

Version of 2.5.99.

CONTRAST COLOURS

Paul Whittle

Department of Experimental Psychology, Cambridge University

pw109@cam.ac.uk

'So the judgements that we hold about the colours of objects seem not to depend uniquely on the absolute nature of the rays of light that paint the picture of objects on the retina; our judgements can be changed by the surroundings, and it is probable that we are influenced more by the ratio of some of the properties of the light rays than by the properties themselves, considered in an absolute manner.' Monge (1789, cited by Mollon, 1995)

'Hue is perhaps the direct perception of the direction of departures of the stimulus from some reference stimulus.' Evans (1974, p234)

'The brain forms colors by comparing objects to their background and not by analyzing their local spectral reflectance. An object is bright or dark, and of a particular color, only in relationship to its background.' Gouras & Zrenner (1981)

Introduction

Most of the science of early vision has made the 'contrast turn' in the last thirty years. That is, it has become generally accepted that *relative* stimulus magnitude, along whatever dimension, is more important than absolute, and that this is so because early stages of the visual pathway code information in terms of contrast.¹ Accounts of colour appearance, as opposed to colour discrimination, have lagged somewhat behind, or followed only in a zig-zag manner. Thus on the one hand I could cover pages with quotations like those above, which occur in the contexts of colour constancy, simultaneous contrast, adaptation, and retinal physiology. On the other hand, much colour science pays little attention to the surroundings of coloured objects. So much so that demonstrations of vivid contrast colours, such as coloured shadows or Land's two-colour projections, often amaze colour scientists as much as anyone. Yet it is 200 years since Goethe, 150 since Chevreul, 100 since Hering, to name only three of the best known figures who have written at length and eloquently on chromatic contrast effects. One would have thought that these observations could at last be assimilated and find their home in a visual science that has made the contrast turn. But would-be comprehensive treatments of colour appearance, such as Kaiser & Boynton's 1996 text or Abraham & Gordon's much cited 1994 review, are still being published which treat effects of contrast as minor side-effects.

This is a curious state of affairs, but it is not without reason, for the idea that colour is always perceived relative to its background is contradicted by the everyday observation that if you move an object against a variegated background, it is often hard to see any changes in its colour at all. In the jargon of colour science, 'simultaneous colour contrast' between neighbouring objects is often a very small effect in ordinary vision, as has been observed at least since Helmholtz.

¹I refer to the 'contrast turn' by analogy with the 'linguistic turn' in philosophy. I use 'contrast' imprecisely to mean 'relative stimulus magnitude', particular expressions being discussed later on. Contrast is a physical quantity. I refer to psychological contrast effects by longer phrases such as 'simultaneous contrast'.

I try in this chapter to understand this situation and to point to ways in which it might be resolved. I discuss conditions in which colour is a function of contrast, in the strong sense that patch and surround lights have equal and opposite influence, along all three dimensions of colour. If we represent colours in a three dimensional space, the perceived colour depends only on the direction and length of the vector from background to object colour, and not on its absolute position. We can conveniently call such colours ‘contrast colours’, using a term current also since Helmholtz. I argue that they reflect a stage of early visual processing through which all colour information passes. But contrast colours are seen clearly only under certain conditions. This can be understood in terms of two kinds of colour constancy that the visual system achieves, one with respect to illumination and the other with respect to neighbouring object colours. Contrast coding facilitates the first but impedes the second. Since both occur, the brain must be able to evaluate the contrast code differently depending on (or as part of) the parsing of a scene into surface colour, lighting and spatial layout. Different stimulus conditions show up different aspects of these processes, some showing contrast dependence, and some, independence.²

I find it helpful to think of contrast colours as analogous to stereo depth. Both reflect a relative or contrast code computed early in the visual pathway. It takes special situations to show them up in their purest form. In ordinary vision contrast combines with other determinants of colour, just as binocular disparity combines with other depth cues. This analogy can be pursued in considerable detail (as Brookes and Stevens, 1989, do for the brightness dimension). I want to draw attention here to one particular aspect. The apparently simple question ‘How far away is x?’ can be answered in many ways. For example, by giving an absolute judgment in meters, by a relative judgement (‘about five meters behind that house’), or by reaching with the hand (‘this far’). It is the same with colour. We can always see at least the colours of both object and lighting. But although this has been pointed out over and over again, we continually fall back into the assumption that each region of the visual field has one and only one colour. To understand contrast colours, however, we have to remain aware of the ambiguity of colour. This is a gain, not a loss. It helps us to remain aware of the true complexity of experience.

The discussion in this chapter is mostly concerned with uniform colours seen against larger uniform backgrounds. Both temporal and spatial interactions may be involved. I discuss their relative roles briefly later in the chapter.

Demonstrating contrast colours

Contrast colours can now be made and studied with great freedom by using a computer monitor and appropriate software. Set up a large uniform gray field as a background, and vary the colour of a small central region. Suppose it can be controlled in a three dimensional space with axes Red-Green, Blue-Yellow and Black-White, with the background gray at the origin. If you vary the central patch along the axes, those are the colours you see. Now change the background to some other colour, and vary the small central region along directions parallel to the axes. The remarkable thing is that it doesn’t make much difference what background colour you start with, provided the eye has been allowed to adapt to it: you still see Red-Green, Blue-Yellow and Black-White. The colours you see are contrast colours, dependent primarily on the *relative* rather than the absolute physical colours of patch and background. If you move along intermediate directions, the colour behaves as you would expect from vector addition. A direction between red and blue will look purple, one between green, yellow and black will look olive green, and so on. The relativity of the contrast colours can be enhanced by various tricks. A simple one is just to flash them, say 1s on, 2 off, with only the uniform background between flashes. Here is another more elaborate but very striking demonstration.

²I put forward an earlier version of arguments in this chapter for the intensive dimensions of colour—brightness and lightness—in Whittle (1994a,b). Here I develop and modify the argument for colour as a whole, present it less cluttered with psychophysical detail, and review some new evidence in the colour domain.

Arrange two rows of eight coloured patches as in Fig 1a, with one row and its background seen by each eye and the two superimposed binocularly so that one row is seen above the other, apparently on the same background.³ Choose the colours of one row to form a circle in the red-orange region of a chromaticity diagram, and those of the other row a circle in the blue-violet region, as in Fig 1b. Make both backgrounds a neutral colour (N) mid-way between the two circles. You now have a row of reds and oranges above a row of blues and violets, against a neutral background. Within each row the colours are quite similar, and the two groups of colours completely disjunct. Now comes the crucial move: change each background colour to the chromaticity at the centre of its corresponding circle (A and B). The backgrounds can be brighter or darker than the patches, provided the luminance contrasts of both rows are the same. The transformation in the apparent colours of the patches, which are physically unchanged, is rapid and startling. Each row now contains a complete gamut of hues—red, purple, blue, green, yellow and orange—and the two rows more or less match, each ‘red-orange’ with its corresponding ‘blue-violet’.

All that is needed to produce the full gamut of hues is to choose background colours so that the vectors from background to patches point in all directions, in any chromaticity diagram. The matches, however, are more satisfactory in some colour spaces than others. The diagram of Fig 1b, in which they are good, is a surface in the three-dimensional space whose axes are LogL, LogM and Log S, where L, M and S are the quantum catches of the long- medium- and short-wave cones. In that space, within quite wide limits, the colour of a light is constant if the vector from background to patch is constant (Fig 2).⁴ The components of the vector are $[\text{Log}L/L_b]$ etc, where the subscript b denotes background. These are the ‘cone contrasts’ for the three types of cone. The vectors are the same in the two rows of colours, so the fact that the rows match shows that colours match if the three corresponding cone contrasts are equal ($L/L_b = L'/L'_b$, etc). I call this the ‘cone contrast rule.’ It is yet another illustration of the Weberian principle that sensation depends on *ratios* of stimulus magnitudes, not on the absolute values. It also supports von Kries’s Law of Coefficients (von Kries, 1905), which asserts that any state of chromatic adaptation alters the sensitivities of the three types of cone simply by three multiplicative factors. In this case the sensitivity coefficients are inversely proportional to the effects of the background on each cone type.

There is nothing special about using red and blue patches. The circles of colours can be anywhere in the chromaticity diagram. But as for all perceptual demonstrations, certain conditions have to be satisfied. In this case the most important is that the patches are spatially separate so that each is entirely surrounded by the background colour. The effect is enhanced by allowing each eye to adapt fully to the background, and presenting the colours in such a way that the background colours cannot be compared. The trick of binocular superimposition is just one way of ensuring this. Other tricks are presenting the backgrounds as a stabilised image (Larimer and Piantanida, 1988), blurring the edge between them (Wuerger, 1996), or using only temporal contrast so that the backgrounds are not separately visible during the matching (Webster and Mollon, 1995).

