximate it locally over sufficiently small neighborhoods, and this will call for local variations in the orientation of the slicing grid.[3] Good orientations for slicing an

---

[3] Even when independence does not hold, any single-peaked input probability density distribution can be mapped into a radially symmetrical one (though not, in this case a Gaussian one) by some continuous one-to-one deformation of the input space, and here just as in the case of independence, free choice of orientation of the coordinate system becomes possible. Failure of independence is, nevertheless, helpful for defining the best direction of the coordinate system for encoding the stimulus set, because the mapping of a distribution that violates independence into a radially symmetric one is more constrained than the mapping to (u,v) in the case of an independent distribution: independence is preserved under separate arbitrary transformations of u and v. Thus it is easier to find a coordinate frame satisfying independence, than it is to find one yielding radial

elongated or more complicated distribution can be determined without iteration in the following way. We fix all but one of the stimulus variables, and determine the pleistochrome for variation in the remaining variable alone. When this is done in turn for a range of values of the other variables, we obtain coordinates of stimuli that give equal response for one of the neural signals of interest, and these trace out a constant response surface for that signal. This generally differs in orientation from the initially adopted stimulus axes. If independence is not satisfied, the surface is curved—appropriately so, since in that case the constant-response contours have to crowd together to improve the precision of the representation in those regions where stimuli occur more frequently than expected from independence . We have found this approach quite effective in reducing the reconstruction error not only for the simple sausage-shaped distribution but also for more challenging stimulus distributions (banana-shaped, L-shaped, bagel-shaped, or the real colour distribution illustrated in Fig. 7). The technique shares with Independent Components Analysis (ICA) (Bell and Sejnowski 1997), the merit that non-orthogonal rotations of the initial axes are permitted, but is more versatile than ICA in allowing non-planar cuts. The spacing rule of equation (3) also differs from the one customarily adopted in ICA, which, like the proposal of Laughlin (Laughlin 1983) aims to maximize the information about the input, given the output, through histogram equalization.

As Figure 7 shows, the spacing of the constant-response contours dictated by error minimization is only slightly non-uniform within the (very limited) chromaticity range of natural colors. The implied relatively gentle nonlinearity allows much better discrimination among colors near the margins of the distribution that would result from histogram equalization.

Optimal treatment of complex distributions would require a procedure where the grid of constant-response contours is allowed to vary locally in orientation and scale in order to minimize the total reconstruction error. An adaptation of Kohonen's iterative procedure for creating self-organizing neural maps (Kohonen 1989), might be appropriate; intriguingly, this procedure has been shown to generate a set of spacings consistent with the cube-root pleistochrome (Equ. (3)) in the 1-dimensional case (Ritter, Martinetz, and Schulten 1989),

---

symmetry in a case where independence is not possible; and by the same token a successful choice is then less likely to be unique.

even though it makes no attempt to evaluate or to minimize error in the representation of the input.

8. Benefits of split range coding: something for nothing

As perspicacious readers may have noticed, our discussion of optimum response nonlinearity has not yet yielded the promised rationale for opponent codes. On the contrary, the sigmoidal nature of the pleistochrome, with its steepest gradient at or near white, is quite 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 encoding is done by a pair of neurons. Clearly 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 increases in firing rate, and with null responses 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). As Fig. 5 shows, the responses of two such neurons (isolated open circles and squares) correspond to the two halves of the single-neuron pleistochrome sigmoid (plain curve, Fig. 5), 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. The optimal implementation of such a 'split range' code using two rectifying neurons is instead to set the cross point, x_0, at the 50% point on the sigmoid, and then replace the integral from $-\infty$ in Equation (4) or (5) with *twice* the integral upward or downward from x_0. In this way the appropriate segment of the input range can elicit the maximum possible response range from each neuron, and the gradients of the response functions are doubled everywhere (filled circles and squares, Fig. 5). This account assumes ideal rectification behavior, with no response on the wrong side of x_0 for each neuron.

By using two neurons in this way the visual system can double the precision in its representation of the input in the presence of output noise, since the gradients of the individual response functions have been doubled and the equivalent input noise thereby halved. 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 led to an improvement of only a factor of the square root of two in average error. 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. The spontaneous activity shown by real neurons without stimulation, or on presentation of the null stimulus, entails some reduction in efficiency within the present framework, but if the spontaneous activity is a small fraction of the maximum firing rate most of the advantage of split range coding is preserved. Parenthetically, we note that reserving high firing rates for unusual stimuli may have other important advantages as well: it usefully facilitates selective response to unusual inputs (Barlow 1972; Field 1987), and by reducing average firing rate and neurotransmitter release it lowers the metabolic cost of perception.

With Poisson noise, low spontaneous activity $n_{min} << n(x)$, and a smoothly peaked $g'(x)$, the split range code with Equ. (5) leads to a threshold-like, approximately quadratic increase in firing rate with stimulus value (here $|x-x_0|$) as x moves away from the null stimulus value x_0 (dashed curves, Figure 5). In conjunction with the output-dependent noise, the effect of this is to make discrimination relatively keen but constant in that neighborhood. Here we have an intriguing possible function for threshold nonlinearity: its role could be not to make the effective precision of the neural representation non-uniform over the relevant part of the input range, but to make it uniform.

The split range code may be viewed as a step from a purely analog representation (the sigmoidal single-neuron pleistochrome) to a hybrid, analog-digital one. Although the benefits of taking that step may seem intuitively surprising—because in adopting the split range code, the visual system appears to get something for nothing—the process could be taken further, with still greater ensuing benefits. In a fully digital encoding of a stimulus dimension, a set of N neurons, each with m reliably distinct outputs, can represent m^N different stimulus levels by allocating successive digits of the digital representation to the different neurons. This compares with just mN distinguishable levels (in the simplest case) for a split range code where the input range is divided into m segments each spanned by the graded firing range of one neuron, or with $mN^{1/2}$ for a parallel averaging of signals from neurons with identical response functions (Fig. 6(a)). But the fully digital representation is dangerous and difficult to implement. Because individual neural outputs depend discontinuously (in a sawtooth manner) on the input value, it creates a risk of large errors.

Encoding schemes known as Gray codes (Savage 1997) avoid this discontinuity by replacing the sawtooth function of Fig. 6(a) with a symmetrical triangle function, but since these schemes still require the individual digit values to vary in a nonmonotonic manner with the stimulus value, they still pose formidable difficulties for a biological system in both encoding and readout.. A split range code with $N > 2$ (Fig. 6(b)) is not subject to any such problems—it allows a simple center-of-gravity or weighted sum of neural firing rates to represent the stimulus value. This may be the only biologically plausible alternative to the simpler dual-opponent-neuron encoding scheme ($N = 2$) that is shown in Fig. 5 and that is usually taken to be representative of physiological findings in colour vision (Derrington, Krauskopf, and Lennie 1984; DeValois and DeValois 1975). But there seems to be no conclusive evidence, from psychophysics or from physiology, for the staggered arrangement of null planes and response functions among different neurons that this scheme would require.

9. Comparison of colour appearance and discrimination with predictions based on pleistochrome

Appearance. If the colour opponent code is designed for optimal characterization of natural colours, the null stimulus should be close to the center of gravity of the distribution of natural colours. And if the null stimulus in the chromaticity diagram is the subjectively achromatic white, typical natural colours should be nearly white. This is of course roughly correct, but as shown in Fig. 7 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. The constant luminance chromaticity diagram of Fig. 7 has axes $r = L/(L+M)$ and $b = S/(L+M)$, which are closely related to the inputs to different classses of colour opponent neurons in the lateral geniculate nucleus under adaptation to equal energy white (Derrington, Krauskopf, and Lennie 1984). The contour map is for Ruderman et al's distribution of colours of natural surfaces, under D65 illumination, which simulates a slightly overcast daylight. The equal energy white stimulus plots at $r=0.70$, $b = 1.0$, and the heavy straight line, $r = 0.723 - 0.0325b$, shows the approximate locus of colours that are subjectively neither reddish nor greenish when presented in the dark (Larimer, Krantz, and Cicerone 1974). Clearly the centroid of the distribution of natural colours is displaced

toward yellow (low *b*) and toward green (low *r*) from the subjective neutral point. The shift is small—small enough to be substantially influenced by the choice of assumed illuminant—but it amounts to several times the colour discrimination threshold. In Brown's sample of haphazardly selected colors, the mean is again yellowish, but is reddish rather than greenish: leaves form a dense concentration in the green, but their influence is outweighed in the case of the *r* axis by the inclusion of many reddish fruits and flowers, leading to mean *(r,b)* coordinates under D65 of (0.715, 0.636) as compared with (0.691, 0.750) for the complete images of Ruderman et al.

