ON THE NATURE OF SIMULTANEOUS COLOUR CONTRAST

Dissertation
zur Erlangung des Doktorgrades
der Philosophischen Fakultät
der Christian-Albrechts-Universität
zu Kiel

vorgelegt von

Vebjørn Ekroll

Kiel
2005

Erstgutachter: . . . . Prof. Dr. Rainer Mausfeld

Zweitgutachter: . . . . PD Dr. Johannes Andres

Tag der Mündlichen Prüfung: 27.06.2005

Durch den zweiten Prodekan, Prof. Dr. Norbert Nübler
zum Druck genehmigt am: 13.07.2005

Our journey from the psychophysics of simple displays toward the perception of natural scenes may call to mind the following well-known story that appears in many guises. Benighted space traveller in impenetrable rain forest on unknown planet, on meeting small person with long ears: ”Sir, can you tell me the way to Alderon?” Person with ears: ”If to Alderon going were I, not from here start would I.”

- PAUL WHITTLE

Contents

Chapter 1

Introduction

The two central squares in Colour Plate I (see page 137) are physically equal, yet they appear rather different since they are embedded in surrounds of different colour. This phenomenon, which is referred to as *simultaneous colour contrast*, has been known since antiquity, and particularly in the two last centuries a rather impressive amount of scientific effort has been directed towards developing a deeper understanding of it (Tschermak, 1903; Whittle, 2003). The motivations behind these investigations are and have been manifold, ranging from immediate practical concerns to pure intellectual curiosity and philosophical questions. In the sciences of vision and perception the major impetus stems from the notion that the study of perceptual illusions is likely to teach us something about the mechanisms which transform the highly ambiguous raw input impinging on our sensory organs into useful and reliable information about the external world (Mach, 1866; Hoffman, 1998; Eagleman, 2001). In accordance with this general approach, it has been proposed that the simultaneous contrast 'illusion' is intimately linked to biologically important mechanisms of colour constancy. According to Hering's (1920) classical theory, for instance, both colour constancy and simultaneous colour contrast are the results of the same neural mechanism of lateral inhibition. Similarly, Helmholtz (1911) proposed that the simultaneous colour contrast 'illusion' can be observed whenever the physical situation is at odds with assumptions that the visual system relies on in order to obtain colour constancy, i.e. to provide the organism with an illumination-independent representation of object colours. Under either hypothesis it is assumed that simultaneous colour contrast and colour constancy are basically two sides of the same coin. Hence, the study of simultaneous contrast can be expected to yield insights into the nature of colour constancy, and vice versa. Indeed, experiments designed to investigate colour constancy are often virtually indistinguishable from experiments performed in order to learn more about simultaneous colour contrast.

By virtue of the postulated connection to colour constancy, understanding simultaneous contrast is regarded an important goal of vision science, and accordingly there is an impressive amount of experimental data available. Yet, although significant insights have been gained, simultaneous colour contrast is still far from a well-understood phenomenon; A reasonably general and commonly accepted quantitative model accounting for the psychophysical data has yet to emerge. The reasons for these difficulties are probably manifold. It is well known that the simultaneous contrast effect depends on a number of parameters, such as viewing conditions, stimulus size, perceptual organisation, and how the observers are instructed. In the present investigations, though, the main focus is on a more basic conceptual problem.

As will be discussed in chapter 3, most quantitative models of simultaneous contrast implicitly rest on the seemingly natural and innocuous presumption that the colour changes induced by a coloured surround are, in principle, of the same nature as those obtained by changing the physical color-coordinates of the target patch. Accordingly, it is assumed that the latter kind of colour change may *compensate* the former, or vice versa. Though this ‘compensation assumption’ is seldom explicitly discussed, it is of profound theoretical significance. The hard facts of Young-Helmholtz trichromatic theory demonstrate that human colour vision is trivariant, both in a well-defined psychophysical sense (Krantz, 1975b) and a corresponding physiological one (Sharpe & Stockman, 1999). If the perceived colour of a local stimulus were independent of the context in which it is embedded, this would imply that the *perceived* colour corresponding to a location of the visual field can vary along no more than three dimensions. The context-dependence of perceived colour demonstrated by such phenomena as simultaneous contrast, however, opens up the theoretical possibility that perceived colour may vary along more than three dimensions, simply because context-dependence means that perceived colour depends on more triplets of cone-excitation values than that of the target itself (Evans, 1974). Thus, that perceived colour should be trivariant does not follow from trichromatic theory alone. Rather, the trivariance of perceived colour would only follow from trichromatic theory *provided that* the compensation assumption is valid (W. S. Stiles, 1961).

That most quantitative models of simultaneous contrast are based on the compensation assumption is evidenced by the fact that they represent the perceived colour of the target by a triplet of numbers. That this may be inappropriate, though, is suggested by the work of several authors: Based on a number of different arguments, it has been conjectured that more than three numbers are necessary in order to represent the perceived colour of the target adequately (e.g. Katz, 1911; Gelb, 1929; Kanizsa, 1966; Evans, 1964, 1974; Mausfeld, 1998; Niederée, 1998).

In chapter 4, I present experimental evidence lending strong support to these claims. The basic finding is that experimental estimates of the neutral point obtained with the conventional method of grey settings differ, in a clear and lawful manner, from estimates of the neutral point obtained with some novel methods, in which the neutral point is measured according to *indirect* criteria. It is found, for instance, that when targets are viewed embedded in a uniform coloured surround, the convergence point for lines of constant hue does not, as one may intuitively expect, coincide with the chromaticity of a patch which is judged to appear achromatic, but instead with the chromaticity of the coloured surround. Though this finding may appear paradoxical, it is shown that it can be accounted for in quite rational terms by giving up the compensation assumption and instead assuming that simple centre-surround stimuli sometimes evoke phenomena reminiscent of perceptual transparency, which is traditionally thought to occur only in more complex stimulus displays (Metelli, 1970; D’Zmura, Colantoni, Knoblauch, & Laget, 1997; Adelson, 1993, 2000; Anderson, 1997).