Since coloured shadows are the subset of contrast colours in which the light in the patch is obtained by blocking out some of the surround light, their hues can therefore also be predicted by the direction of the vector from background to patch in a chromaticity diagram. This was pointed out by Evans in his posthumously published book, *The Perception of Color*:

As far as hue is concerned these effects [colored shadows] can all be predicted by a single generalization as startling in its implications as in its predictions: if the chromaticities of the two sources are plotted on the CIE diagram and the line connecting them extended to the spectrum locus in both directions, the intercepts of this line indicate the wavelengths (or complementary wavelengths) that, seen as *isolated* colors, would be of roughly the same hues as the two patches,

³The binocular superimposition can be achieved by altering eye vergence or by using mirrors.

⁴This is approximately but not exactly true for a vector translated in the logarithmic MacLeod-Boynton diagram of Fig 1b.

regardless of where the source points lie on the diagram. Since the mixture chromaticity of the two sources lies between them on this same line, this leads to a definition of the complementary with respect to this mixture point. (Evans, 1974, p222)

Notice that Evans finds this 'startling.' His book is a striking example of the claims I made in my opening paragraphs. It is all about 'related' colours, that is, colours seen in an illuminated surround and strongly affected by it, or what I call contrast colours. Yet Evans had not made the contrast turn. The word 'contrast' is not indexed, and he always described his stimuli in terms of luminance and wavelength, never contrast.

We can therefore say that in some situations any light can be made to appear any colour whatsoever by choosing a suitable background, provided that the background light is available within the gamut offered by the physical set-up and the limitations of the eye.⁵ This striking characteristic of our colour vision must surely be of profound functional importance.

Experimental evidence for the cone contrast rule

Chichilnisky and Wandell (1995) reported the results of hundreds of matches in a haploscopic set-up similar to Fig 1a, though with only a single patch shown to each eye. This is a type of 'asymmetric matching' between lights seen with eyes or retinal regions in different states of adaptation. I will call displays like Fig 1a 'haploscopic superimposed displays', or 'HSDs'. They were introduced by Hering (1890). Chichilnisky and Wandell's data obeyed the cone contrast rule to a good approximation. They thus provided the most direct and fullest confirmation of the thesis that colour depended on the three cone contrasts rather than simply on the numbers of photons absorbed by each class of cone.⁶

Various precursors of this idea have been proposed over the past 150 years. These have varied in their choice of colour dimensions (CIE dimensions, cone channels, retinexes), in the contrast expressions (ratios, delta-contrasts etc), and in their context (adaptation, constancy, simultaneous contrast). An early one is Rollett (1867; colour systems, contrast colours). McDougall (1901) argued that contrast colours depend on interactions within the three (R,G,B) colour systems rather than on Hering's opponent colour mechanisms. Von Kries's (1905) statement of the coefficient rule is the best known older one (cones, ratios, adaptation). Ives (1912) applied the von Kries rule to colour constancy. Spencer (1943) proposed a revision of colorimetry to make it more compatible with colour appearances, because 'the trichromatic system ignores adaptation.' She expressed colours as tensors in CIE XYZ space, which were equivalent to vectors whose components were the three delta-contrasts $\Delta X/X_b$ etc, and showed that her scheme fitted the data of Helson (1938). This work was described in Le Grand's (1948) influential text. Both Spencer and Le Grand saw it as subsuming all of adaptation, simultaneous contrast and constancy. Alpern (1964) suggested that simultaneous colour contrast could be explained by interactions within each Stiles 1 -mechanism. The 'retinex' theory of Land and McCann (1971) explained colour constancy in terms of the computation of image contrasts in each cone mechanism. Walraven (1976, 1981) showed that contrast colours could be described by delta contrasts computed within each 1 -mechanism. Mausfeld and Niederžse (1993) proposed a general scheme for colour coding in terms of the three cone delta contrasts. Webster and Mollon (1995) showed von Kries transformations of colours in chromatic adaptation.

It is striking how neglected most of the early papers are. Rollet (1867) was known by McDougall (1901) who had independently reached the same conclusions, reviewed by Rivers (1900) and Tschermak (1903), but then

⁵A calculation set as an exercise for students in Le Grand (1948, p462, Ex 23).

⁶It is of interest in this connection that Evans (1974) also wrote: 'What's needed are 'inter-eye comparison methods but these have been considered respectable only in quite recent years. I am sure that a systematic study of perceived hue by this technique wd unravel most of the mysteries of color perception.' p122. Asymmetric matching between the eyes exploits the fact that the states of adaptation of the eyes are almost completely independent. Matching in the HSD is particularly easy and reliable, for reasons discussed later.

forgotten, at any rate in the English language literature.⁷ Ives's (1912) paper was neglected (see Brill, 1995). Spencer's (1943) work was described by Le Grand, but I have come across no other citation. Apparently these were ideas whose time had not yet come. I discuss why in a later section.

This brief history shows the same basic idea repeatedly being put forward. This is the Weberian relativity of colour perception shown by constancy and contrast colours. These are robust phenomena, and can be described in various ways. The description must contain ratios between light and background, but given that, the precise choices of expressions and variables are for many purposes unimportant. But we can also ask, as many have, what is the most accurate mathematical description? and where exactly in the visual pathway is the relativity imposed?

With regard to the mathematical description, an important distinction is between ratios and contrast expressions that also involve differences. This concerns what is meant by 'contrast'. Weber contrast and Michelson contrast are the commonest expressions for it, not the simple ratio L/L_b . Weber contrast = $(L - L_b)/L_b$, usually written $\Delta E/L_b$. When the backgrounds are weak a 'dark light' constant must be included: $\Delta E/(L_b + L_0)$. Michelson contrast = $|\Delta E|/(L + L_b)$. When people talk of contrast coding in the early stages of the visual pathway, they usually mean that the firing rate of sensory neurons is a function of, perhaps proportional to, Weber contrast or Michelson contrast. Such a code combines differencing (calculating $L - L_b$) with normalisation (attenuating by some function of the absolute stimulus level). The two components probably serve different functions. Normalisation copes with the wide range of illuminations under which we can see, and contributes to constancy with respect to illumination changes both in intensity and colour. Differencing seems to be associated with enhanced discrimination about an adaptation level.⁸

What is the evidence for differencing? The most direct is that many responses fall to zero when the difference is zero. The object disappears or the firing rate of a neuron is at resting level. Other bits of evidence come from the analysis of psychophysical data. Walraven (1976, 1981) provided one. His subjects saw a test patch made of superimposed red and green lights, on various coloured backgrounds. When they set the ratio of red to green to produce 'unique' yellow, they kept the ratio of $1/4$ to $1/5$ (effectively M and L cones) Weber contrasts constant. Here the differences were important; simple ratios L/L_b and M/M_b didn't do the trick. Another line of evidence comes from studying the enhanced discriminability (or rapid change of appearance) around the background colour, described by von Bezold (1874) and called by Takasaki (1966, 1967) the Crispening Effect. For luminance, the form of the data imply that differences are computed (Semmelroth, 1970; Whittle, 1992, 1994a). The story is not yet fully worked out for colour, but the fact that there is also a chromatic Crispening Effect (Takasaki, 1967; Ovenston and Whittle, 1996), suggests a comparable importance for chromatic differences. The implications of matching experiments obeying the cone contrast rule depend on the values of the dark light constants L_0, M_0, S_0 . If these are zero, matches do not distinguish equating ratios from equating Weber (or Michelson) contrasts, because those expressions are equal if the ratios are equal and vice versa. Chichilnisky et al (1995) fit their data with L_0 etc > 0 . I am not clear whether the constants were sufficiently large to provide strong evidence for differencing, though they argue on other grounds (p251) that differential stimuli were being equated.

Walraven's result had a remarkable implication, which I discuss briefly both because it is taken up later, and because it gave rise to a controversy which seemed to impugn the validity of the cone contrast story. His data implied that the visual system was comparing only the cone *increments*, ΔE and $\Delta E/M$. But these were superimposed on bright red or green backgrounds, which of course added physically to them and could be clearly seen. Nevertheless neither the physical addition nor the subjective appearance seemed to influence

⁷Tschermak's long review, with 180 references, is a valuable and little known resource.

⁸It is tempting to associate them with the subtractive and multiplicative components of adaptation that have been demonstrated by many workers (see Walraven, Enroth-Cugell, Hood, MacLeod and Schnapf, 1990), but I doubt whether this subtractive component does actually play the required role (Whittle, 1994a, p104).

subjects' settings. It was as though the background was subtracted out by the visual system and influenced the colour of superimposed stimuli only by setting the adaptational state of the retina. Walraven called this 'discounting the background', echoing the older description of colour constancy as 'discounting the illuminant'.

This finding did not go uncontested. Shevell (1978), using the same technique, found that there was also an additive effect of the background hue, particularly when the patch was near threshold. But the additive effect seems, rather surprisingly, to depend on the subjective appearance of the background, not on its physical addition to the test patch. This is implied by an ingenious experiment by Nerger, Piantanida and Larimer (1993). They put an annulus of a different colour round the background and stabilised its inner edge on the retina. When this stabilized edge faded, as such edges always do, the inner background disappeared and its colour was filled in by the colour of the annulus. With an ordinary long-wavelength red background, they found Shevell's additive effect. But when this background was made to appear yellow by filling-in, the effect disappeared. Therefore, it depended on the *appearance* of the background, not its physical composition, which was unchanged by stabilisation.