There are a number of more or less plausible *post hoc* rationalizations for the somewhat unexpected placement of the white point. 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. It is not clear to what extent exclusion of the sky in the determination of the parameters of the opponent code is desirable. And while complete exclusion is doubtless within the powers of evolution, it might be difficult if the opponent code is adaptively determined by accumulated stimulation during development.

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 (Osorio and Bossomaier 1992) suggest that discrimination among the greens of vegetation is not particularly important, whereas discrimination of reddish fruits from vegetation is. A null point, with optimal discrimination, near white might usefully promote those discriminations at the expense of the less important ones.

In Fig. 7, the *b* and *r* chromaticity coordinates show some negative correlation ($r = -.18$). This is not unexpected given the spectral sensitivities of the cones, which create a negative correlation between *r* and *b* across the spectrum. Weak as the correlation is for natural surfaces, it indicates that the *(r,b)* coordinate frame is not quite the optimal one for slicing colour space.

As noted above, the orthogonal vectors in terms of which the natural colours in Fig. 7 can be represented with least average error by neurons with limited range, are roughly the principal component vectors, for which the independence relation $p(x,y)=p(x)p(y)$ is best approximated. Owing to the negative correlation between *r* and *b*, the principal component coordinate frame in Fig. 7 is tilted anti-clockwise from the *(r,b)* fame. The red/green equilibrium axis of the Hering opponent scheme, which connects colours that are neither reddish nor greenish, is likewise tilted anti-clockwise (straight line, Fig. 7). It is therefore possible to suggest that the subjective redness of violets (located above the red/green equilibrium locus near the middle of Fig. 7), and the associated near-circularity of the spectrum in phenomenal colour space, is needed for optimal discrimination among natural colours. These phenomena of colour appearance result from an alliance of the short-wavelength cones with the long-wavelength cones in the psychophysically defined red/green opponent system. No such an alliance is found in the cells of the lateral geniculate nucleus (Derrington, Krauskopf, and Lennie 1984). The LGN therefore embodies non-optimal—because correlated—signals. In the LGN representation, it is blue colours that most strongly polarize the M-L ('red/green') colour opponent signal in the M ('green') direction and are the most likely to overload it; in certain multiple stage models of the colour system (DeValois and DeValois 1993; Müller 1924), this tendency is counteracted in the third and final stage by a short-wave cone input, synergistic with the long-wave cones and antagonistic to the midspectral cones. The linear decorrelation principle does not, however, predict the red/green equilibrium locus exactly, as the experimental red-green equilibrium axis is tilted nearly 5 times more than the principal component direction for the natural stimulus distribution in Fig. 7, and about twice as much as the principal component direction for Brown's similar data. As Fig. 7 shows, it is also somewhat more tilted than the system of curved constant-response contours generated by the nonlinear algorithm of §7, but here the correspondence is closer.

Discrimination. We have suggested that the opponent split-range code has evolved in the interests of minimizing errors in colour perception, through the adoption of an optimally designed nonlinearity in neural response. That proposal leads to definite quantitative predictions for the discriminability of stimuli that differ in colour or intensity. The predictions depend simply on the form of the distribution of natural colours that the system has evolved to deal with, rather than on known or estimated physiological parameters of the

visual system. If errors of discrimination are indeed distributed across stimulus dimensions in such a way as to minimize the average error, then the root mean square error implicit 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 $p(x)$. Conveniently enough, as noted in §6, non-uniformity in output noise (embodied in the factor $\sigma(y))^2$ in Equation (3)) does not affect this prediction at all, provided that the nonlinearity $g(x)$ that couples the stimulus x to the response y is optimized for the prevailing dependence of $\sigma(y))$ on y.

Data for evaluating this prediction are available. The mean stimulus difference needed to make a test stimulus just noticeably different from a standard depends upon the colour or intensity of the standard (Krauskopf and Gegenfurtner 1992; Miyahara, Smith, and Pokorny 1993). In typical experiments the test and standard stimuli are intensive or chromatic modulations that appear in separate regions within a steady, generally white adapting field. Discrimination is most acute if the standard and test are both very close in colour and intensity to the adapting white; and quite small differences of the standard from white seriously impair the precision of the comparison. Current experiments by A. Leonova (in preparation) quantify this for difference directions in colour space, using as a metric for stimulus differences the cone contrast between the standard stimulus and the adapting white. For achromatic contrasts and achromatic intensity differences between test and standard, comparison error is doubled for a standard contrast of about 20%; 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 (Fig. 8).

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. As a rough approximation, this prediction is borne out: the distribution of Ruderman et al. is indeed extremely narrow in the red/green direction, although it is about 25% wider than would be required for a best fit to the data of Fig. 8. Further, since discriminations in luminance, or in the blueness-related chromaticity coordinate b, are both maintained for a range of standard colours an order of magnitude greater than are discriminations in r, the distribution of

environmental colours for these coordinates should be an order of magnitude or so broader than for *r*. And indeed it is wider, by an order of magnitude or more (see §2).

These comparisons suggest that 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. By encoding the *r* dimension of colour space with a particularly steep and narrow sigmoid (with halved discrimination at a deviation of 2% from the null value of *r*)—one that is matched to the small environmental range in *r*—the system is able to make correspondingly finer discriminations within that narrow range. It would be disastrous, however, to encode only the same narrow range for *b* as for *r*, since this would result in almost incessant overloading of the blueness signal, with a purely categorical classification of the visual scene into intensely bluish or yellowish colours.

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. The variation in luminance greater than the variation in *r* by a factor of 15 (Brown 1994) or as much as 60 (Ruderman, Cronin, and Chiao 1998) as opposed to the factor of 10 or so expected.[4] This exacerbates the deviation from prediction that we noted in the analysis of red/green discrimination data: in the presence of a white adapting field, the range of standard colours allowing good discrimination of colour and intensity is distinctly narrower than expected on the basis of the distribution of environmental inputs. Why should the operating range of the visual system be narrower than 'optimal' in this way?

Adaptation, local contrast and the pleistochrome. One answer points to a deficiency in our theoretical framework, which does not incorporate the important phenomenon of visual adaptation. Prior to the extraction of a colour opponent signal, the cone photoreceptors individually take up a sensitivity inversely related to their short term average intensity of stimulation [Chaparro, 1993 #18; Chichilinsky, 1995 #20; He, 1997 #19; Boynton, 1970 #21; Valeton, 1983 #22; MacLeod, 1992 #45}. For this and other reasons, retinally stable images fade in perception (Ditchburn 1973). This implies that the effective stimuli for

---

[4] It is not clear which of these two divergent estimates of the ratio of achromatic to chromatic environmental variance is to be preferred. The ratio in the whole scene-based data could be inflated by the uncompensated effects of local variation in illumination within the scene (Brown 1994). On the other hand, Brown may have favored highly chromatic objects in selecting his samples.

postreceptoral cells are temporal transients, that are generated by small eye movements in conjunction with spatial gradients in the image. It is therefore relevant to consider the distribution of the spatial differences in cone excitation across a neighborhood small enough to be traversed by the gaze of a fixating observer. In the images of Ruderman et al, the differences in luminance, in *b*, and in *r*, between adjacent 3 min arc pixels, are suited for this purpose. Those differences have standard deviations of 30%, 13% and 0.6% respectively; the distributions are thus 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 a narrow but dynamically shifting operating range (Craik 1938), and thus adopt a roving null point, rather than a fixed one, for the colour-opponent code (Krauskopf and Gegenfurtner 1992; Thornton and Pugh 1983), if its objective is the precise representation of local contrast.[5]

The advantage of a roving null point is that the range of input values spanned by the neural response functions can be more restricted—it 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. We have generated local-contrast pleistochromes from the images of Ruderman et al.on this basis; these are the contrast-response functions that lead to least error in the representation of pixelwise spatial differences. They are consistent with psychophysical results in the case of the chromatic variables. For luminance, however, the contrast operating range implicit in the discrimination results remains narrower (by a factor of about three) than the theoretically optimal one.