Since the compensation assumption turns out to be unwarranted, it is clear that framing models of simultaneous contrast in terms of it is bound to impede the development of better and more adequate models. Thus, from this perspective, it appears less surprising that the vast amount of previous research efforts, which have implicitly presupposed the validity of the compensation assumption, have been but moderately successful in producing a commonly accepted model. On the constructive side, one may hope that once the traditional theoretical framework based on the compensation assumption is traded for a more adequate one, making progress will be easier. As mentioned above, the findings reported in chapter 4, which otherwise must appear rather paradoxical, can be easily understood within a framework based on perceptual transparency. A reappraisal of previous models and findings in terms of this framework has several further advantages. To begin with, it offers a natural explanation for the inconsistent and conflicting results reported in course of the long-standing Walraven-Shevell controversy (Walraven, 1976, 1979; Shevell, 1978, 1980; Drum, 1981; Adelson, 1981; Davies, Faivre, & Werner, 1983). Furthermore, as will be shown, one of the candidate models discussed in this controversy, viz. the contrast coding model (or, ‘full discounting model’) put forward in the seminal contributions of Whittle and Walraven (Whittle & Challands, 1969; Whittle, 1994b, 2003; Walraven, 1976), is in excellent agreement with the relevant findings provided that it is reconceptualised in terms of the new framework. Understood in the classical 'literal' reading implied by the compensation assumption, the contrast-coding model is problematic in many respects: Some of the empirical evidence suggests that it is only valid to a certain approximation (e.g. Shevell, 1978), and it is quite quite clear that it makes a paradoxical prediction in the case of zero-contrast centre-surround stimuli (Whittle, 1994b, p. 451, Mausfeld & Niederee 1993, p. 451). When the contrast-coding model is reconceptualised in terms of perceptual transparency, though, these arguments against its validity are clearly inappropriate: Perceptual transparency means that the total colour impression consists of two phenomenally distinct layers or components, and the contrast-coding model is found to describe the colour of one of these very accurately.

The inappropriateness of the compensation assumption has a further important consequence. The compensation assumption provides the rationale for classical psychophysical methods for measuring simultaneous colour contrast, such as grey settings and asymmetric colour matching. Thus, the invalidity of the compensation assumption implies that these methods should yield results flawed with artifact. This matter is pursued in the asymmetric colour matching experiments reported in chapter 5. Here, a rather complex pattern of results grossly consistent with some previous reports (Smith & Pokorny, 1996; Miyahara, Smith, & Pokorny, 2001) is obtained. The strong nonlinearities in these data are not accounted for by classical models of simultaneous contrast. However, the present experiments reveal that it is not only impossible to account for the complexity of the data curves in the traditional manner implied by the compensation assumption, but also that the complexities can be accounted for in a quite simple manner. Basically, the characteristic and sometimes very large step-shaped portions of the data-curves are found to be due to the fact that well-defined sets of colour impressions cannot be realised at all when the target is embedded in a uniform surround. Which set of colour impressions are realisable depends on the colour of the surround in a simple and lawful manner, and since the sets of missing colours are different for differently coloured surrounds, establishing a colour match between two targets embedded in different surrounds is sometimes impossible. Whenever a match is impossible, the subjects must by necessity resort to choosing the best match among poor alternatives, and this accounts for the complexity of the data curves. Based on these results, a simple descriptive model based on the notions of 'saturation scale truncation and extension' is put forward. This very simple model, which systematises the way in which a surround of a given colour prohibits the perception of a set of target colour impressions, is found to account very well for the strong nonlinearities of the data curves. Interestingly, once the portions of the data curves which can be accounted for as simply due to the impossibility of establishing a true match are disregarded, the remaining data points, which represent a rather small induction effect, can be accounted for by simple von Kries scaling (Kries, 1905). Hence, it would appear that the total effect observed is due to two distinct mechanisms, one which leads to saturations scale truncation and extension, and one which leads to von Kries scaling. This hypothesis is substantiated in further experiments, in which it is found that when comparable variegated surrounds are used instead of uniform ones, the von Kries effect remains, whereas the other effect is abolished.

The latter result suggests that there are fundamental qualitative differences between uniform and variegated surrounds. Thus, the much-cited idea, according to which any variegated surround can be replaced with a uniform one which is functionally equivalent in the sense that it has the same effect on the colour appearance of targets embedded in it (Valberg & Lange-Malecki, 1990; Brill, 2000) is clearly false, as suggested also in a number of previous studies (e.g. Schirillo & Shevell, 1996; Brown & MacLeod, 1997; Shevell & Wei, 1998; Barnes, Wei, & Shevell, 1999). Interestingly, though, the present findings suggest that uniform and variegated surrounds with the same spatial average of cone excitation vectors can nevertheless be functionally equivalent with respect to a common von Kries mechanism: The reason why functional equivalence is not observed in psychophysical experiments appears to be that a second mechanism, leading to saturation scale extension and truncation, is only operating when uniform surrounds are used.

The major conclusions that are suggested by the present findings can be summarised as follows:

- There are fundamental differences between uniform and variegated surrounds.
- In the case of uniform surrounds, the compensation assumption is unwarranted. Instead, perceptual transparency appears to play a central role.
- Though uniform surrounds are geometrically simpler than variegated ones, they appear to evoke more complex visual processing than variegated ones.

Furthermore, the rather simple descriptive model based on saturation scale extension and truncation is shown to describe different findings and observations related to simultaneous contrast, including Brown and MacLeod's (1997) gamut expansion effect and Meyer's (1855) effect. As argued in the discussion, the present findings also suggest a new perspective on the role of contrast-coding in colour vision. Contrary to the widespread view according to which contrast-coding constitutes a fundamental peripheral mechanism providing the input for all later processing of colour in the visual pathway, it is suggested that absolute coding plays at least an equally important role. Thus, contrast-coding may perhaps better be conceived of as a special-purpose mechanism upon which the visual system relies only under certain circumstances.

The next chapter of this thesis gives a synopsis of basic colour theory. Then, in chapter 3, classical models of simultaneous contrast are reviewed. In chapters 4 and 5 the experiments are described, and in chapter 6 the findings are discussed in relation to a number of current research issues.

Part I

Theoretical background

Chapter 2

Basic colour theory

The basis for our current scientific understanding of colour perception is provided by Young-Helmholtz trichromatic theory and Hering's opponent colours theory. The former describes the intitial coding of colour by the receptors of the retina, whereas the latter is thought to describe the processing and recoding of the initial colour signal at a higher level of the visual system. Since knowledge of these theories is indispensible for understanding the models of simultaneous contrast to be discussed in the next chapter, an introductory synopsis is provided here. Readers familiar with basic colour theory may skip this chapter, those interested in a more thorough treatment are referred to the standard reference work of Wyszecki and Stiles (1982), the very readable expositions of Schrödinger (1920a, 1920b), or the penetrating but somewhat more demanding formal treatments provided by Krantz (1975b, 1975a). The historical origins of basic color theory are discussed by Mollon (2003).

2.1 Trichromatic theory

It is a well-established fact that the colour vision of normal human observers is trichromatic (Maxwell, 1860; Helmholtz, 1911). Although the immediate stimuli for vision – light beams with radiant energy distributed in the wavelength interval between approximately 400 and 700 nm – may vary along infinitely many dimensions (one for each real value of wavelength λ), human colour vision depends, in a certain sense, on only three variables. The generally acknowledged explanation for this is that the retina contains only three kinds of receptors which are active under daylight illumination conditions, the long (L), middle (M), and short (S) wavelength sensitive cones. The spectral sensitivity curves of these receptors – $l(\lambda), m(\lambda)$ and $s(\lambda)$, respectively – were first estimated in the classical psychophysical experiments of König and Dieterici (1886), and modern estimates, such as those of Stockman, MacLeod, and Johnson (1993), shown in Figure 2.1, do not differ substantially from the original ones.