The measurements of Chichilnisky and Wandell (1995) strengthened the evidence for the cone contrast rule in two ways. First, they were able to explore more of the colour domain than Walraven, because they were not restricted to 'unique' hues. Second, the HSD produced particularly good cone contrast matching. If Nerger et al are right, one reason was that it prevented subjects seeing the separate monocular backgrounds, thus removing Shevell's additive effect.⁹

Although equating cone-contrasts provides a good first order description of contrast colours, it is not the whole story even in situations like the HSD that favour it. Whittle and Arend (1991) adjusted a patch in a grey surround in the HSD to match 'homochromatic' standards. These were patches which differed from the background only in luminance, though both were intensely coloured, being produced by single monitor phosphors, red, green or blue. When there is only a luminance difference, the cone contrasts are all equal to the luminance contrast, so a matching patch should be set with the same characteristics, differing from its grey background only in luminance. This was true for weak decrements, but increments were set to a desaturated version of the coloured background, and higher contrast decrements to a colour approximately complementary. This resembles the Helson-Judd effect for judgements of colours under chromatic illumination (Helson, 1938; Judd, 1940). Chichilnisky and Wandell (1996) found evidence for the same effect in experiments in which subjects set a patch to achromatic in various coloured backgrounds. They also found evidence for adaptation in opponent colour mechanisms, the stage beyond 'receptor adaptation', which was their preferred description for what I call cone contrast matching. Finally, Chichilnisky and Wandell (1997) admitted that they also found small differences between increments and decrements in the experiments of their 1995 paper. All these results imply deviations from exact cone contrast matching.

There is another type of experiment in which patches were set to neutral in an illuminated surround, whose results cannot be explained in terms of keeping the ratios of cone contrasts constant. These are the experiments on the black threshold by Werner and colleagues (see, for example, Shinomori, Scheffrin and Werner, 1997). They found the luminance at which lights of different wavelength in a constant surround became black, and also the luminance of surround lights which induced black in a constant patch. The spectral sensitivities were quite different, contradicting the symmetry implied by a cone-contrast scheme. The sensitivity in the patch showed the multiple peaks characteristic of opponent colour mechanisms, whereas that in the surround followed the single-peaked luminosity function.

⁹This has been confirmed in my lab by Nick Blaker, in a more direct comparison of matching in the HSD with setting equilibrium hues (Blaker, 1997).

Cone contrasts and opponent colours

The notion of colour opponency was partly motivated by contrast colours, and indeed it could have been designed expressly for them. They have an ‘opponent’ structure which is more general and in one way more clearly defined than the opponency of ‘unrelated’ colours in a dark surround. If an unrelated colour is varied along a line through neutral in a chromaticity diagram, the series will contain just two hues, such as green and magenta, divided by the neutral point into two saturation series: saturated magenta to no hue, and no hue to saturated green.¹⁰ The magenta and green are alternatives, ‘opponent’ colours. They can never be seen together as components of a mixture, as red and yellow can be seen in orange. The neutral point of the magenta-green series can be judged reliably, but it takes care. If the line is tilted out of the chromaticity (constant luminance) plane a black-white component is added. The colours might now run from bright magenta to dark green. If the line is vertical—black-white—the division point becomes arbitrary.

Contrast colours generalise this structure in that the background provides a balance point which can be any colour, not just neutral. If the eye is adapted to it then any line through any background colour contains just two categorically different colours, such as bright magenta and dark green, dividing at the background. The balance point is easy to judge and never arbitrary. It is where patch and background are equal: zero contrast, no object. The background colour may be visible, but that belongs to a different object, not in the same continuum.

It would therefore be somewhat paradoxical if contrast colours were to be explained in terms of cone contrasts rather than the opponent colour mechanisms that phenomenology and physiology imply. But what is the relationship between these two frameworks? The question is sometimes mentioned in the literature (eg Mollon, 1987, p35; Wyszecki & Stiles, 1967, p556), but is remarkably little discussed considering the long history of the use of both frameworks. One of the interests of contrast colours is that they open it up.

Two different types of colour space have been mentioned in passing. One in which the axes represent some quantity—quantum catch, contrast or whatever—corresponding to each cone type, and one in which two axes represent opponent chromatic dimensions and the third axis the light-dark dimension. A popular version of the latter type is the ‘DKL’ space proposed by Derrington, Krauskopf and Lennie (1984), who found it a convenient framework within which to represent the responses in different classes of retinal ganglion cell. The x and y axes are the MacLeod-Boynton axes $L/(L+M)$ and $S/(L+M)$, and the z axis is luminance $(L+M)$. What is the relationship between these two types of colour space, and which best represents contrast colours? We are of course interested in functional localisation as well as representation. Just as we can say that metameric matching is determined by the photopigments, we would like to know what level of the visual pathway determines contrast colours.

Since contrast colours match, to a first approximation, if their three cone-contrasts are all equal—the cone contrast rule—one might expect a cone-contrast space, with axes $\Delta E/L$, ΔB etc, to be a good representation for them. But the DKL space also has advantages. Both represent colours as vectors from a background, but the DKL space represents colour appearance more intuitively. It segregates brightness, corresponding to the vertical luminance axis, from chromaticity, represented in horizontal planes. Within any horizontal plane, polar coordinates r, θ map onto saturation and hue.

The planes and axes of cone-contrast space, on the other hand, do not map conveniently onto familiar dimensions of colour. Brightness corresponds not to an axis, but to the diagonal through the points (1,1,1) and (-1,-1,-1). None of the axes map simply onto hue. For example, although varying S-cone contrast varies hue along a violet-chartreuse dimension, excellent violet or chartreuse can be produced with zero S-cone

¹⁰I ignore the curvature of constant-hue lines, and in the case of contrast colours, Shevell’s additive effect, which can give low contrast stimuli a tinge of the background colour.

contrast.¹¹ This follows immediately from the fact that the violet-chartreuse axis in the MacLeod-Boynton chromaticity diagram is the *ratio* $S/(L+M)$, so you can manipulate this dimension by changing $(L+M)$ with S constant just as well as by changing S . Similarly cherry is produced by M -cone decrements just as well as by L -cone increments, and vice versa for teal. The *relative* stimulation of cone types, which is what is computed by opponent colour mechanisms, maps more simply onto hue than the stimulation of separate cone types. This is well known, but it is necessary to remind ourselves of it in the present context.

Furthermore, there is empirical evidence that opponent colour mechanisms are involved in generating contrast colours. I mentioned Chichilnisky et al (1966) above. Poirson and Wandell (1993) matched a uniform region to the bars of square wave gratings of various frequencies. The latter were contrast colours in my sense because their hue would be determined as much by the spectral composition of neighbouring bars as by their own. Poirson et al showed that opponent colour spectral sensitivities could be derived from their data on the simple assumption of pattern-colour separability—that the relative chromatic sensitivities are independent of spatial frequency and vice versa.

It is not difficult to find a variant of the DKL space that can accommodate both cone contrast and opponent colour findings.¹² One possibility is to take the axes as $x = C_L - C_M$, $y = C_S - (C_L + C_M)/2$, $z = (C_L + C_M)/2$, where C_L , C_M and C_S are the Weber or Michelson cone contrasts for the L , M and S cones. As in the DKL space, the x axis depends on the relative L and M cone signals, although now they are contrast signals, and is independent of S . The z axis is a bright-dark dimension: the average of the L and M cone contrasts. The y axis still represents the S cone signal (now contrast) relative to the bright-dark dimension. These coordinates have several advantages for representing contrast colours. The background colour is at the origin $(0,0,0)$ because all cone contrasts are zero there. Two points coincide if their cone contrasts are equal: the cone contrast rule. Therefore graphs plotted on these axes immediately show the matches to be predicted from that rule, and the deviations from it. Finally, such coordinates express the increasingly accepted idea that most adaptation—the computation of contrast—occurs prior to the opponent stage. This is the ‘receptor adaptation’ of Chichilnisky et al, 1995. For recent evidence see He & MacLeod, 1998.

Differences of cone contrasts have been used extensively in the literature on chromatic discrimination (eg Friele, 1961, see Wyszecki & Stiles, 1967, p558; Vingrys and Mahon, 1998).¹³ In that literature the C_s are usually Weber contrasts. For describing colour appearance, however, signed Michelson contrast $(L - L_b)/(L + L_b)$ has two major advantages. First, it provides a compressive transform of the S -cone axis that yields approximately equal subjective intervals.¹⁴ Whereas Weber contrast preserves the linear S scale which is a poor subjective scale because it compresses the chartreuse end relative to the blue violet end (Le Grand, 1949). We see this in the MacLeod Boynton diagram, for example, which puts neutral near the bottom, red-green, edge. Second, Fig 3a shows that the difference of Michelson contrasts behaves well as luminance contrast is varied with chromaticity constant. They have maximum magnitude near isoluminance, and drop to zero at high positive or negative contrasts. This accords with experience. Differences of Weber contrasts, on the other hand, increase monotonically from black, tending to infinity at high positive contrasts (Fig 3b), which is not at all how subjective hue and saturation behave. However, Michelson contrasts also have a

¹¹I refer to the colours of the MacLeod-Boynton axes as violet versus chartreuse and cherry versus teal (Abramov and Gordon, 1994) to emphasize the fact that the relation of these ‘geniculate’ axes to the ‘basic colours’ red, green, blue and yellow is a major unsolved problem.