A second limitation in our initial framework, also connected with the role of adaptation, may account for this remaining discrepancy for luminance. 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 differences in lightness and colour are also

---

[5] An important problem for a system operating in this way is: how can the spatial and temporal differential signals be used to construct a precise representation of color in absolute terms? This is discussed in other chapters and elsewhere (Arend 1973; Land 1964). The influence of nonlinearity of the code in this context has not been much considered, but has been

indispensable for the detection of spatial features (Morgan, Adam, and Mollon 1992). Precision in the representation of surface elements that are already recognizably distinct is not useful for that purpose—less useful at any rate than for identifying and characterizing surfaces. What is most critical for spatial vision is that 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, as in the inset to Fig. 8(a) 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 detection of spatial features on the other. The unexpectedly abrupt nonlinear saturation of the psychophysical signal for luminance contrast could have this as its *raison d'etre*. As we will see in the following section (§10), this is consistent with the common view that the luminance system (or in physiological terms, the magnocellular pathway) is more concerned with form and with detection of spatial structure than are the chromatic ones (e.g. Boynton, Hayhoe, and MacLeod 1977; Gregory and Heard 1979; Livingstone and Hubel 1987).

Anisotropy of color space. Thus far we have considered only how discriminative power is allocated (or should be allocated) along a stimulus continuum, but our ecological framework also raises issues concerning the relative precision with which different stimulus dimensions are represented. If we continue to suppose that discrimination is limited mainly by fluctuations in firing rate at the neural output, then if the same number of opponent units, with the same range and the same random variability, are used to encode the three dimensions of colour space, discrimination errors along each dimension will be scaled by in the same way as the total operating range for that dimension. That is, the mean error will be tenfold less for *r* than for *b* or for luminance. To a very rough approximation the data support this: Leonova's lowest threshold values are 0.8% for luminance, 0.1% for *r* and 3% for *b*. Red/green contrasts are indeed 'what the eye sees best' (Chaparro and others 1993)—but not in the sense that natural environments provide more detectable differences for the red/green dimension than for luminance. Rather, the sensitivity difference is as expected for two otherwise comparable systems, limited by output noise, that have very different contrast

responses, matched to the very different range of environmental inputs they receive and spanning that range with about the same number of reliably distinguishable signal levels in each case.

But the higher threshold for *b* than for luminance, despite similarity in operating range for these two dimensions, indicates that the visual system invests less in discrimination for *b* than for the other two dimensions. If small differences in *b* between surfaces were as important as small differences in *r*, it would be worthwhile to use more slices, of the same thickness (rather than the same number, of increased thickness) to span the more extended environmental range of *b* values. This might require using 100 times as many neurons for *b* as for *r*, to compensate for their tenfold shallower response function and maintain the same precision in estimates of *b* as of *r*. In reality, though, the visual system makes fewer slices in the *b* than in the *r* direction, as if by allocating *fewer* neurons to *b* than to *r*. There are physical reasons why small differences in *b* should be less important than differences in *r*. The *b* coordinate has a higher inherent gain than does *r*, in the sense that physically characterized differences of similar magnitude between surfaces—notably in the slope of their spectral reflectance functions—typically generate much larger contrasts in *b* than in *r*; this happens because the spectral separation of the S cone sensitivity from L and M is roughly sixfold greater than the L/M separation.[6]

---

explored in current experiments by Brown, Leonova and MacLeod that use non-uniform surrounds to diagnose the nonlinearity in the contrast response (Brown and MacLeod 1997; Brown and MacLeod 1991). Here we note only one point that is critical for our discussion: to construct a metric representation of color and brightness on the basis of spatio-temporal contrast signals generated at borders, those signals must themselves have a metrically meaningful dependence on border contrast, rather than being all-or-none. Our analysis of optimal nonlinearity thus remains applicable, requiring only the modification considered here—that it is error in the representation of spatial contrast that must be minimized.

[6] Thus while the differences reviewed here in postreceptoral processing of the three dimensions of color space may generally be regarded as a successful internalization of the statistical regularities of the external chromatic environment, it is more accurate, in the case of the b versus r comparison, to say that the opponent system has internalized a regularity that is "prior" rather than strictly "external", since the regularity in question (unequal environmental variation for r and b as cone inputs to the opponent stage) originates in the interplay between the population of external surfaces and the cone receptor sensitivities. Considerations of spatial resolution also dictate a reduced number of S cones and of "yellow/blue" postreceptoral cells that take input from the S cones. Here too, what makes the sparseness of S cones a good design choice is their own spectral isolation at

10. Comparison with physiological input-output functions: variations on a theme of Fechner

Much of our discussion makes use of ideas traceable to G. T. Fechner (Fechner 1860), who realized that stimulus-dependent variation in discriminative precision can be understood by invoking a nonlinear function that relates stimulus to physiological response or to sensation. Having found satisfactory theoretical connections between lightness and colour discrimination data and the form of the environmental stimulus distribution, relating as it were the input and the output of the visual system, we now embark on a neo-Fechnerian analysis of the same data, in order to compare the implied nonlinearity of opponent codes for lightness and colour with physiological data on the response functions of single neurons in the optic nerve.

In Fig. 8, predictions for the two extreme theoretical cases that were introduced in Fig. 3 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 $\Delta r = |r - ro|$, where ro is the colour category boundary. In Fig. 8 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 (Fig. 8(b)). The linear variation of the discrimination threshold with r on each side of the null point $r = .7$ in Fig. 8, with an abscissa intercept at $\pm r_0$, leads to a response-intensity function of the form

---

the short-wavelength end of the spectrum, where chromatic aberration prevents them from receiving a sharp image (Boynton 1980).

$$N = ln|(r\text{-}ro)/(0.7\text{-}r_0)| \qquad (6)$$

where r_o has a value of about .714 for the 'green' response and .68 for the 'red' response, equivalent to an L cone contrast of about 2% with respect to the null stimulus in each case. By reflection around the null point, $r = 0.7$, this specifies the value of r associated with a doubling of threshold, or a halving of differential sensitivity.

The logarithmic nonlinearity of Equ. (6) and Fig. 8(b) has no strict response limit, so the assumptions underlying Equs (3) and (4) are not applicable. Modifications of that framework can, however accommodate the Fechnerian nonlinearity. It might be desired, for instance to keep the firing rate below some practical limit, even though the logarithmic function does not entail such a limit. Or one could consider a revised optimization criterion: to minimize average error within the constraint of a given average response to all stimuli; this yields optimal prescriptions only subtly different from the pleistochrome of Equ (4). 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. (6)—optimal? The answer is $p(x) = p_{max} / (1+|(x/x_0)|)^3$. Although this peaked function does not fit distributions of individual surface colors, it does fit fairly 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.

The nonlinearity implied by Equ. (6) with the experimentally determined parameter values is quite severe. The gradient of the response function, assumed to be directly proportional to differential sensitivity, is halved at a cone contrast of 2%. No physiological data suggest quite so severely compressed a response function for responses to chromatic stimuli: half- saturation L cone contrasts of around 10% appear to be more typical, for the red-green sensitive P cells of the parvo-cellular stream (Lee and others 1990). Thus although the psychophysically estimated visual operating range is efficiently matched to the range of environmental inputs, the physiological one apparently is not.