The cone excitation values L, M, S of a given light stimulus is obtained by multiplying its spectral energy distribution $x(\lambda)$ with each of the spectral sensitivity curves and then integrating over the visual spectrum:

$$\begin{aligned} L(x) &= \int l(\lambda)x(\lambda)\,d\lambda \\ M(x) &= \int m(\lambda)x(\lambda)\,d\lambda \\ S(x) &= \int s(\lambda)x(\lambda)\,d\lambda. \end{aligned} \tag{2.1}$$

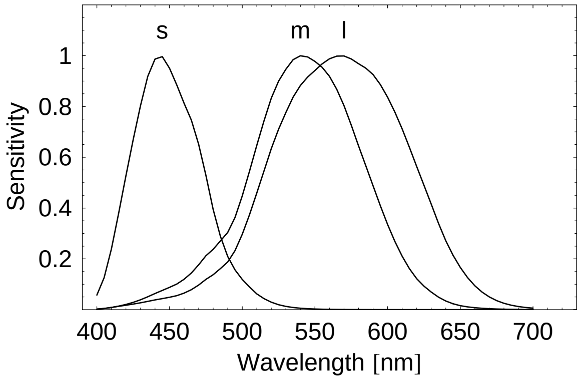

Figure 2.1: The spectral sensitivities of the cones according to the estimates of Stockman, MacLeod and Johnson (1993).

Since colour perception is assumed to depend only on the excitation of these three kinds of receptors, any two light stimuli $x(\lambda)$ and $y(\lambda)$ should appear identical in colour whenever they evoke identical triplets $\mathbf{e}(x), \mathbf{e}(y)$ of cone-excitations, that is

$$x(\lambda) \sim y(\lambda) \Leftrightarrow \mathbf{e}(x) = \mathbf{e}(y) \tag{2.2}$$

where the symbol $\sim$ denotes perceptual equivalence, and

$$\mathbf{e}(x) := \begin{pmatrix} L(x) \\ M(x) \\ S(x) \end{pmatrix}. \tag{2.3}$$

The Grassmann Laws The basic psychophysical evidence for trichromatic theory stems from simple colour matching experiments in which the observer typically views a bipartite stimulus consisting of two semi-circular fields (see Figure 2.2). In this classical experimental situation, which is commonly referred to as the **two-degree-paradigm**, or **symmetric colour matching**, several fundamental observations, commonly referred to as the Grassmann Laws of Colour Mixture, can be made. According to **the first Grassmann Law**, two light beams which appear equal in spite of having different spectra will continue to look equal if the same third light stimulus is added to both of them, i.e.

$$x \sim y \Rightarrow x + z \sim y + z \tag{2.4}$$

for all light beams x, y and z. The light beams are here identified with their spectral power functions, which provide a measure of the radiant power density for each value of wavelength λ. For our purposes, the proper domain of these functions is the visual spectrum, which extends from about $\lambda = 400$ to $\lambda = 700$ nm. In Figure 2.3, two examples of such functions x and y, as well as their sum $x + y$, are shown: The sum of two functions is simply obtained by wavelength-wise addition, i.e. $f = x + y$ is defined by $f(\lambda) = x(\lambda) + y(\lambda)$ for all values of λ. Physically, light beams can be added in this way simply by superimposing them, say, by projecting them on the same screen.

According to **the second Grassmann Law**, two light beams which appear equal will continue to look equal if their intensities are changed by the same amount, i.e.

$$x \sim y \Rightarrow \alpha \cdot x \sim \alpha \cdot y \tag{2.5}$$

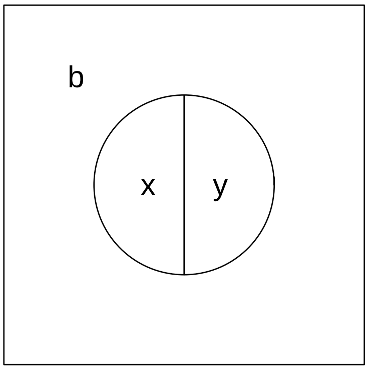

Figure 2.2: Symmetric colour matching: The observer views two semi-circular fields against a common background. Typically, one of the stimuli x or y is fixed and the other is a mixture of three component lights, the amounts of which the observer may adjust. The task of the observer is to find a setting such that the semi-fields appear equal in colour. Whenever this is the case, we write $x \sim y$. The semi-disks are viewed foveally, and together subtend a visual angle of 2 degrees or less, such that only the fovea, which contains only cones and no rods, is stimulated. This is the basic experimental situation to which the laws of trichromatic theory refer. For the laws to be valid, it is of fundamental importance that the two semi-fields are viewed against a common background (this is why one speaks of *symmetric* colour matches). Typically, the background consists of complete darkness, but this is merely a matter of convention; The important point is that it is constant. Note that the observer is never asked *what* colour the semi fields have, only whether they have the *same* colour, or, equivalently, whether the dividing line between the two semi-circles becomes imperceptible.

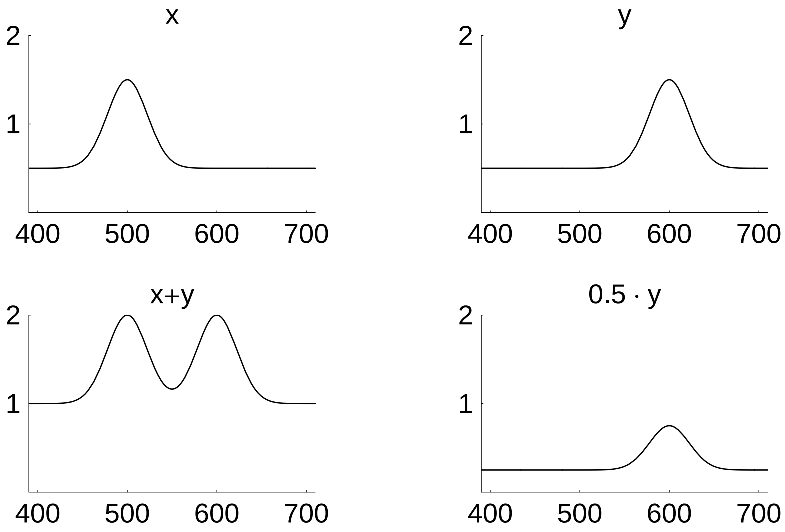

Figure 2.3: Addition and scalar multiplication of (spectral) functions. Top panels: Two arbitrarily chosen functions x and y. Bottom left: The function that is obtained by adding x and y. Bottom right: The function that is obtained by multiplying the function y with the scalar 0.5.

for all light beams x, y and scalar factors α. Mathematically, the multiplication of a function x with a scalar α amounts to wavelength-wise multiplication of the function values $x(\lambda)$ with the scalar α, i.e. $f = \alpha \cdot x$ is defined through $f(\lambda) = \alpha \cdot x(\lambda)$ for all values of λ. For illustration, the result of multiplying a spectral power function y with the scalar factor 0.5 is shown in Figure 2.3. One way in which scalar multiplication of a light beam can be physically realised is by inserting a neutral density filter (e.g. 'untinted sunglasses') in the light beam's path.