¹² Simply making the axes logarithmic works quite well, though it does not give a convenient origin.

¹³ ‘For modest signals under a constant adaptation state, single-cell responses and psychophysical sensitivity are consistent with mechanisms that respond to simple sums or differences of the cone contrasts.’ Webster, 1996.

¹⁴ $(L - L_b)/(L + L_b)$ is approximately equal to $\text{Log}(L/L_b)$, for contrasts between about $\pm 80\%$.

major disadvantage. Consider a coloured light in a dark surround. All three Michelson cone contrasts are 1.0, so the chromatic x and y coordinates defined above are both zero. Yet such a light can look intensely saturated.¹⁵ How can this be? Using modified Weber contrasts incorporating 'dark light' constants, $\sqrt{L}/(L_b+L_0)$ etc, could avoid the paradox. If $L_0 = M_0 = S_0$, for example, with a dark surround, for which $L_b = M_b = S_b = 0$, the three cone contrasts would be proportional to L , M and S , and strong hues would be expected. These questions obviously remain very open. I offer these remarks mainly as an example of a level of discussion which I feel is needed, intermediate between analysing particular data sets and would-be comprehensive models of colour vision.¹⁶

If the matching rule in this space were just that x , y and z were each equated, this would be empirically indistinguishable from equating cone contrasts, because each would imply the other. But in fact results like the Helson-Judd effect found by Whittle and Arend (1991), and the other deviations from cone contrast matching discussed above, require postulating a linear mapping between the coordinates of the stimuli being matched, where the mapping coefficients may vary between octants (increment versus decrement, cherry versus teal, violet versus chartreuse). The focus of my current research is on determining these mappings. Note that the diagonal terms (constants a in $x=ax'$, etc) can be called 'contrast gain' values because the coordinates are contrast expressions. So this may allow a rapprochement between 'second site adaptation' and the 'contrast gain' of, for instance Chubb, Sperling and Solomon (1989), although the present 'contrast gains' are set by uniform adapting fields, not by contrast stimuli (see Webster, 1996).

One further point about the suggested space should be briefly mentioned because it raises the vexed question of the dimensionality of contrast colours (Evans, 1974; Mausfeld et al, 1993). Since three variables are required to express the visual effect of the surround, and another three that of the focal patch, this dimensionality could be as much as six. I proposed $(C_L+C_M)/2$ just now as the expression for the intensity dimension, because of the algebraic convenience of treating contrast colour as a function of only the three variables C_L , C_M and C_S . But if, for example, luminance contrast C_{L+M} , were to describe matches better, as is sometimes true for my data, then since this is not predictable from the three cone contrasts alone, no three-dimensional space would give a complete representation, though it might provide a useful approximate one.

Contrast colours form a natural space for colour appearance

This section is trying to articulate a hunch, a strong feeling that I have when working with contrast colours. The main component is that spaces like the DKL space or the variant of it just discussed are the *right representation* for this work in so many ways (notwithstanding the problem of dimensionality just raised). One feels at home in them. Sometimes I feel that they know more than I do: they seem to lead me into good experiments or fruitful analyses of the data. This is not a feeling that one often has in experimental psychology, so I think it is worth trying to articulate it, to unpack it a little. Furthermore, these spaces are also the right representation for the behaviour of primate retinal ganglion cells, and are closely related to the colour solids that express the structure of surface colours. This three-fold coincidence is surely of significance.¹⁷ I want to say 'surely such a space is the natural space for representing colour appearance.' This is in contrast particularly to spectrum-based representations or the CIE xy diagram, which are right for

¹⁵I am indebted to Donald MacLeod for this point. Note that the inverse stimulus, a dark patch in a saturated surround, has all three contrasts equal to -1.0, and looks black, as that would imply.

¹⁶General discussions like this of the merits of different opponent colour expressions seem remarkably rare, given the importance of the topic, and the fairly obvious points that can be made. For example, one still finds the simple difference $L-M$ being proposed. But this would vary with luminance from zero to infinity, which would be absurd for an index of chromaticity.

¹⁷One could speculate on evolutionary links between the items.

metamerism or other phenomena that depend on the visual pigments, but are clumsy for representing colour appearance. None of this is new, but we seem to have difficulty grasping the whole pattern.

This class of spaces is characterised by a vertical (z) axis representing intensity (luminance, luminance contrast, etc), and horizontal (xy) planes representing chromatic colour (hue, saturation, purity, etc) at constant intensity. The x and y axes are opponent-colour axes, loosely speaking. Any horizontal plane has colour circles containing all the hues. Colours can be represented as vectors from the reference colour at the origin, which is neutral physically or psychologically or both. Such a space accords with the phenomenology of colour, as Hering pointed out. Colours are all represented. There are no arbitrary absences like brown or the extra-spectral purples. Further, and most importantly, we can navigate easily in such a space. Subjects who are given a joystick connected so that its lateral movements take them round the colour circle, its front-back movements vary the radius of the circle (chromatic contrast), and which has two buttons for varying intensity, learn to make accurate three-dimensional colour matches within a few minutes. This is the most concrete way in which colours form a space: that we can find spatial representations which translate bodily movements into operations on colours.

The origin can be treated just as a computational reference point. But in work with contrast colours, it can represent the background colour, which is physically present. The vector from origin to colour is then physically realised as the contrast at the edge of the patch and is *seen* as a contrast colour. The demonstrations I described, and measurements like Chichilnisky and Wandell's, show that the background can be any colour. The eye adapts and makes it the current neutral. So these spaces accord particularly well with the phenomenology of *contrast* colour. The length of the vector represents a generalised contrast including both colour and luminance. It can be resolved into chromatic (saturation or colourfulness) and light-dark components. The azimuth represents hue, and the altitude luminance contrast. In this representation we can have negative light (decrements), which generalizes some colour algebra. The bipolar z axis dividing at the origin, where the object is absent (zero contrast), represents light-dark, or 'contrast brightness' (comprising lightness of surfaces and the brightness of light sources; Whittle, 1994a). In fact all straight lines through the origin represent just two (complementary or opponent) hues with their category boundary at the origin. Thus most qualities find a convenient representation in this space.¹⁸

Although the space accords well with colour appearance, the axes of course represent physically defined quantities, albeit 'anthropocentric' ones (Hilbert, 1987). One important move has been to make them reflect the biology of the visual system. That is the main difference between the DKL space and variants whose axes are also functions of the cone signals, and the three-dimensional spaces—colour solids—used in colour order systems (reviewed by Derefeldt, 1991). Spaces whose axes are physiologically meaningful are therefore powerful research tools in studying the relationships between physical, physiological and psychological variables.

I have tried to unpack the intuition of 'right representation'. It has at least the following components:-

Physical, physiological and subjective quantities are all represented.

The vertical axis and the horizontal chromaticity planes are easy to interpret (compare, say, the x axis of the CIE xy diagram).

All colours are represented.

Easy to navigate.

Allows flexibility, crucial for a research tool, in the choice of the precise quantities to plot on the axes. See previous section.

In my experience, though I cannot document it here, it is also serendipitous in suggesting new connections, and cleanly distinguishing different influences on data.

¹⁸An exception is the 'brilliance' dimension of Evans (1974), the bipolar dimension comprising grayness and fluorence.

It will have occurred to many readers that there is a dialectical relationship between much of what I am saying and the ubiquitous—in vision labs and out—current technology of colour CRT devices such as TVs and computer monitors. The spaces I have been discussing are particularly appropriate to them. (Though not only to them: Hering's opponent colour space antedates them by decades and colour solids antedate them by centuries. Deryfeldt lists a version by Forsius from as early as 1611.) It is no coincidence that the DKL space and the use of a CRT to present stimuli for colour research developed together. The CRT makes it particularly easy to generate and study contrast colours, to modulate colours about a mean and so on. For instance increments and decrements are equally easy to produce on a CRT, which is not at all the case using projectors or optical systems of lenses and mirrors. This ease is reflected in the bipolar intensity axis of the DKL space. Further, monitor colours are often textureless and always coplanar, which are features that enhance simultaneous contrast (e.g. Woodworth, 1938), that is, seeing contrast colours. Much of this chapter could be framed as a contribution to an emerging technology of colour representation suited to the CRT.¹⁹ This doesn't mean that the science is just technology in disguise. New technologies are successful only if they fit (and enlarge) our capacities, and they in turn allow us to better understand those capacities.