This would not have alarmed Fechner, who located the logarithmic compression at the brain-mind interface and assumed that physiological processing would be linear. Even if we reject this particular reconciliation of nonlinear psychophysics with (relatively) linear retinal physiology as metaphysically unsound, it remains possible that later stages of

processing could be implementing the severe logarithmic contrast compression that Fechner's integration entails. But the theoretical link between discrimination and physiological nonlinearity is uncertain for at least two other reasons. First, physiological measures of half-saturating contrast may be made with stimuli that are not optimal in spatio-temporal structure. This leads to underestimation of half-saturating contrast, because the nonlinearity of retinal ganglion cells is a function of response rather than of the stimulus contrast *per se*. Second, in applying Fechner's integration to estimate hypothetical neural response functions we assume that just detectable differences correspond to equal differences in the neural signal. This amounts to assuming fixed output noise, whereas physiological observation suggests instead that the standard deviation in firing rate increases with mean rate. In some experiments the increase is almost linear with the square root of the mean, as expected for a Poisson process (Levine, Cleland, and Zimmerman 1992; Levine, Zimmerman, and Carrion-Carire 1988; Tolhurst, Movshon, and Dean 1983); in others, the increase in variability is small, and the fixed-noise idealization of cell behaviour is more appropriate (Croner, Purpura, and Kaplan 1993). If we choose to model neural firing with a Poisson process rather than a fixed-noise assumption, then the compressive nonlinearity required to model the red/green discrimination data becomes far less severe; spontaneous activity rate now becomes an important free parameter, but with plausible estimates of that, the r_o value needed to model the data becomes an order of magnitude greater than in the fixed noise analysis—more than enough to match roughly the physiologically measured nonlinearity. Psychophysical chromatic discrimination data and (P cell) retinal physiological data are therefore consistent after all, if the assumptions made about neural noise are tailored for a good fit between them. Fortunately the theoretical connection between discrimination and the distribution of environmental inputs, explored in §9, is not subject to these uncertainties.

Turning to the achromatic axis of color space, we saw in §9 that the psychophysical operating range in cone contrast along that axis is some tenfold greater than along the red/green one, and that the ratio of the dispersions of the environmental inputs is at least that large. If physiologically documented nonlinearities were consistent with visual performance on the one hand and with the stimulus statistics on other, the cells mediating judgments of achromatic intensity should likewise have a tenfold greater contrast range than those mediating red/green sensitivity. But which are these cells that construct the

achromatic axis of color appearance? A popular answer would be: the M cells of the magnocellular pathway (Livingstone and Hubel 1987). These cells, however saturate at very *low* achromatic cone contrasts, with half-saturation values of around 5%, even lower apparently than the 10% L cone contrast value quoted for red/green P cells (Kaplan and Shapley 1986; Lee and others 1990; Wachtler, Wehrhahn, and Lee 1996). The M cells, then, deviate more than 10 fold from the optimal behavior embodied in the pleistochrome, and as a result 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 in the sense discussed in §9. The metric representation of the achromatic axis could be the job of the color system (Allman and Zucker 1990). In support of this idea, physiological investigations such as those cited above have shown that the P cells have an almost linear response to achromatic contrast, consistent with the pleistochrome for achromatic inputs.

11. Concluding summary

We have shown how visual nonlinearity can be optimized for the precise representation of environmental inputs. Such optimization leads to the adoption of opponent split-range codes, and the recognition of this provides a new ecological justification for opponent codes. The key points in our account are:

- Nonlinearly compressed neural signals are needed in order to form the most precise representation of stimulus values from a peaked frequency distribution, using neurons of limited response range (§3).
- The optimal form for the nonlinear response function (the *pleistochrome*) can be determined given the distribution of inputs (§5).
- The treatment can be extended to multidimensional stimulus domains, notably to colour space (§7).
- When a single stimulus dimension can be represented by more than one neuron, a dual opponent or '*split range*' code, of the type familiar from the physiology of colour vision, is much more efficient that the optimal single-neuron code (§8).

- Some aspects of the phenomenology of colour vision, and of data on colour discrimination, are understandable on the assumption that the relevant neural codes have been selected for minimizing error in the perceptual estimation of stimulus parameters for natural colours (§9). In particular, for different dimensions of colour space the neural response function spans a range well matched to the environmental distribution of natural colours.
- Physiological data from the parvocellular pathway are also roughly consistent with the idea that these cells are optimized for precision representation of color. But cells in the magnocellular pathway have a much stronger than optimal saturating nonlinearity, and this supports the view that their function is mainly to detect boundaries rather than to specify contrast or lightness (§10).

Acknowledgements: We are grateful to Richard Brown, Anya Leonova, Dan Ruderman, Tom Cronin and C.C.Chiao for permission to use the data we cite from them (some if it not yet fully published). For useful conversations or comments we thank Brown, Leonova, Ruderman, Simon Laughlin, Markus Meister, and Harvey Smallman. Supported by NIH grant EY01711.

Allman, J., and S. Zucker. "Cytochrome oxidase and functional coding in primate striate cortex: a hypothesis." *Cold Spring Harb Symp Quant Biol* 55 (1990): 979-82.

Arend, L.E. "Spatial differential and integral operation in human vision: implications of stabilized retinal image fading." *Psychological Review* 80 (1973): 374-395.

Atick, J. J., Z. P. Li, and A. N. Redlich. "Understanding Retinal Color Coding From 1st Principles." *Neural Computation* 4, no. 4 (1992): 559-572.

Barlow, H.B. "Optic nerve impulses and Weber's law." *Cold Spring Harbor Symposium on Quantitative Biology* 30 (1965): 539-546.

Barlow, H.B. "Single units and sensation: A neuron doctrine for perceptual psychology?" *Perception* 1 (1972): 371-394.

Barlow, H. B., W. R. Levick, and M. Yoon. "Responses to single quanta of light in retinal ganglion cells of the cat." *Vision Research Supplement* 3 (1971): 87-101.

Baylor, D. A. "Photoreceptor signals and vision. Proctor Lecture." *Invest. Ophthalmol. Vis. Sci.* 28 (1987): 34-49.

Bell, A.J., and T.J. Sejnowski. "The "independent components" of natural scenes are edge filters." *Vision Res.* 37 (1997): 3327-3338.

Bialek, W., and F. Rieke. "Reliability and information transmission in spiking neurons." *Trends Neurosci* 15, no. 11 (1992): 428-34.

Boynton, R.M. "Design for an Eye." In *Neural mechanisms in behavior*, ed. D. McFadden: Springer-Verlag, 1980.

Boynton, R.M., M.M. Hayhoe, and D.I.A. MacLeod. "The gap effect: chromatic and achromatic visual discrimination as affected by field separation." *Optica Acta* 24 (1977): 159-177.

Boynton, R. M., and N. Kambe. "Chromatic difference steps of moderate size measured along theoretically critical axes." *Color Research and Application* 5 (1980): 13-23.

Brown, R.O. "The world is not grey." *Invest. Opthalmol. Vis. Sci. (Suppl.)* 35/4 (1994): 2165.

Brown, R. O., and D. I. MacLeod. "Color appearance depends on the variance of surround colors." *Curr Biol* 7, no. 11 (1997): 844-9.

Brown, R.O., and D.I.A. MacLeod. "Induction and constancy for color saturation and achromatic contrast variance." *Invest. Opthalmol. Vis. Sci. (Suppl.)* 32/4 (1991): 1214.

Buchsbaum, G., and A. Gottschalk. "Trichromacy, opponent colours coding and optimum colour information transmission in the retina." *Proc. Roy. Soc. Lond. B* 220 (1983): 89-113.

Chaparro, A., C.F. III Stromeyer, E.P. Huang, R.E. Kronauer, and R.T. Jr. Eskew. "Colour is what the eye sees best." *Nature* 361 (1993): 348-350.

Craik, K. J. "The effect of adaptation on differential brightness discrimination." *Journal of Physiology (London)* 92 (1938): 406-421.

Croner, L. J., K. Purpura, and E. Kaplan. "Response variability in retinal ganglion cells of primates." *Proc Natl Acad Sci U S A* 90, no. 17 (1993): 8128-30.