The validity of the above psychophysical laws is what one would expect based on the physiological theory behind equation set 2.1 and the linking proposition (Teller, 1984) expressed in equation 2.2. It is easily deduced from equation set 2.1 that the cone excitation vector of a light stimulus should be a linear function of its spectral density function, i.e.

$$\mathbf{e}(x + y) = \mathbf{e}(x) + \mathbf{e}(y) \tag{2.6}$$

and

$$\mathbf{e}(\alpha \cdot x) = \alpha \cdot \mathbf{e}(x) \tag{2.7}$$

holds for all spectra x, y and scalar factors α.[1]

The third Grassmann Law is a bit trickier to formulate in an immediately comprehensible yet correct manner. Basically, though, it states that any light stimulus can be matched in appearance by mixing just three reference lights in appropriate proportions. Strictly speaking this is sometimes not possible, but this is merely due to the extensive overlap of the spectral sensitivity curves.

For trichromatic observers, it is always possible to find three reference lights a, b and c so that neither can be matched by mixing the two others. This should be the case whenever the the corresponding cone excitation vectors $\mathbf{e}(a), \mathbf{e}(b)$ and $\mathbf{e}(c)$ are linearly independent. Once such a set of reference lights has been chosen, it is possible to find mixture coefficients α, β and γ for many test stimuli x such that

$$x \sim \alpha \cdot a + \beta \cdot b + \gamma \cdot c. \tag{2.8}$$

In the case of test stimuli for which this is not possible, it will be possible to establish a match by adding one or two of the reference lights to the test stimulus instead of to the remaining reference light(s), e.g.

$$x + \alpha \cdot a \sim \beta \cdot b + \gamma \cdot c. \tag{2.9}$$

Using the linking proposition and linearity this colour equation simply means that

$$\mathbf{e}(x) + \alpha \cdot \mathbf{e}(a) = \beta \cdot \mathbf{e}(b) + \gamma \cdot \mathbf{e}(c) \tag{2.10}$$

which, by solving for $\mathbf{e}(x)$, can be rewritten as

$$\mathbf{e}(x) = -\alpha \cdot \mathbf{e}(a) + \beta \cdot \mathbf{e}(b) + \gamma \cdot \mathbf{e}(c). \tag{2.11}$$

This motivates the formal convention introduced by Maxwell (1860) of rewriting equation 2.9 as

$$x \sim -\alpha \cdot a + \beta \cdot b + \gamma \cdot c. \tag{2.12}$$

Though a negative mixture coefficient makes no physical sense, this way of writing equation 2.9 highlights the fact that it can be interpreted as implying the relation between cone excitation vectors expressed in equation 2.11.

---

[1] By the linking proposition $x \sim y$ means $\mathbf{e}(x) = \mathbf{e}(y)$, hence $\mathbf{e}(x) + \mathbf{e}(z) = \mathbf{e}(y) + \mathbf{e}(z)$. By linearity this implies $\mathbf{e}(x + z) = \mathbf{e}(y + z)$, which by the linking proposition means that $x + z \sim y + z$. The second Grassmann Law is obtained by analogous reasoning.

This convention being made, where the reference lights are always written on the right side of the colour equation, and a negative sign of the mixture coefficients denotes that the corresponding reference light was actually added to the test stimulus, we may formulate the third Grassmann Law as follows: For any test stimulus x, real and unique mixture coefficients α, β and γ can always be found such that equation 2.8 holds.

The third Grassmann Law, then, also referred to as the **psychophysical trivariance of color matching**, implies that as far as normal human color vision is concerned, it is possible to specify any light beam uniquely through just tree numbers α, β and γ. Whenever equation 2.8 holds, the latter are said to be the tristimulus coordinates of the stimulus x with respect to the primaries a, b and c.

A very important corollary[2] of the Grassman Laws is that the tristimulus vector $\mathbf{t} := (\alpha, \beta, \gamma)^t$ of a light stimulus is a linear function of its spectral density function, i.e.

$$\mathbf{t}(x+y) = \mathbf{t}(x) + \mathbf{t}(y) \tag{2.13}$$

and

$$\mathbf{t}(\alpha \cdot x) = \alpha \cdot \mathbf{t}(x) \tag{2.14}$$

for all light stimuli x, y and scalar factors α. This linearity is of fundamental importance. To begin with, since any light beam can be regarded as a sum of monochromatic light beams of unit radiant power, it ensures that it is possible to compute the tristimulus values of any arbitrary light beam once the tristimulus values have been determined for all monochromatic stimuli of unit radiant power in the visual spectrum.

In order to appreciate this point, first consider that, in actual practice, measuring a spectral density function boils down to partitioning the visual spectrum into N disjunct wavelength intervals of width $\Delta\lambda$ and then measuring the radiant power within each of them. This yields a *discrete* spectral density function, such as the one in Figure 2.4. In this particular example, it is assumed that the visual spectrum from 300 to 700 nm is partitioned into N=31 wavelength intervals of width $\Delta\lambda$ = 10 nm. Clearly, any spectral density function x measured at this spectral resolution can be represented by a vector

$$\mathbf{x} := \begin{pmatrix} x_1 \\ x_2 \\ x_3 \\ \cdot \\ \cdot \\ x_N \end{pmatrix}$$

with N entries, where the entry x_i stands for the radiant power in the i-*th* wavelength interval. A so-called **monochromatic light** is a light having all of its radiant power confined to a single

---

[2] In order to prove this corollary, it is necessary to assume that colour matches are symmetric and transitive, i.e. $a \sim b$ implies $b \sim a$ (symmetry) and $a \sim b$ together with $b \sim c$ imply $a \sim c$ (transitivity) for all lights a, b, c ('Zero'th Grassmann Law'). Using this, it follows from the first Grassmann Law that $a \sim b$ together with $c \sim d$ implies $a + c \sim b + d$. To see this, consider that by $a \sim b$ and $c \sim d$ we have $a + c \sim b + c$ and $c + b \sim d + b$, respectively, which yields $a + c \sim b + c \sim d + b \sim b + d$, in short $a + c \sim b + d$. We are now ready to prove that Eq. 2.13 follows from the first Grassmann Law: Let $(\alpha_x, \beta_x, \gamma_x)^t$ and $(\alpha_y, \beta_y, \gamma_y)^t$ be the tristimulus vectors of x and y, respectively. By definition, this means that the colour equations $x \sim \alpha_x a + \beta_x b + \gamma_x c$ and $y \sim \alpha_y a + \beta_y b + \gamma_y c$ hold. By the above corollary of the first Grassmann Law, this means $x + y \sim \alpha_x a + \beta_x b + \gamma_x c + \alpha_y a + \beta_y b + \gamma_y c$. Rearranging terms yields $x + y \sim (\alpha_x + \alpha_y)a + (\beta_x + \beta_y)b + (\gamma_x + \gamma_y)c$, that is, $(\alpha_x + \alpha_y, \beta_x + \beta_y, \gamma_x + \gamma_y)^t$ is the tristimulus vector of $x + y$. This proves that Eq. 2.13 follows from the first Grassmann Law. Proving that Eq. 2.14 follows from the second Grassmann Law is even simpler.