With regard to the importance of contrast colours, the argument of this section is as follows. First, these spaces reflect the structure of colour appearance and of a stage of early visual processing. Second, contrast colours exemplify the structure particularly clearly. They enable you to 'see' the vectors which for isolated surface colours are merely an ordering device. Therefore contrast colours are likely to be revealing something important.

Why contrast colours are not accepted as fundamental; two types of colour constancy

If one was acquainted only with the striking phenomena of contrast colours, such as the demonstrations described above, and their simple mathematical structure, the cone contrast rule, a form of the well known coefficient rule of von Kries, it would seem surprising that they have not been generally accepted as a basic fact of colour vision, on a par with metamerism. But in fact, in spite of strong statements like the epigraphs to this chapter, their significance has remained controversial and often unappreciated. This was vividly illustrated by the stir created by Edwin Land's two-colour projection demonstrations, which it was quickly shown could be predicted from known laws of chromatic contrast and adaptation (Judd, 1960; Walls, 1960; Wilson and Brocklebank, 1960). Only in a climate of ignorance of the behaviour of contrast colours could Land have made such a stir, and have believed that he was revolutionizing colour science. One could argue that this lack of acceptance is because the right experiments had not been done until recently, and that they were dependent on accurate knowledge of the cone spectral sensitivities, but I do not think those are the main reasons. The HSD and other ways of generating strong contrast colours were available in the nineteenth century, and Kšnig's or Fick's fundamentals were sufficiently accurate to allow Le Grand in 1949, to take just one example, to perform analyses of psychophysical data of lasting validity. Le Grand had also, in describing Spencer's (1943) work, stated clearly the dependence of contrast colours on the cone Weber contrasts (Le Grand, 1948, p222).

I think a much more important reason is that it has been and still is quite reasonably assumed that if there are strong effects of contrast on colour then they should be clearly visible whenever two colours are juxtaposed. However, that is simply not so. If you place saturated red and green papers side by side, and place a grey square on each, the squares do not generally acquire clear contrast colours. Similar illustrations of 'simultaneous colour contrast' in textbooks are often rather unconvincing.

Furthermore, if you move a coloured object above a variegated background, you will usually see little change in its colour as it moves over different colours of floor or furniture. This constancy of colour in the presence

¹⁹Indeed I fear that in colour vision labs which habitually use CRTs and the DKL space, much of what I am saying will seem to be just spelling out the obvious. My defences are that not everyone has such experience, and that the obvious usually merits more reflection.

of different surrounding surface colours or different backgrounds is a second type of colour constancy in addition to the classical type with respect to changes in illumination.

The most convincing evidence that there are two separate types of constancy, as opposed to a single complex algorithm that takes all the visual field into account, as in retinex theory, is provided by situations where the apparent colours change markedly as a result of changing only the perceptual interpretation, not the stimuli on the retina. An informal demonstration was described by Evans (1948, p166). Gilchrist, Delman and Jacobson, (1983) achieved a good approximation to the same situation. These were both in the lightness domain. See Whittle, 1994b, for a fuller account. Both kinds of constancy have recently been demonstrated for colour in a series of experiments by Brainard and colleagues. (Brainard, 1998; Brainard, Brunt and Spiegle, 1998). Subjects saw largish (4x6) coloured rectangles mounted on a wall whose illumination could be changed. They were not CRT simulations as in so many current studies. In one experiment, subjects set a rectangle to look achromatic under different colours of ambient illumination.²⁰ Their settings approximately tracked the illuminant, so that if it was red light, for example, the patch was set so that it sent red light to the eye. This is classical colour constancy: the patch looked grey when the eye received almost the same light that a grey surface would reflect under each ambient illuminant (80% constancy according to their index). But when the surroundings were varied in colour not by ambient illumination but by placing around the patch either a large coloured board (16x20) or a sheet which covered the whole wall, there was very little effect on the achromatic settings. The patch was kept close to neutral reflectance. This is constancy with respect to neighbouring objects.

Clearly if the arguments of this chapter are correct, that the dependence of colour on contrast is mediated by low-level, perhaps retinal, mechanisms, this second kind of constancy is very puzzling.

The classical type is one context in which the role of contrast in determining perceived colour has been most strongly asserted. The argument has been that since contrast is approximately independent of overall illumination, the constancy of perceived colour with respect to illumination, also approximate, would be explained if perceived colour depended on contrast. This argument is less straightforward for chromatic than for intensity changes, because the cone contrasts are not in fact at all independent of the spectral composition of illumination for arbitrary surfaces and illuminants—consider the extreme case of narrow-band reflectors and illuminants. However, they are approximately so for most actual surfaces and natural illuminants (Foster and Nascimento, 1994).

But while a contrast code is an ideal mechanism for maintaining constancy with respect to illumination, it ought to produce strong ‘simultaneous contrast’ effects (ie contrast colours) when the colours of surrounding objects are changed. But instead, we find the second type of constancy: an object is little affected by the colours of surrounding objects. There are two major puzzles here: first, how do we distinguish illumination changes from surrounding-object changes? Second, in the latter case, how is constancy maintained if a contrast code is all the brain receives?

We don’t know the answer to either question, and this is not the place to speculate at length. The distinction between the two situations might require a global parsing into objects, lighting and spatial layout (many phenomena show linkages between these three). But it is also possible that it could be done on the basis of relatively low level cues such as the structure of intersections, which can generate different percepts of transparency or opacity (Metelli, 1974; Gerbino, 1994).

The colours of objects under a changing illumination can be contrast colours, derived directly from a contrast code. But when it is the surrounding surfaces that are changing around an object, if the object’s colour is coded in terms of contrast with respect to those surfaces, then to recover the constant colour of the object, these changes in contrast must somehow be compensated for. This suggests a process of integration, and

²⁰The adjustment was by concealed projectors illuminating just the test rectangle, and so set up so that the rectangle always looked like a pigmented surface, except for its strange variability.

various ideas of ‘edge integration’ have been put forward, sometimes after sorting the edges into reflectance and illumination edges. See Gilchrist (1994a) for some discussion of these ideas. The constant of integration (the ‘anchoring problem’) may be supplied by an average of the colours in the scene: the ‘grey world’ assumption. Or perhaps the brain does not receive only contrast signals: there might be absolute information such controls pupil size, though I do not know of any evidence for that in the colour domain. Another possibility is that the brain receives only contrast signals, but that there are many types of them, transmitted by different retinal ganglion cell types (which is likely), and that these could somehow be solved to extract absolute information.²¹

At present, this is a major area of ignorance in our understanding of colour vision. It may be that only when these problems are solved will contrast colours really find a secure, and probably fundamental, place in colour science.

The important point in the present context, which it seems to me cannot be overstated, is that since we know from both everyday experience and experimental evidence that both these constancies exist, we cannot expect strong contrast colours (or von Kries adaptation) to show up every time we juxtapose different coloured lights or present lights in different states of adaptation. The outcome will depend on how it is done, and in a rather subtle way. Specifically, on whether the variations in surround colours *are seen as* changes in illumination or changes in neighbouring surface colours. When the surround colours cannot be seen, or at any rate cannot be compared, as in the HSD, this choice is unavailable, and subjects are constrained to respond only on the basis of physical contrast. But the stimulus displays commonly used are in effect abstract pictures which leave the perceptual interpretation indeterminate. The equivocal results of experiments on simultaneous contrast or on von Kries adaptation should be expected.

The fact is that asymmetric matching experiments in which the background colours can be seen and compared, as they cannot be in the HSD, often do not produce strong contrast colours. Varying the background colour generally has a much weaker effect than varying patch colour, so measurements do not follow the cone contrast rule. But this is not an invariant result. The only accurate generalisation about these experiments is that the results are variable. Sometimes they give results following the cone contrast rule (for example, some results of Lucassen, 1993), and sometimes not. There have been an immense number of such experiments, dating back at least to Kirschmann (1890). Wyszecki and Stiles (1982) and Wandell (1995) review some of them. The indeterminacy of their results is presumably one reason why each new arrival in the field feels they have to make their own measurements. But it is not enough to control only the spectral composition of patches and surrounds. New measurements will serve no purpose unless the *perceptual interpretation* of the surround colour is controlled.

Spatial versus temporal interactions

The concept of ‘contrast colours’ is short-hand for much of what is discussed under the headings of simultaneous and successive colour contrast, chromatic adaptation and chromatic induction.

It is often asked whether contrast colours are the product of spatial or temporal interactions or both. We need here to distinguish interactions between neighbouring or successive parts of the *stimulus*, from the neural mechanisms mediating them. For example, to produce a dark colour as a steady object rather than a transient appearance requires a bright surround. In that sense spatial interactions are required. But the effect of the surround could be mediated entirely by temporal neural interactions produced by eye-movements jiggling the image of the edge back and forth across the photoreceptors rather than by the spatially opponent colour receptive fields which we know exist.

Good contrast colours can certainly be produced by temporal stimulus and neural interactions alone, as is shown, for example, by Webster & Mollon (1995). The instantaneous appearance of coloured shadows and

²¹Eg if L/L_b and $\Delta EL/(L_b+L_0)$ are signalled, and the constant L_0 is known, then L and L_b can be derived.

other contrast effects is often cited as evidence that contrast colours can be produced by spatial neural interactions alone. However, I do not know of really convincing experimental evidence, for example from stabilized image experiments. The results of most studies could be due eye-movements plus temporal neural interactions.