Derrington, A.M., J. Krauskopf, and P. Lennie. "Chromatic mechanisms in lateral geniculate nucleus of macaque." *Journal of Physiology (London)* 357 (1984): 241-265.

DeValois, R.L., and K.K. DeValois. "Neural coding of color." In *Handbook of Perception*, ed. E.D. Carterette and M.P. Friedman, 5, 117-166. New York: Academic Press, 1975.

DeValois, R. L., and K.K. DeValois. "A multi-stage color model." *Vision Res.* 33 (1993): 1053-1065.

Ditchburn, R. W. *Eye-movements and visual perception*. Oxford: Clarendon Press, 1973.

Donner, K. "Noise and the absolute thresholds of cone and rod vision." *Vision Res* 32, no. 5 (1992): 853-66.

Eisner, A., and D.I.A. MacLeod. "Blue sensitive cones do not contribute to luminance." *Journal of the Optical Society of America* 70, no. 1 (1980): 121-123.

Fechner, Gustav Theodor. *Elemente der Psychophysik*. Leipzig,: Breitkopf und Härtel, 1860.

Field, D. J. "Relations between the statistics of natural images and the response properties of cortical cells." *Journal of the Optical Society of America A* 4 (1987): 2379-2394.

Fukurotani, K. "Color information coding of horizontal cell responses in fish retina." *Col. Res. Appl.* 7 (1982): 146-148.

Gregory, R. L., and P. Heard. "Border locking and the Cafe Wall illusion." *Perception* 8, no. 4 (1979): 365-80.

Kaplan, E., and R. M. Shapley. "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 (1986): 2755-2757.

Kohonen, Teuvo. *Self-organization and associative memory*. 3rd ed. Springer series in information sciences ; 8. Berlin ; New York: Springer-Verlag, 1989.

Körner, T. W. *The pleasures of counting*. Cambridge, UK ; New York: Cambridge University Press, 1996.

Krauskopf, J., and K. Gegenfurtner. "Color discrimination and adaptation." *Vision Res* 32, no. 11 (1992): 2165-75.

Land, E. H. *Scientific American* 52 (1964): 247.

Larimer, J., D.H. Krantz, and C.M. Cicerone. "Opponent-process additivity -- I. Red/green equilibria." *Vision Research* 14 (1974): 1127-1140.

Laughlin, S.B. "Matching coding to scenes to enhance efficiency." In *Biological processing of images*, ed. O.J. Braddick and A.C. Sleigh, 42-52. Berlin: Springer Verlag, 1983.

Le Grand, Y. "Les seuils differentiells de couleurs dans la theorie de Young." *Revue d'Optique* 28 (1949): 261-278.

Lee, B. B., J. Pokorny, V. C. Smith, P. R. Martin, and A. Valberg. "Luminance and chromatic modulation sensitivity of macaque ganglion cells and human observers." *J Opt Soc Am [A]* 7, no. 12 (1990): 2223-36.

Lee, B. B., C. Wehrhahn, G. Westheimer, and J. Kremers. "Macaque ganglion cell responses to stimuli that elicit hyperacuity in man: detection of small displacements." *J Neurosci* 13, no. 3 (1993): 1001-9.

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

Levine, M. W., B. G. Cleland, and R. P. Zimmerman. "Variability of responses of cat retinal ganglion cells." *Vis Neurosci* 8, no. 3 (1992): 277-9.

Levine, M. W., R. P. Zimmerman, and V. Carrion-Carire. "Variability in responses of retinal ganglion cells." *J Opt Soc Am [A]* 5, no. 4 (1988): 593-7.

Livingstone, M.S., and D.H. Hubel. "Psychophysical evidence for separate channels for the perception of form, color, movement, and depth." *Journal of Neuroscience* 7, no. 11 (1987): 3416-3468.

Luther, R. "Auf dem gebeit der farbreizmetrik." *Z. Tech. Phys.* 8 (1927): 540-558.

MacLeod, D.I.A., and R.M. Boynton. "Chromaticity diagram showing cone excitation by stimuli of equal luminance." *Journal of the Optical Society of America* 69, no. 8 (1979): 1183-1186.

Marr, D. "The computation of lightness by the primate retina." *Vision Research* 14, no. 12 (1974): 1377-1388.

McMahon, M.J., and D.I.A. MacLeod. "Dichromatic vision at high light levels: red/green discrimination using the blue-sensitive mechanism." *Vision Res.* 38 (1998): 973-983.

Miyahara, E., V. C. Smith, and J. Pokorny. "How surrounds affect chromaticity discrimination." *J Opt Soc Am [A]* 10, no. 4 (1993): 545-53.

Morgan, M.J., A. Adam, and J.D. Mollon. "Dichromats detect colour-camouflaged objects that are not detected by trichromats." *Proc. R. Soc. Lond. B.* 248 (1992): 291-295.

Müller, G.E. *Darstellung und Erklarung der verschiedenen Typen der Farbenblindheit.* Gottingen: Vandenhoeck-Ruprecht, 1924.

Osorio, D., and T. R. J. Bossomaier. "Human Cone-Pigment Spectral Sensitivities and the Reflectances of Natural Surfaces." *Biological Cybernetics* 67, no. 3 (1992): 217-222.

Ritter, H. J., T. M. Martinetz, and K. J. Schulten. "Topology-Conserving Maps For Learning Visuo-Motor-Coordination." *Neural Networks* 2, no. 3 (1989): 159-168.

Ruderman, D.L., T.W. Cronin, and C.C. Chiao. "Statistics of cone responses to natural images: implications for visual coding." *J. Opt. Soc. Am. A* 15 (1998): 2036-2045.

Savage, C. "A survey of combinational Gray codes." *Siam Review* 39, no. 4 (1997): 605-629.

Stockman, A., D.I.A. MacLeod, and N.E. Johnson. "Spectral sensitivities of the human cones." *J. Opt. Soc. Am. A* 10 (1993): 2491-2521.

Thornton, J. E., and E. N. Pugh, Jr. "Red/Green color opponency at detection threshold." *Science* 219, no. 4581 (1983): 191-3.

Tolhurst, D. J., J. A. Movshon, and A. F. Dean. "The statistical reliability of signals in single neurons in cat and monkey visual cortex." *Vision Research* 23, no. 8 (1983): 775-786.

Wachtler, T., C. Wehrhahn, and B. B. Lee. "A simple model of human foveal ganglion cell responses to hyperacuity stimuli." *J Comput Neurosci* 3, no. 1 (1996): 73-82.

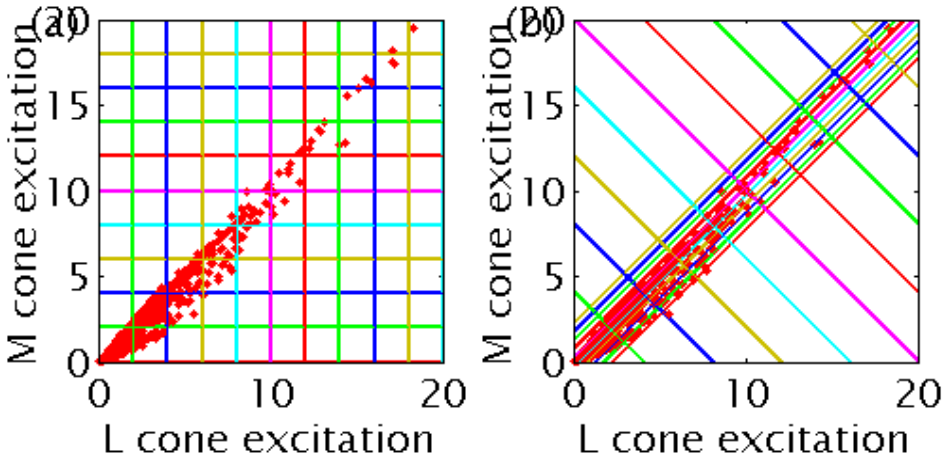

Figure 1: A scatter plot for natural surface colour stimuli measured by R. O. Brown in the (L,M) plane, divided by a grid into (a) 10 distinguishable levels of L and of M cone excitation; (b) 10 levels of L+M and 10 levels of L - M. (In this plot the values of M have been scaled up by 2.5 relative to a luminance basis, hence true constant luminance contours deviate from the negative diagonal shown.)