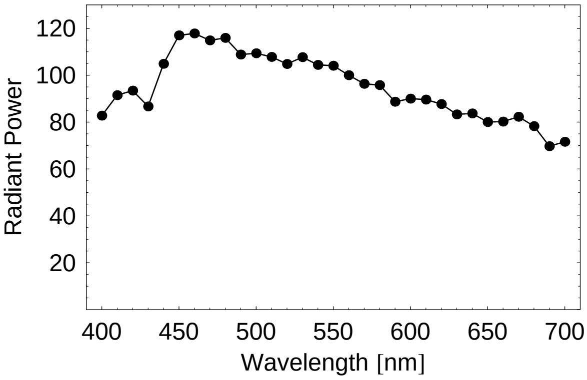

Figure 2.4: Discrete representation of a spectral density function.

wavelenth interval.[3] Hence, in the above representation, a *monochromatic* light corresponds to a vector with only one non-zero entry. For instance, a monochromatic light of wavelength λ_2, where λ_2 is the mid-point of the $2nd$ wavelength interval, would be given by a vector of the form

$$\mathbf{m}_2 := \begin{pmatrix} 0 \\ x_2 \\ 0 \\ \cdot \\ \cdot \\ 0 \end{pmatrix} = x_2 \cdot \begin{pmatrix} 0 \\ 1 \\ 0 \\ \cdot \\ \cdot \\ 0 \end{pmatrix},$$

where x_2 is the radiant power of the monochromatic light. Thus, from

$$\mathbf{x} := \begin{pmatrix} x_1 \\ x_2 \\ x_3 \\ \cdot \\ \cdot \\ x_N \end{pmatrix} = x_1 \cdot \begin{pmatrix} 1 \\ 0 \\ 0 \\ \cdot \\ \cdot \\ 0 \end{pmatrix} + x_2 \cdot \begin{pmatrix} 0 \\ 1 \\ 0 \\ \cdot \\ \cdot \\ 0 \end{pmatrix} + x_3 \cdot \begin{pmatrix} 0 \\ 0 \\ 1 \\ \cdot \\ \cdot \\ 0 \end{pmatrix} +, \cdots, +x_N \cdot \begin{pmatrix} 0 \\ 0 \\ 0 \\ \cdot \\ \cdot \\ 1 \end{pmatrix},$$

it is obvious that any spectral density function can be regarded as a linear sum of monochromatic lights of unit radiant power. In more compact form, this can be written as

$$\mathbf{x} = \sum_{i=1}^{N} x_i \cdot \mathbf{u}_i, \tag{2.15}$$

where $\mathbf{u}_i$ is the spectral density function of a monochromatic light having unit radiant power in the i-*th* wavelength interval, and zero radiant power in all of the others. Thus, we have

$$\mathbf{t}(\mathbf{x}) = \mathbf{t}\left(\sum_{i=1}^{N} x_i \cdot \mathbf{u}_i\right) = \sum_{i=1}^{N} x_i \mathbf{t}(\mathbf{u}_i), \tag{2.16}$$

---

[3] Partitioning the spectrum into smaller wavelength intervals yields a finer representation of the spectral distribution, and the *continuous* spectral distribution, which is a purely theoretical entity, can be thought of as the function which is obtained when $\Delta\lambda$ approaches zero. In actual practice, though, one obviously has to work with wavelength intervals of a finite width, and for the purposes of colour science working with wavelength intervals of 1, 5 or 10 nm is often considered to yield sufficient accuracy.

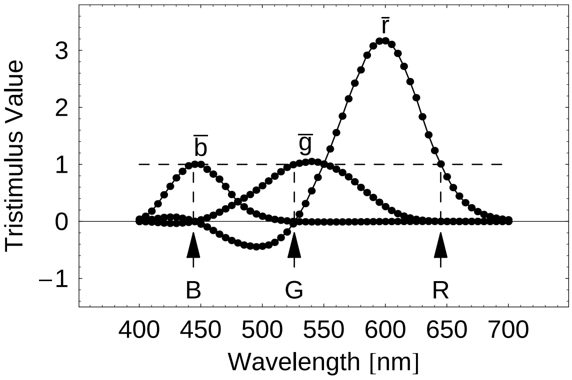

Figure 2.5: A set of colour matching functions provided by Stiles and Burch (1955). The numerical values of these, as well as many other important colourimetric functions are freely available from from Stockman and Sharpe's Colour & Vision Database (http://cvision.ucsd.edu/).

where the second equality is due to the Grassmann linearity laws (equations 2.13 and 2.14). We now see that the tristimulus vector $\mathbf{t}(\mathbf{x})$ of *any arbitrary* light is obtained by the straightforward computation defined by equation 2.16. The only pieces of information necessary for performing this computation are a) the spectral power function $\mathbf{x} = (x_1, x_2, \cdots, x_N)^t$ of the light beam of interest, and b) the tristimulus vectors $\mathbf{t}(\mathbf{u}_i)$ of all monochromatic lights of unit energy. The former is easily obtained through purely *physical* measurements, for instance by means of a radiospectrometer, whereas the latter is provided by the results of *psychophysical* colour matching experiments.

Colour matching functions The latter piece of information is obtained by performing the kind of colour matching experiment described earlier (see Figure 2.2, page 15). For each monochromatic stimulus $\mathbf{u}_i$ in the visual spectrum, the intensity factors r_i, g_i and b_i of three fixed reference lights $\mathbf{R}, \mathbf{G}$ and $\mathbf{B}$ which are necessary in order to establish a colour match are determined. Then, the entries of the tristimulus vector of the monochromatic light with respect to the chosen set of reference lights are r_i, g_i and b_i, i.e.

$$\mathbf{t}(\mathbf{u}_i) = \begin{pmatrix} r_i \\ g_i \\ b_i \end{pmatrix} \tag{2.17}$$

Once such a colour match has been made for all monochromatic lights in the spectrum, that is for all values i from 1 to N, we have all the information that we need. If all these experimentally determined tristimulus values are tabulated in an 3 by N matrix

$$\begin{pmatrix} r_1 & r_2 & \cdot & \cdot & r_N \\ g_1 & g_2 & \cdot & \cdot & g_N \\ b_1 & b_2 & \cdot & \cdot & b_N \end{pmatrix} \tag{2.18}$$

where the $i-$*th* coloumn corresponds to the tristimulus vector of the $i-$*th* monochromatic stimulus, then the rows of this matrix are called colour matching functions. Figure 2.5 shows a set of colour matching functions $\overline{r}, \overline{g}, \overline{b}$ established by W. Stiles and Burch (1955). In this particular case, monochromatic reference lights $\mathbf{R}, \mathbf{G}$ and $\mathbf{B}$ with wavelengths of 645, 526 and 444 nm, respectively, were used. Note that, as is particularly prominent in the case of $\overline{r}$, the colour matching functions have negative branches. This corresponds to cases in which one of the reference lights had to be added to the test light in order to establish a match.