One piece of evidence comes from the Crispening Effect (see above). For contrast *brightnesses* this effect is markedly dependent on the exact edge profile, suggesting that spatial interactions across the edge are important. The evidence is that the effect is abolished by a thin black ring round the edge of the patch (Whittle, 1992). However, a black ring makes no difference at all to the chromatic version of the effect (Ovenston, 1998). This suggests, though not conclusively, that spatial neural interactions may be more important for brightness than colour.

It should be noted that even if contrast colours are produced predominantly by temporal interactions rather than spatial, this cannot account for the weakness of simultaneous colour contrast with respect to neighbouring objects (ie to the second type of colour constancy). Retinal generation of chromatic contrast signals, whether by spatially opponent colour receptive fields or receptor adaptation or whatever, encodes the proximal stimulus, not the distal, and will therefore be indifferent to whether the neighbouring colour is a surface or an illumination colour.

Further complications: colour is not single-valued

The phenomenon of contrast colours has surprising implications. It is easy to say ‘situations under which colour is determined by contrast’, but if such situations are common in everyday life they radically upset common beliefs about colour. Take a large bright red field. Make a small region of it slightly darker, keeping its physical colour (the spectral distribution of energy) the same, so that the relative stimulation of the three cone classes is unchanged. Therefore the three cone contrasts are all the same, and equal to the luminance contrast. So according to the cone contrast rule it should match a patch of the same luminance contrast in a background of any other colour, including grey. Another way to put it is to say that since it has no *chromatic* contrast with respect to its surround, if colour is determined by chromatic contrast, it should look colourless, grey. But it was bright red a moment ago. Can the slight darkening or lightening have so dramatically changed its colour? We know this sort of discontinuity in colour appearance doesn’t happen. What we will have is the bright red field with a small region a bit darker or lighter, but still red. Yet in the HSD it does happen, in the sense that the small region turns out to match a grey patch in a grey surround (Whittle et al, 1991). The matches seem to imply a discontinuity that you do not see in the colour as normally viewed. But in fact there is a corresponding discontinuity in normal vision. The slight darkening gives birth to *totally new ways to see* the small region: for example as a grey shadow on the red background, or as a grey object in a lighter grey surround both illuminated by intense red light. Contrast creates boundaries, and it is boundaries that organise the field and allow different ways of seeing: create ‘the structure “lighting-object lighted”’ (Merleau-Ponty, 1945, p307 of English translation). In this sense it can suddenly ‘look grey’ in ordinary vision.²² The corresponding mode of seeing—relative colour—is the only one allowed by the HSD, in which the background is masked.²³

²²Dejan Todorovich notes: ‘The predicted grayness of its appearance is in accord with our phenomenal impression, noted by Kardos (1934), that, if we ascribe a color to the shadow at all, it is usually gray and not chromatic (the case of genuine “colored shadows” excluded, but it is notable that for the layman that is a counterintuitive phenomenon). Kardos notes that is very hard for us to see a shaded portion of a chromatic surface as a darker shade of chromatic color.’ (Todorovich, personal communication.)

²³The psychophysics of contrast colours has a dual relation to normal vision. On the one hand it describes some mechanisms of early vision. On the other hand it measures, fixes, expresses, particular ways of seeing coloured patches: primarily as object colours, but also as shadows. To be good psychophysics it has to remove the ambiguity of colour.

So the phenomena of contrast colours lead us to the ambiguity of colour: that regions of the visual field do not just 'have a colour', but can be seen in different ways. Colours can vary dramatically, depending on how one parses the scene into objects, lighting and transparent media. Jakobsson, Bergström, Gustafsson and Fedorovskaya, 1998, give a striking example of a pattern of saturated colours that looks achromatic when the saturated colour is parsed (and therefore discounted) as the illuminant. And it is quite normal to see two colours in the same place. Here is Evans again:

Perhaps the greatest contribution Katz made in his book [Katz, 1911/1935] was his insistence that the light that is seen to be illuminating objects is perceived as *separate from* the objects. This is a concept that is obvious to the naive and obscure to scientists; it is a basic tenet of the present book that that this is so fundamentally true that no perception of a complex scene can be analyzed without taking it into account. (Evans, 1974, p91)

Both these kinds of ambiguity contradict the notion that perceived colour is just a transduction or transformation of the physical colour. This notion, which Gilchrist (1994b) has called the 'photometer metaphor', and Mausfeld (1998) the 'measurement device conception of perception' has dominated psychophysics at least since Fechner. Although many people have persuasively argued against it, it is so temptingly simple, allows just enough space for some neural complexities (like lateral inhibition and receptive fields), and meshes with so many of our habitual Cartesian assumptions, that it has great staying power. Its dominance is perhaps the deepest reason why contrast colours remain in their uneasy position in colour science: invoked when convenient, but much of the time ignored.

References

- Abramov, I., and Gordon, J. (1994) Color appearance: on seeing red—or yellow, or green, or blue, *Ann Rev Psych* 45, 451-485
- Alpern, M. (1964) Relation between brightness and color contrast, *JOSA* 54 1491-2.
- Brainard, D.H. (1998) Color Constancy in the nearly natural image. 2. Achromatic loci, *JOSA A* 15, 307-325.
- Brainard, D.H., Brunt, W.A., and Speigle, J.M. (1998) Color constancy in the nearly natural image. 1. Asymmetric matches, *JOSA A* 14, 2091-2110.
- Brainard, D.H., and Wandell, B.A. (1986) Analysis of the retinex theory of color vision, *JOSA A* 3, 1651-1661.
- Brill, M.H. (1995) Commentary on Ives, 1912, *Color Res Applctns* 20, 70-71.
- Brookes, A., and Stevens, K.A. (1989) The analogy between stereo depth and brightness, *Perception*, 18 601-614.
- Chichilnisky, E.J., and Wandell, B.A. (1995) Photoreceptor sensitivity changes explain color appearance shifts induced by large uniform backgrounds in dichoptic matching, *Vis Res* 35, 239-254
- Chichilnisky, E.J., and Wandell, B.A. (1996) Seeing gray through the ON- and OFF-pathways, *Visual Neuroscience* 13, 591-596.
- Chichilnisky, E.J., and Wandell, B.A. (1997) Increment-decrement asymmetry in adaptation, *Vis Res* 37, 616.
- Chubb, C., Sperling, G., and Solomon, J.A. (1989) Texture interactions determine perceived contrast, *Proc Natl Acad Sci USA* 86 9631-5.
- Derefeldt, G. (1991) Colour appearance systems, in Gouras P, (Ed) *The Perception of Colour*, Vol 6 of *Vision and Visual Dysfunction*, Cronly-Dillon J, (Ed), MacMillan
- Derrington, A.M., Krauskopf, J., and Lennie, P. (1984) Chromatic mechanisms in lateral geniculate nucleus of macaque, *J Physiol* 357 241-65.
- Evans, R.M. (1948) *An introduction to color*, Wiley
- Evans, R.M. (1974) *The perception of color*, New York: Wiley
- Fairchild, M.D. (1998) *Color Appearance Models*, Reading MA: Addison Wesley.
- Foster, D.H., and Nascimento, S.M.C. (1994) Relational colour constancy from invariant cone-excitation ratios, *Proc Roy Soc B* 257, 115-121.
- Friele, L.F.C. (1961) Analysis of the Brtowen and Brown-MacAdam colour discrimination data. *Die Farbe*, **10**, 193.
- Gilchrist, A.L. (1994a) (Ed), *Lightness, Brightness and Transparency*, Hillsdale, NJ: Erlbaum.
- Gilchrist, A.L. (1994b) Introduction: Absolute versus Relative theories of lightness perception, In Gilchrist A.L., (Ed), *Lightness, Brightness and Transparency*, Hillsdale, NJ: Erlbaum.
- Gilchrist, A.L., Delman, S., and Jacobsen, A. (1983) The classification and integration of edges as critical to the perception of reflectances and illumination, *Perc Psychophys* 33 425-36
- Gouras, P., and Zrenner, E. (1981) Color vision: a review from a neurophysiological perspective, *Progress in Sensory Physiology* 1, 139-179.
- Gerbino, W. (1994) Achromatic transparency, In Gilchrist AL, (Ed), *Lightness, Brightness and Transparency*, Erlbaum.
- He, S., and MacLeod, D.I.A. (1998a) Contrast-modulation flicker: dynamics and spatial resolution of the light adaptation process, *Vis Res* 38, 985-1000.
- Helson, H. (1938) Fundamental Problems in color vision I. The principle governing changes in hue, saturation and lightness of non-selective samples in chromatic illumination, *J Exp Psy* 23 439-476.
- Hering, E. (1890) Beitrag zur Lehre vom Simultankontrast, *Z Psychol* 1, 18-28.
- Hilbert, D.R. (1987) *Color and Color Perception: A Study in Anthropocentric Realism*, Stanford CA: CSLI Lecture Notes.
- Ives, H.E. (1912) The relation between the color of the illuminant and the color of the illuminated object, *Trans Illum Eng Soc* 7, 62-72.
- Jakobsson, T., Bergström, S.S., Gustafsson, K.- A., and Fedorovskaya, E. (1998) Ambiguities in colour constancy and shape from shading, *Perception* 531 - 541.
- Judd, D.B. (1940) Hue, saturation and the lightness of surface colors with chromatic illumination, *JOSA* 30 2-32
- Judd, D.B. (1960) Appraisal of Land's work on Two-Primary Color Projection, *JOSA* 50 254-268
- Kaiser, P.K., and Boynton, R.M. (1996) *Human Color Vision*, 2nd edtn of Boynton, 1979, Opt Soc Am.
- Kardos, L. (1934) Ding und Schatten, *Z Psychologie, Ergänzungsband* 23, 1-184.
- Katz, D. (1935) *The world of colour*, London: Kegan Paul, Trench, Trubner. (First German edition, 1911)