:

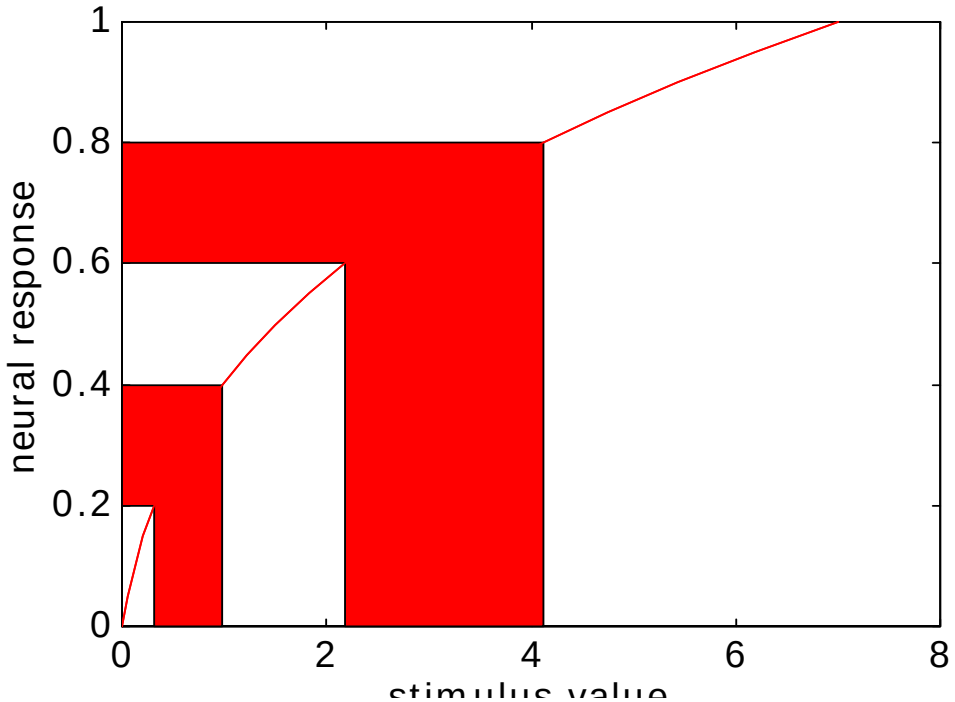

Figure 2. Compressively nonlinear function for neural firing rate vs. input stimulus value. Random fluctuations in firing rate occur with constant standard deviation σ, spanning bands of equal height. The associated RMS errors in the estimation of the input value are shown by the widths of the horizontal bands, and vary in inverse proportion to the derivative of the compressive response function.

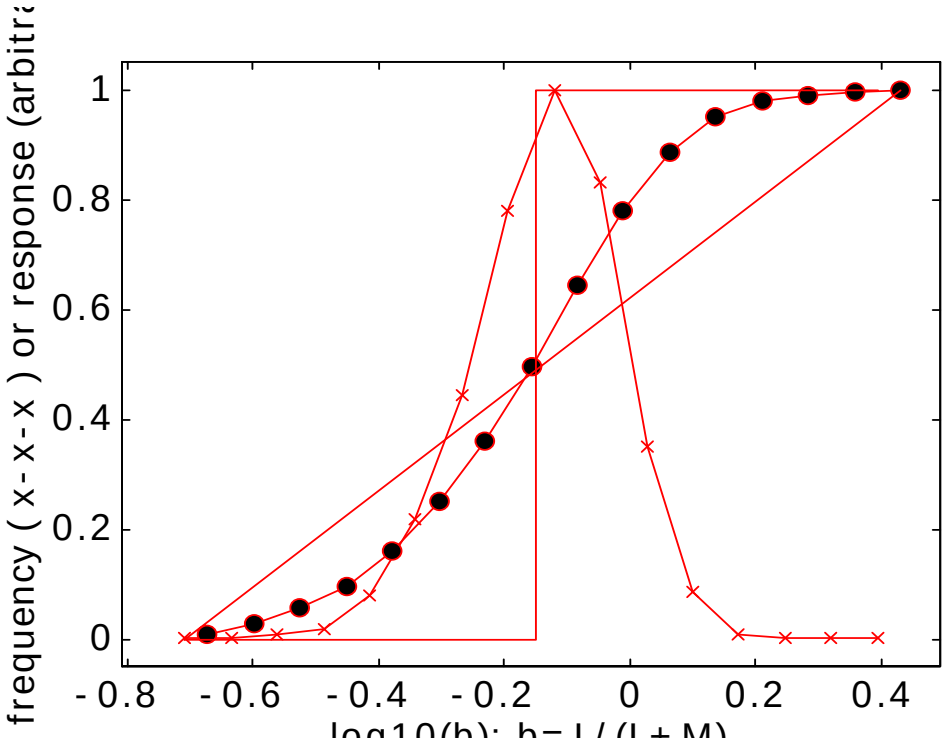

Figure 3. Crosses: frequency distribution of log10(*b*) for Ruderman et al.'s set of natural colours; *b* specifies S cone excitation per unit luminance, i.e. *b* = S/(L+M). Whites and greens are near the middle of the distribution, with equal energy white at log10(*b*)=0. To the right lie bluish colours; to the left, generally yellowish or reddish ones. Candidate input-outputs functions: pleistochrome (circles), compared with linear and stepwise alternatives (curves).

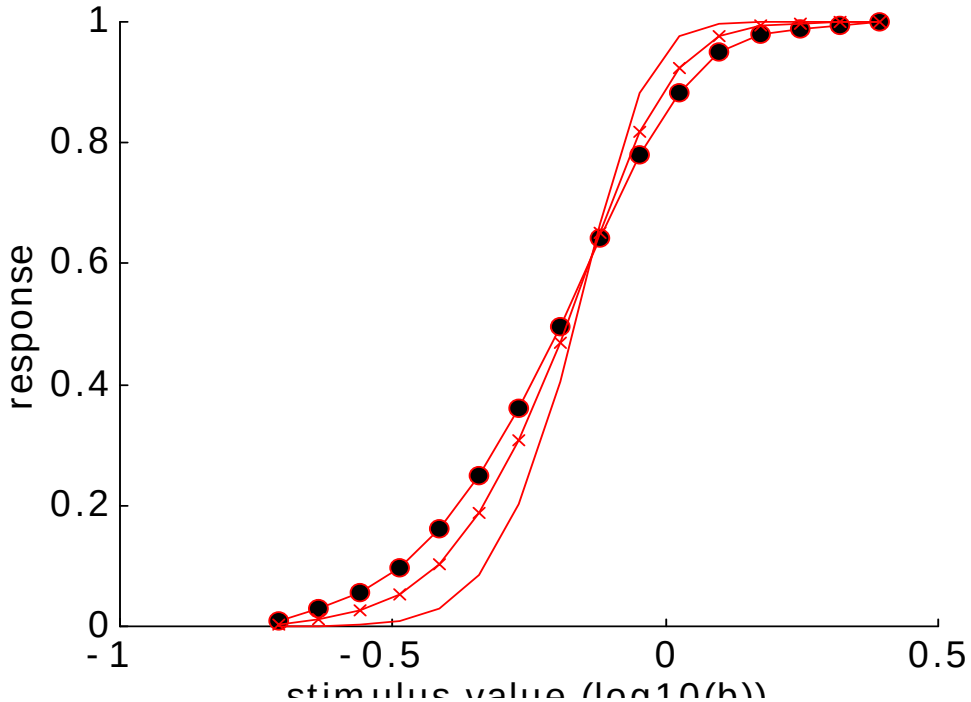

Figure 4. Pleistochromes based on minimization of mean squared error (circles), or of mean absolute error (crosses), compared with the function that maximizes mutual information through histogram equalization (plain curve).