The colour matching functions obtained by the above procedure play a fundamental role in colour science and technology. The colour matching functions are generally found to be very similar for the majority of observers (i.e. those with normal trichromatic colour vision). Thus, the mean colour matching functions of a small number of observers can be used as a universal standard which is approximately valid for most observers, and thanks to the validity of the Grassmann Laws, these standard colour matching functions can be used to compute the tristimulus values of any light beam, and thus provide the basis for the ‘objective’ specification or ‘measurement’ of colour.

Change of reference lights The three primary stimuli used to determine the colour matching functions can, in essence, be chosen arbitrarily[4]. However – again due to the Grassmann Laws – two different sets of colour matching functions $\bar{r}(\lambda), \bar{g}(\lambda), \bar{b}(\lambda)$ and $\bar{r}'(\lambda), \bar{g}'(\lambda), \bar{b}'(\lambda)$ obtained by using two different sets of primaries $\mathbf{R}, \mathbf{G}, \mathbf{B}$ and $\mathbf{R}', \mathbf{G}', \mathbf{B}'$, respectively, are related by a simple linear transformation, i.e. there exists an invertible[5] 3-by-3 matrix $\mathbf{M}$ such that

$$\begin{pmatrix} \bar{r}(\lambda) \\ \bar{g}(\lambda) \\ \bar{b}(\lambda) \end{pmatrix} = \mathbf{M} \cdot \begin{pmatrix} \bar{r}'(\lambda) \\ \bar{g}'(\lambda) \\ \bar{b}'(\lambda) \end{pmatrix}, \tag{2.19}$$

and it is easily verified that $\mathbf{M}$ is given as the matrix which has the tristimulus vectors of the primaries $\mathbf{R}'$, $\mathbf{G}'$ and $\mathbf{B}'$ with respect to the primaries $\mathbf{R}$, $\mathbf{G}$ and $\mathbf{B}$ as columns:

Let X be an arbitrary light beam and $(R'_X, G'_X, B'_X)^t$ its tristimulus vector in terms of the ‘primed’ set of reference lights, which by definition means that the colour match

$$X \sim R'_X \mathbf{R}' + G'_X \mathbf{G}' + B'_X \mathbf{B}' \tag{2.20}$$

will hold. Furthermore, let $(R_{R'}, G_{R'}, B_{R'})^t, (R_{G'}, G_{G'}, B_{G'})^t$, and $(R_{B'}, G_{B'}, B_{B'})^t$ be the tristimulus vectors of the primed reference lights with respect to the unprimed ones, which means that the colour matches

$$\mathbf{R}' \sim R_{R'} \mathbf{R} + G_{R'} \mathbf{G} + B_{R'} \mathbf{B} \tag{2.21}$$

$$\mathbf{G}' \sim R_{G'} \mathbf{R} + G_{G'} \mathbf{G} + B_{G'} \mathbf{B} \tag{2.22}$$

$$\mathbf{B}' \sim R_{B'} \mathbf{R} + G_{B'} \mathbf{G} + B_{B'} \mathbf{B} \tag{2.23}$$

are valid. The Grassmann Laws ensures that these expressions can be substitued directly into equation 2.20. A simple rearrangement of terms then yields

$$\begin{aligned} X \quad \sim \quad & (R'_X \cdot R_{R'} + G'_X \cdot R_{G'} + B'_X \cdot R_{B'}) \cdot \mathbf{R} + \\ & (R'_X \cdot G_{R'} + G'_X \cdot G_{G'} + B'_X \cdot G_{B'}) \cdot \mathbf{G} + \\ & (R'_X \cdot B_{R'} + G'_X \cdot B_{G'} + B'_X \cdot B_{B'}) \cdot \mathbf{B}. \end{aligned} \tag{2.24}$$

By definition, the expressions in the above parentheses are the tristimulus values R_X, G_X, B_X of the stimulus X with respect to the unprimed reference lights. Thus, in matrix notation we have

$$\begin{pmatrix} R_X \\ G_X \\ B_X \end{pmatrix} = \begin{pmatrix} R_{R'} & R_{G'} & R_{B'} \\ G_{R'} & G_{G'} & G_{B'} \\ B_{R'} & B_{G'} & B_{B'} \end{pmatrix} \cdot \begin{pmatrix} R'_X \\ G'_X \\ B'_X \end{pmatrix}. \tag{2.25}$$

As this relationship holds for the primed and unprimed tristimulus vectors of any arbitrary light beam X, it also holds for the primed and unprimed tristimulus vectors $(\bar{r}(\lambda), \bar{g}(\lambda), \bar{b}(\lambda))^t$ and $(\overline{r'}(\lambda), \overline{g'}(\lambda), \overline{b'}(\lambda))^t$ of any monochromatic light used in the determination of the colour matching functions. Thus, the matrix $\mathbf{M}$ equals the above transformation matrix, the columns of which, as was postulated, correspond to the tristimulus vectors of the primed reference lights

---

[4] They must, however, be chosen such that neither of them can be matched by a mixture of the others, i.e. their tristimulus vectors must be linearly independent. Apart from this, there is absolutely no limitation on the choice of primaries. Specifically, there is no profound reason dictating that one should use monochromatic reference lights. In practice, though, this is generally done because it reduces the number of cases in which negative mixture coefficients are necessary.

[5] It is presupposed that the tristimulus vectors of the three reference lights in each set are linearly independent.

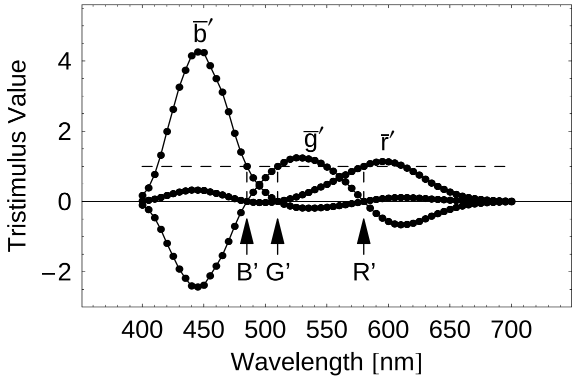

Figure 2.6: The colour matching functions of Figure 2.5 are here expressed in terms of monochromatic reference lights $\mathbf{R}', \mathbf{G}'$ and $\mathbf{B}'$ of wavelengths 580, 510 and 485 nm instead of those originally used. Although these colour matching functions look different from those in Figure 2.5, they provide exactly the same information.

with respect to the unprimed ones.