- Kirschmann, A. (1890) Ueber die quantitativen Verhältnisse der simultanen Helligkeits- und Farbencontrastes, *Philosophisches Studien* 6, 417-491
- Land, E.H., and McCann, J.J. (1971) Lightness and retinex theory, *JOSA* 61 1-11
- Larimer, J., and Piantanida, T.P. (1988) The impact of boundaries on colour: stabilized image studies, *SPIE, 901-Image processing, Analysis, Measurement and Quality*, pp241-247.
- Le Grand, Y. (1948) *Light, Colour and Vision*, Eng transltn 1957, London: Chapman & Hall.
- Le Grand, Y. (1949) Les seuils différentiels de couleurs dans la théorie de Young, *Rev d'Optique* 28, 261-278; Eng transltn (1994) in *Color Research and Applications*, 19, pp 296-309
- Lucassen, M.P. (1993) Quantitative Studies of Color Constancy, Dissrttn, TNO
- Mausfeld, R. (1998) Colour perception: from Grassmann codes to a dual code for object and illumination colours, in Backhaus WGK, Kliegl R, Werner JS, (Eds) 1998, *Color Vision*, De Gruyter
- Mausfeld, R., and Niederž, R. (1993) An inquiry into relational concepts of colour based on incremental principles of colour coding for minimal relational stimuli, *Perception* 22, 427-462.
- McDougall, W. (1901) Some new observations in support of Thomas Young's theory of light and colour-vision, III, *Mind*, 10, 347-382
- Merleau-Ponty, M. (1945) *Phenomenology of Perception*, Paris: Tel; English transltn, Smith, C., 1962, London: Routledge & Kegan Paul.
- Metelli, F. (1974) The perception of transparency, *Sci Amer* 230 90-8.
- Mollong, J.D. (1987) On the nature of models of colour vision. *Die Farbe*, 34, 29-46.
- Mollon, J.D. (1995) Seeing Colour, In Lamb T, Bourriau J, (Eds) *Colour: Art & Science*, Cambridge: University Press
- Monge, G. (1789) Mmoire sur quelques phénomènes de la vision, *Annales de Chimie* 3, 131-147
- Nerger, J.L., Piantanida, T.P., and Larimer, J. (1993) Color appearance of filled-in backgrounds affects hue cancellation, but not detection thresholds, *Vis Res* 33, 165-172
- Ovenston, C.A., and Whittle, P. (1996) Improved discrimination near the background field: transfer of the 'crispness effect' to colour, *Perception* 25 supp p16
- Poirson, A.B., and Wandell, B.A. (1993) The appearance of colored patterns: pattern-color separability, *JOSA* 12, 2458-2471
- Pugh, E.N. Jr, and Mollon, J.D. (1979) A theory of the 1 and 3 color mechanisms of Stiles, *Vision Res* 19 293-312
- Ragona-Schina ()
- Rivers, W.H.R. (1900) Vision, In Schřfer EA (Ed) *Textbook of Physiology Vol 2*, Edinburgh: Young J. Pentland, 1026-1148
- Rollet, A. (1867) Zur Physiologie der Contrastfarben, *Akad Wiss, Math-naturwiss Klasse; Sitzungsb*, 55 Band, Abt 2, 741-66.
- Semmelroth, C.C. (1970) Prediction of lightness and brightness on different backgrounds, *JOSA* 60 1685-1689.
- Shevell, S.K. (1978) The dual role of chromatic backgrounds in color perception, *Vis Res* 18 1649-61.
- Shinomori, K., Scheffrin, and Werner, J.S. (1997) Spectral mechanisms of spatially induced blackness: data and quantitative model, *JOSA A* 14, 372-387.
- Spencer, D.E. (1943) Adaptation in color space, *JOSA* 33 10-17.
- Takasaki, H. (1966) Lightness change of grays induced by change in reflectance of gray background, *JOSA* 56 504-509
- Takasaki, H. (1967) Chromatic changes induced by changes in chromaticity of background of constant lightness, *JOSA* 57 93-6
- Tschermak, A. (1903) Über Kontrast und Irradiation, *Ergebn. Physiol* 2, 726-798.
- Vingrys, A.J., and Mahon, L.E. (1998) Color and luminance detection and discrimination asymmetries and interactions, *Vision Research* 38, 1085-1096.
- von Bezold, W. (1874) *Die farbenlehrer im Hinblick auf Kunst und Kunstgewerbe*, Brunswick: Westermann. English translation by S.R. Koehler (1876), *The Theory of Color and Its Relation to Art and Art-Industry*, Boston:Prang.
- von Kries, J. (1905) Influence of adaptation on the effects produced by luminous stimuli, Eng trans in MacAdam, DL (Ed) 1970 'Sources of Color Science', MIT
- Walls, G.L. (1960) Land! Land!, *Psych Bull* 57 29-48
- Walraven, J. (1976) Discounting the background - the missing link in the explanation of chromatic induction, *Vis Res* 16 289-295

- Walraven, J. (1981) Perceived colour under conditions of chromatic adaptation: evidence for gain control by ¹ mechanisms, *Vis Res* 21 611-20
- Walraven, J., Enroth-Cugell, C., Hood, D.C., MacLeod, D.I.A., and Schnapf, J. (1990) The control of visual sensitivity: receptor and post-receptor processes, In Spillmann L, Werner JS *Visual Perception: Neurophysiological Foundations*, San Diego, CA: Academic Press
- Wandell, B.A. (1995) *Foundations of Vision*, Sunderland MA: Sinauer.
- Webster, M.A. (1996) Human colour perception and its adaptation, *Network: Computation in Neural Systems* 7, 587-634.
- Webster, M.A., and Mollon, J.D. (1995) Colour constancy influenced by contrast adaptation, *Nature* 373, 694-698.
- Werner, J.S., and Walraven, J. (1982) Effect of chromatic adaptation on the achromatic locus: the role of contrast, luminance and background color, *Vis Res* 22 929-43
- Whittle, P. (1992) Brightness, discriminability and the "Crispening Effect", *Vis Res* 32 1493-1507
- Whittle, P. (1994a) The psychophysics of contrast-brightness, In Gilchrist AL, (Ed), *Lightness, Brightness and Transparency*, Erlbaum
- Whittle, P. (1994b) Contrast-brightness and ordinary seeing, In Gilchrist AL, (Ed), *Lightness, Brightness and Transparency*, Erlbaum.
- Whittle, P., and Arend, L.E. (1991) Homochromatic colour induction, *Perception* 20, 99 (A37), abstract of talk at 14th ECVF, Vilnius, Lithuania.
- Whittle, P., and Challands, P.D.C. (1969) The effects of background luminance on the brightness of flashes, *Vis Res* 9 1095-1110
- Wilson, M.H., and Brocklebank, R.W. (1960) Two-colour projection phenomena, *J Photograph Sci* 8 142-50.
- Woodworth, R.S. (1938) *Experimental Psychology*, New York: Henry Holt
- Wuerger, S. (1996) Color appearance changes resulting from iso-luminant chromatic adaptation, *Vis Res* 36, 3107-18
- Wyszecki, G., and Stiles, W.S. (1967) *Color Science*, New York:Wiley. 2nd edtn, 1982.

Figure captions

Figure 1 (a) A version of the haploscopic superimposed display ('HSD') with eight patches in each eye. (b) Two sets of eight colours in a constant luminance diagram (the logarithmic version of the MacLeod-Boynton diagram). The clown's hat shape is the gamut of colours available on the computer monitor.

Figure 2 A rigid translation of a vector in the LogL, LogM plane preserves cone contrasts.

Figure 3 Showing how the differences of contrasts, as opponent colour expressions, vary with luminance contrast (between ± 1 Michelson contrast), for saturated purple (squares) and green (circles) of constant chromaticity. (a) For signed Michelson contrasts; (b) for Weber contrasts. Filled symbols show differences between S and the average of L and M cone contrasts, hollow ones between L and M cone contrasts.