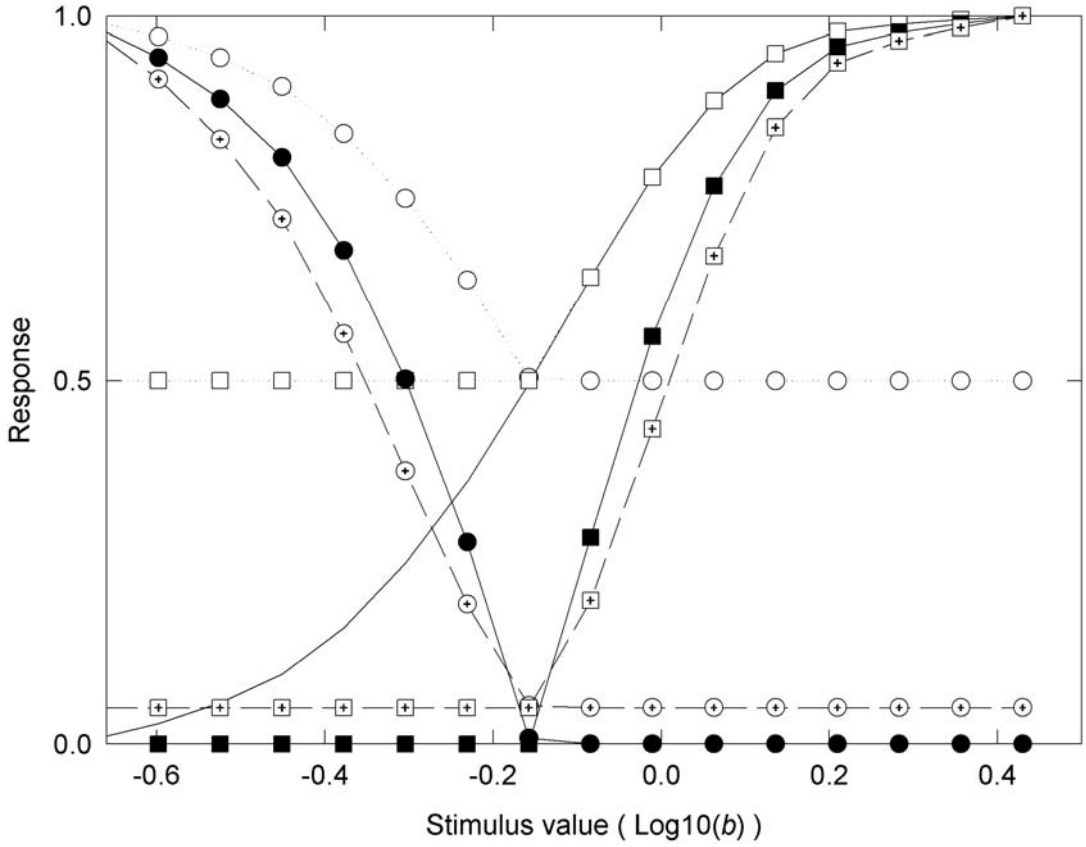

Figure 5. From single-neuron pleistochrome (continuous sigmoidal curve, from Fig. 3) to 2-neuron split range opponent code. Circles and squares show, for two rectifying neurons, how using the full output range for each neuron to cover only half the full input range ('split range' code) allows doubled response function gradients.

Open isolated circles and squares use only the upper half of the output range, with no gain in efficiency over the original sigmoidal neural response function. But using the full response range (filled circles and squares) allows a two-fold vertical expansion, hence doubled differential sensitivity. Dashed curves (with open circles and squares) show how the full-range response functions indicated by the filled symbols are modified when Poisson noise and spontaneous activity are assumed.

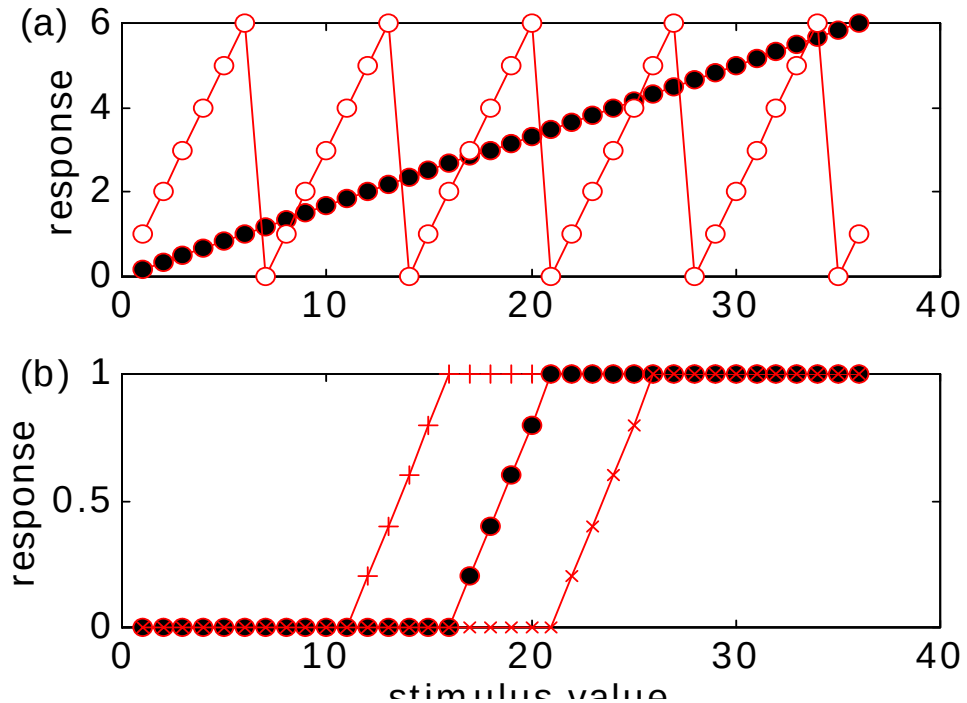

Figure 6. Other candidate multi-neuron encoding schemes. (a) Digital encoding, with one neuron (filled circles) for the more significant digit of the stimulus value, and a second neuron (open circles) for a second digit. (b) Split range encoding with n > 2; multiple neurons (pluses, circles, crosses) have monotonic, but staggered, response functions.

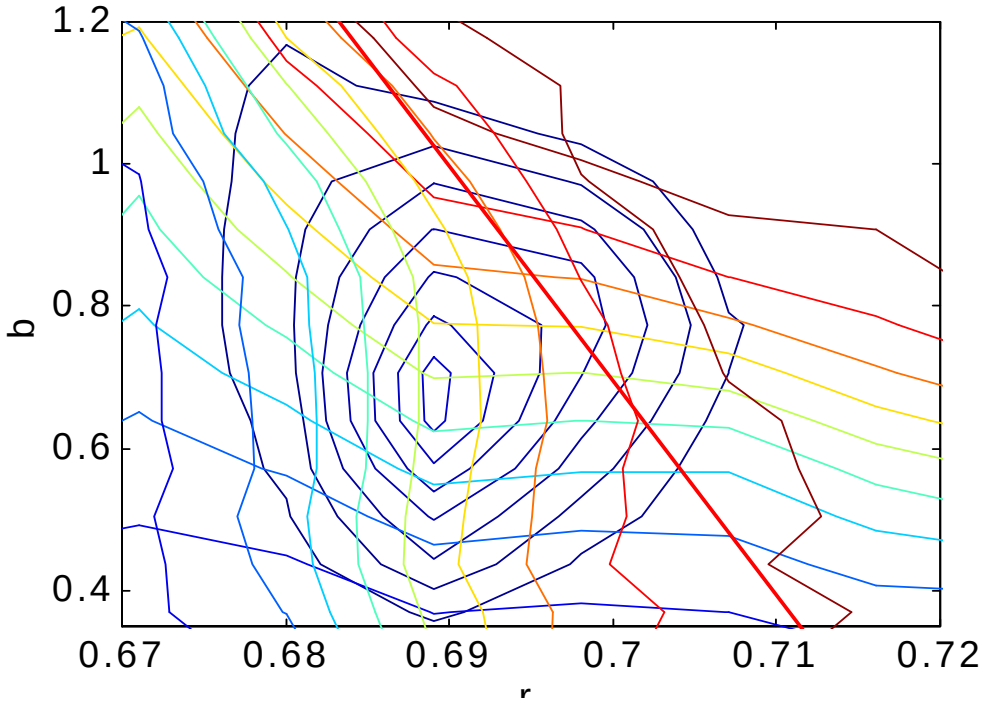

Figure 7. Constant luminance chromaticity diagram, with axes *r* = L/(L+M) and *b* = S/(L+M), with contours of Ruderman et al's distribution of colours of natural surfaces under D65 illumination. The equal energy white plots at *r*=0.70, *b* = 1.0. The straight line shows the locus of colours that are subjectively neither reddish nor greenish. The grid shows equally spaced constant-response contours for a two-neuron nonlinear code optimized for the natural colour distribution.