As an illustration, the set of colour matching functions that one would expect to obtain if reference lights $\mathbf{R}'$, $\mathbf{G}'$ and $\mathbf{B}'$ of wavelengths 550, 500 and 400 nm were used instead of the reference lights $\mathbf{R}$, $\mathbf{G}$ and $\mathbf{B}$ actually used by W. Stiles and Burch (1955) is shown in Figure 2.6. Though these colour matching functions look rather different from those in Figure 2.5, they contain exactly the same information: A given set of colour matching functions provides information about the tristimulus values of all monochromatic lights with respect to *any* conceivable set of reference lights, explicitly for one of them and implicitly for all the others. Thus, there is no profound reason why one set of reference lights should be regarded as more natural than the other.

Colour spaces and chromaticity diagrams The tristimulus vectors of light stimuli can be plotted as points in a three-dimensional space, which is commonly referred to as **colour space** or **tristimulus space**. Only a subset of the points in this space, commonly referred to as the **colour cone** represents tristimulus vectors of physically realisable stimuli. A useful two-dimensional representation of the three-dimensional colour cone is obtained by intersecting it with a suitably chosen plane. In the case of a colourimetric system based on the reference lights $\mathbf{R}$, $\mathbf{G}$ and $\mathbf{B}$ used by Stiles and Burch, it is empirically true that $R + G + B > 0$ for all tristimulus vectors corresponding to real stimuli. Accordingly, the tristimulus vector of any real stimulus, or its positive extension, intersects the unit plane $R+G+B = 1$ in a point $\mathbf{c}$ (see Figure 2.7). The vector $\mathbf{c}$ is obtained by multiplying the tristimulus vector $\mathbf{t} := (R, G, B)^t$ with the scalar factor $1/(R+G+B)$, that is, if r, g and b are the coordinates of $\mathbf{c}$ we have

$$\begin{aligned} r &= R/(R + G + B) \\ g &= G/(R + G + B) \\ b &= B/(R + G + B). \end{aligned} \tag{2.26}$$

These values r, g and b are called the **chromaticity coordinates** of a stimulus with the tristimulus values R, G and B. Clearly, since the point $\mathbf{c}$ is always located in the same two-dimensional plane, no more than two numbers are necessary in order to specify its position uniquely. It is customary to specify just the two first coordinates r and g. Since we always have $r + g + b = 1$, the third coordinate b is then implicitly given as $1 - r - g$. Typically, the chromaticity coordinates r and g are then plotted in a rectilinear coordinate system. In geometrical terms, dropping the b coordinate

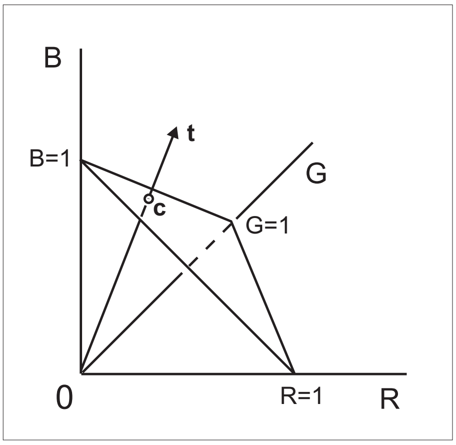

Figure 2.7: In the RGB system, the tristimulus vector $\mathbf{t} := (R, G, B)^t$ corresponding to any real stimulus, or a positive extension thereof ($\alpha \cdot \mathbf{t}$, $\alpha > 0$) intersects the unit plane $R + G + B = 1$ in a point $\mathbf{c}$ (This is true because $R + G + B > 0$ for all real stimuli, such that $\mathbf{t}$ must be located 'above' the plane $R + G + B = 0$).

in this manner amounts to an orthogonal projection of the chromaticity vector $\mathbf{c}$ onto the $B = 0$ plane.

A very important property of chromaticity coordinates is that they remain unchanged when only the intensity of a stimulus is changed, i.e. $\mathbf{c}(X) = \mathbf{c}(\alpha \cdot X)$ for all stimuli X and (positive) scalars α: By the Grassmann linearity law we have $\mathbf{t}(\alpha \cdot X) = \alpha \cdot \mathbf{t}(X)$ and from the above it is clear that $\mathbf{c}(\mathbf{t}) = \mathbf{c}(\alpha \cdot \mathbf{t})$. An immediately useful consequence of this is that the chromaticity of a monochromatic stimulus depends only on its wavelength, and not on its intensity. Thus, the rg-chromaticities of the unit energy monochromatic stimuli used in the determination of the colour matching functions, which are plotted in the upper left panel of Figure 2.8, can be taken to represent the chromaticities of *all* monochromatic stimuli, not only those of unit intensity. Together, these chromaticities representing monochromatic stimuli make up a horseshoe-shaped path commonly referred to as the **spectral locus**. The reason that it consists of single points instead of a continuous path is simply to be sought in the fact that the colour matching functions are based on a partitioning of the spectrum into a finite number of discrete wavelength intervals. The opening between the points at the two endpoints of the horse-shoe, which represents the two endpoints of the visual spectrum, however, is real, and not due to this discrete sampling. No monochromatic stimulus whatsoever will have a chromaticity located between these two endpoints. Binary mixtures of the two monochromatic lights at the endpoints of the visual spectrum, however, do. The monochromatic stimulus at the lower end of the visual spectrum (approx. 400 nm) tends to appear violet, that at the higher end (approx. 700 nm) tends to appear red, and mixtures of these two monochromatic stimuli, which tend to look purple, have chromaticities located on the line connecting the endpoints of the spectral locus. Together with the spectral locus this so-called **purple line** forms a closed path (see upper left panel of Figure 2.8) within which the chromaticity of any real stimulus is located. The region defined by this path is just the intersection of the colour cone with the unit plane $R + G + B = 1$ (as projected orthogonally onto the $B = 0$ plane).

To see that the chromaticities of all real stimuli, whatever their spectral distribution, must be located within this region, one may rely on the fact that the chromaticity of a mixture of two stimuli (with positive intensity coefficients) is located somewhere on the line segment connecting their chromaticities. This is geometrically illustrated in Figure 2.9. Formally, this statement means