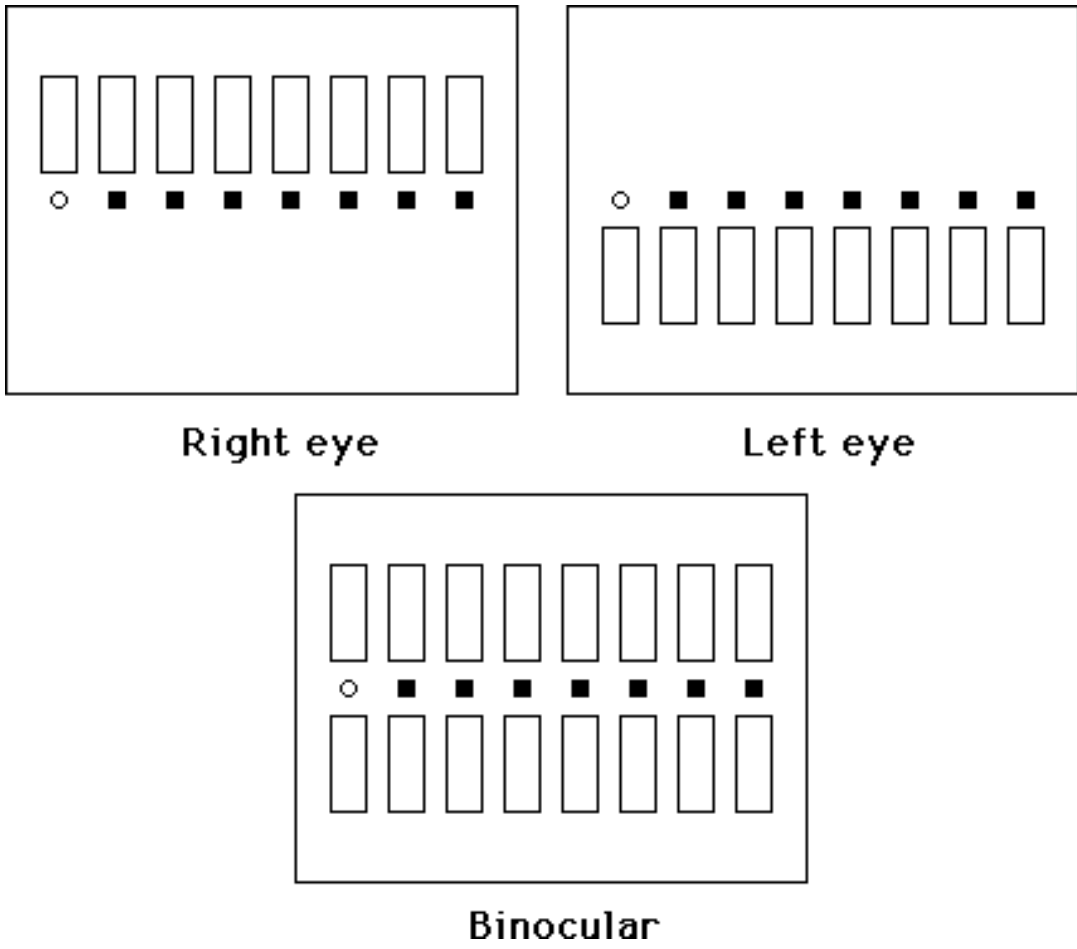


Figure 1a




Figure 1b

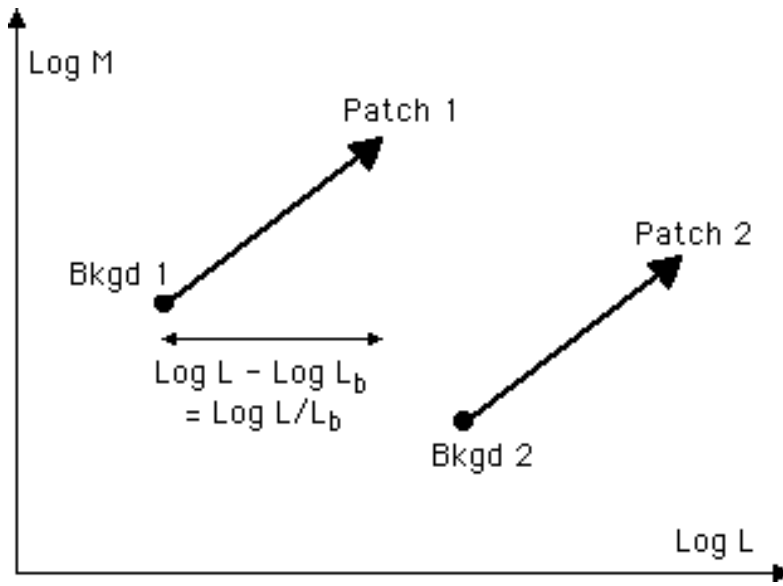


Figure 2

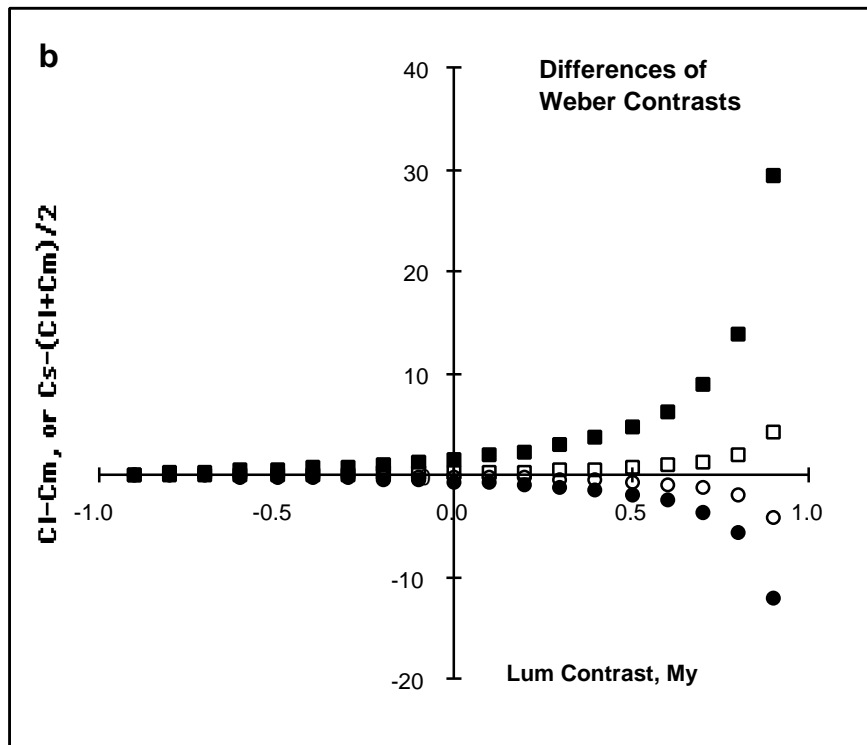
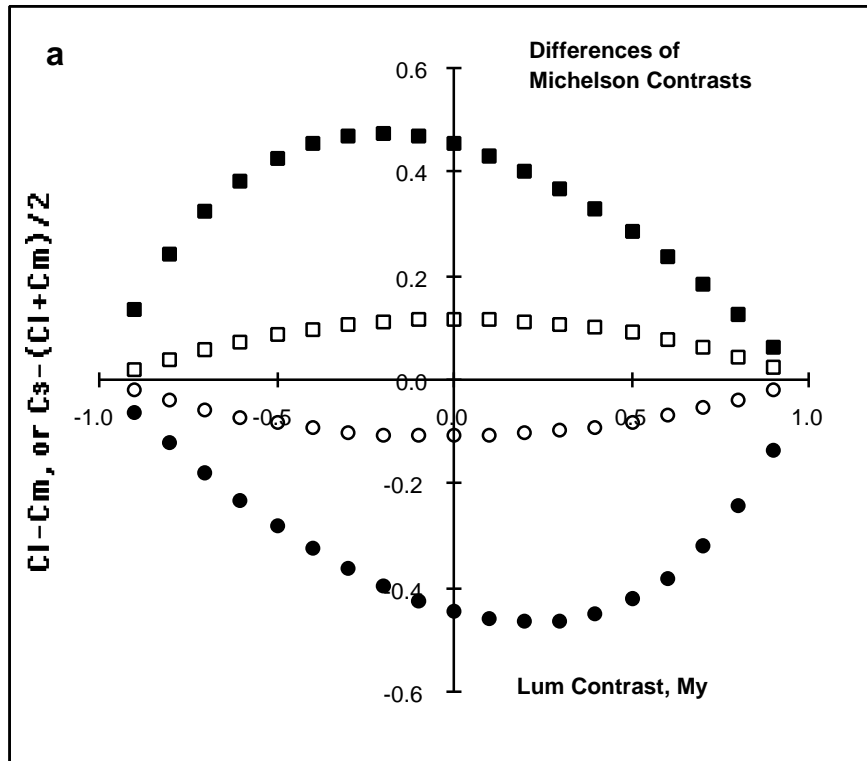


Figure 3