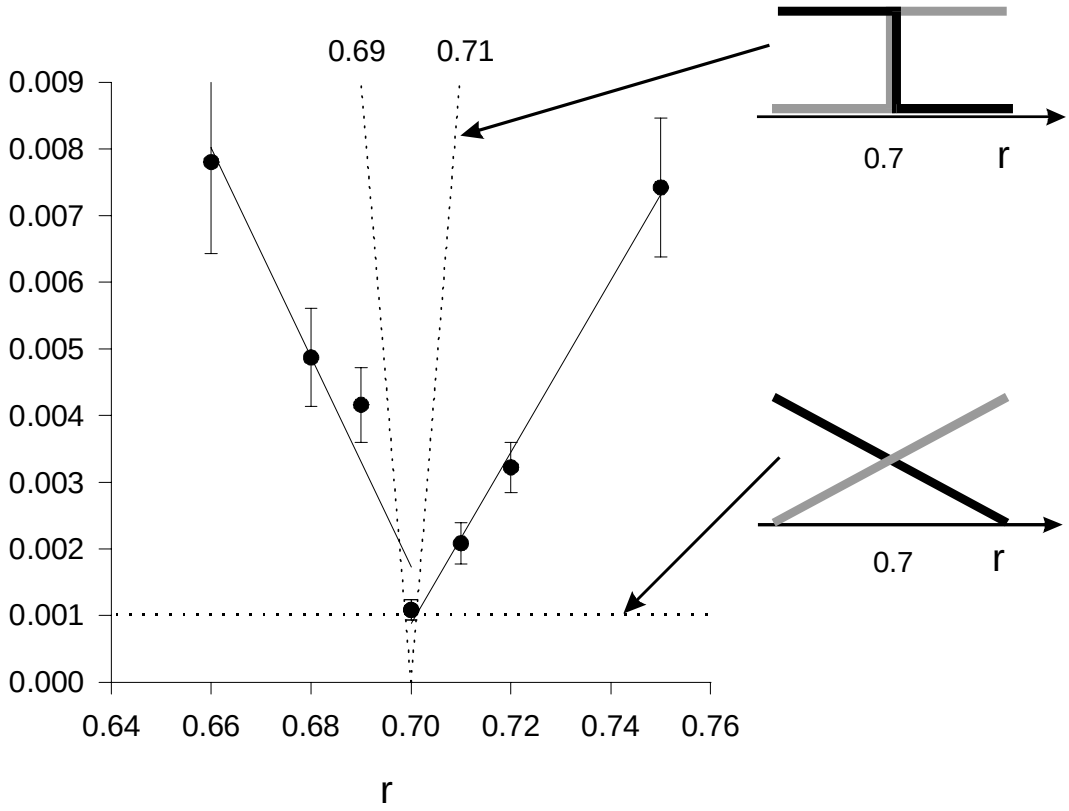

Figure 8 (a) Circles, with straight lines fit, show Δr, the difference in *r* =L/(L+M) just sufficient for 84% correct discrimination between isoluminant test and standard stimuli, as a function of value of *r* for the standard stimulus. The surround was an equal energy white, for which *r* = 0.70; hence 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) and for step function encoding (steep dashed line). The inset illustrates these encoding schemes.

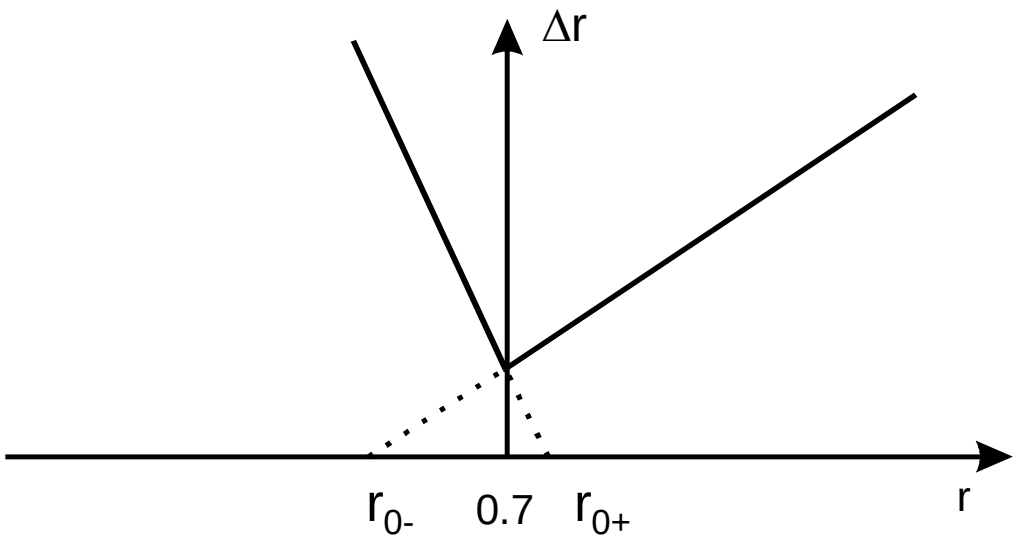

Threshold $\Delta r \sim dr/dN$

$dr/dN = |(r-r_0)/(0.7-r_0)|$

Neural signal $N = a\, ln|(r-r_0)/(0.7-r_0)|$

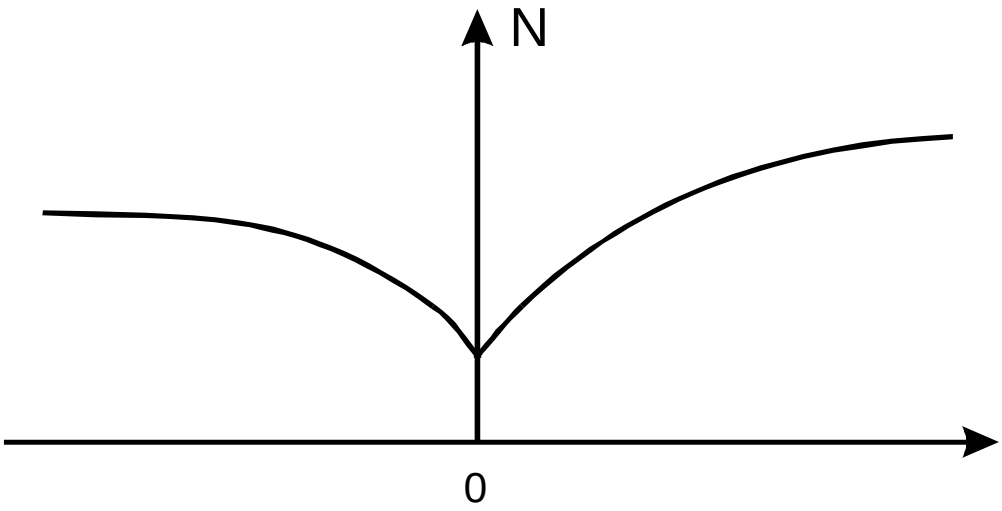

(b) Nonlinear response functions generated by the data of (a) on the basis of Fechner's integration (effectively assuming fixed output noise). The required value of the half-gradient chromaticity r_0 is roughly 2% in L cone contrast.