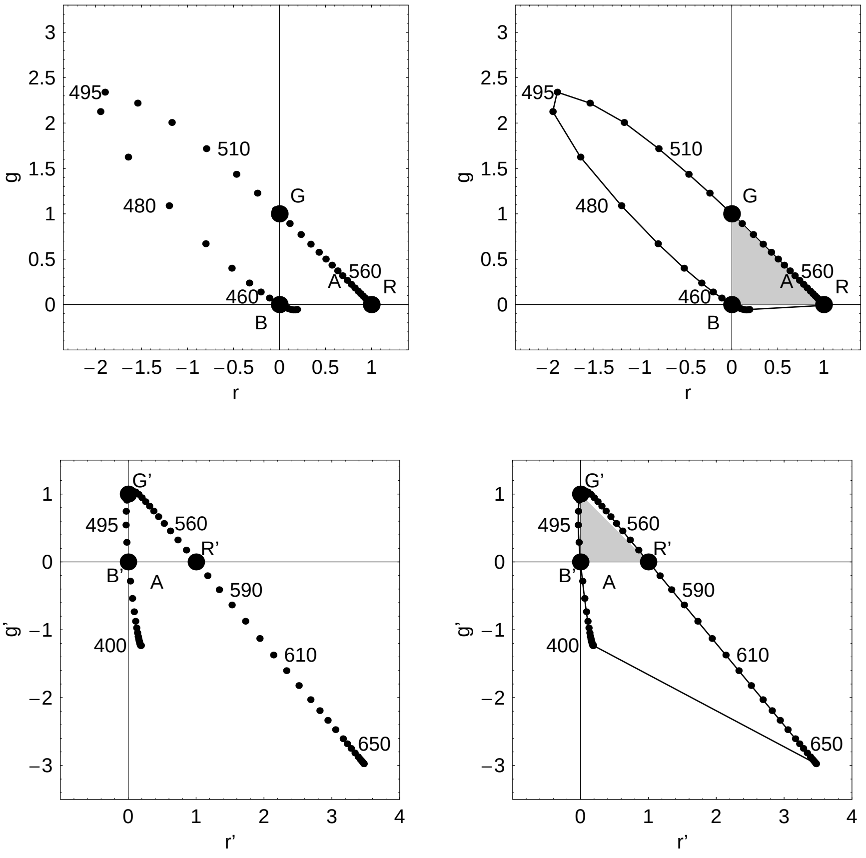

Figure 2.8: Top left: The rg-chromaticities of all monochromatic stimuli ($\Delta\lambda = 5$ nm) are plotted here as small points. This horseshoe-shaped region is called the *spectral locus*. The larger points represent the chromaticities of the reference lights $\mathbf{R}$, $\mathbf{G}$ and $\mathbf{B}$. For selected points the wavelength of the corresponding monochromatic stimulus is shown. Top right: The chromaticites of all physically realisable stimuli are located within the closed path, which consists of the spectral locus and the purple line (the purple line joins the endpoints of the spectral locus). The grey region shows the chromaticities of stimuli which are positive mixures of the reference lights. Points outside of this region corresponds to tristimulus vectors which have a negative component. The letter A shows the chromaticity of an equal energy spectrum. Bottom panels: The same is shown in terms of another set of primaries $\mathbf{R}'$, $\mathbf{G}'$ and $\mathbf{B}'$ (the same ones as those upon which the colour matching functions in Figure 2.6 are based).

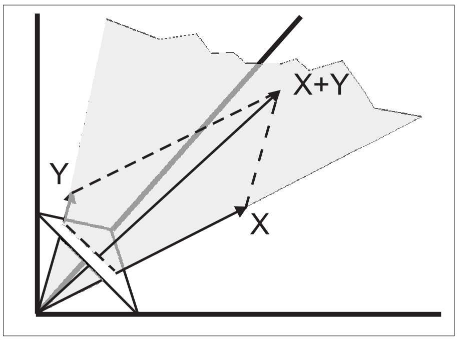

Figure 2.9: The sum $X + Y$ of two vectors intersects the unit plane somewhere on the line segment connecting the intersection points of X and Y. The same holds for any weighted sum $x \cdot X + y \cdot Y$ with positive weights x, y, since this vector must be located somewhere in the sector of a plane ‘between’ X and Y, shown here in gray.

that for any lights X, Y and positive intensity factors x, y we have[6]

$$\mathbf{c}(x \cdot X + y \cdot Y) = \alpha \cdot \mathbf{c}(X) + (1 - \alpha) \cdot \mathbf{c}(Y), \tag{2.30}$$

where $0 \leq \alpha \leq 1$. Exploiting this fact, it is easy to see that the chromaticity $\mathbf{c}(Y)$ of a stimulus Y which is a mixture $Y = x_1 \cdot X_1 + x_2 \cdot X_2 + x_3 \cdot X_3$, $x_i > 0$ of three stimuli must be located somewhere in the triangle (see Figure 2.10) defined by their chromaticities $\mathbf{c}(X_1), \mathbf{c}(X_2)$ and $\mathbf{c}(X_3)$: Defining $P := x_1 \cdot X_1 + x_2 \cdot X_2$, we have $Y = P + x_3 \cdot X_3$, hence $\mathbf{c}(Y)$ is located on the line segment between $\mathbf{c}(P)$ and $\mathbf{c}(X_3)$, and we know that $\mathbf{c}(P)$ is located somewhere on the line segment between $\mathbf{c}(X_1)$ and $\mathbf{c}(X_2)$.

This triangle is the *convex hull* of the points $\mathbf{c}(X_1), \mathbf{c}(X_2)$ and $\mathbf{c}(X_3)$. Informally, the convex hull of a set of points may be understood as the polygon-shaped region that is obtained by spanning a rubber band tightly around them. More generally, it can be shown that the chromaticity of a mixture $Y = x_1 \cdot X_1 + x_2 \cdot X_2 +, \cdots, +x_N \cdot X_N$ of N stimuli, where $x_i > 0$ for all i, must be located in the convex hull of the chromaticities $\mathbf{c}(X_i)$. What needs to be shown for that, is that

---

[6] To prove this, start with the equation $R(A+B) = R(A)+R(B)$, which is valid for all lights A, B by the Grassman linearity law. By eqn. 2.26 we have $R(Z) = r(Z) \cdot S(Z)$, for any light Z, where $S(Z) := R(Z) + G(Z) + B(Z)$. Thus, we have $r(A + B) \cdot S(A + B) = r(A) \cdot S(A) + r(B) \cdot S(B)$. Simple rearrangement yields

$$r(A + B) = \frac{S(A)}{S(A) + S(B)} \cdot r(A) + \frac{S(B)}{S(A) + S(B)} \cdot r(B). \tag{2.27}$$

Defining $\alpha := \frac{S(A)}{S(A)+S(B)}$, this can be rewritten as

$$r(A + B) = \alpha \cdot r(A) + (1 - \alpha) \cdot r(B). \tag{2.28}$$

Strictly analogous reasoning shows that this equation holds for g and b as well, so that all in all we have the vector equation

$$\mathbf{c}(A + B) = \alpha \cdot \mathbf{c}(A) + (1 - \alpha) \cdot \mathbf{c}(B). \tag{2.29}$$

Since this equation holds for all stimuli, we may substitute A by $x \cdot X$ and B by $y \cdot Y$, which, remembering that $\mathbf{c}(x \cdot X) = \mathbf{c}(X)$, yields eqn. 2.30. It remains to show that $0 \leq \alpha \leq 1$. Since $\alpha = \frac{S(x \cdot X)}{S(x \cdot X)+S(y \cdot Y)}$, it suffices to show that numerator and denominator are of the same sign. They are indeed, since $S(X) \geq 0$ for all real stimuli X: In the system based on the reference lights $\mathbf{R}, \mathbf{G}$ and $\mathbf{B}$ of Stiles and Burch, this is empirically true for all monochromatic lights, and therefore also for any mixture thereof, whereby positive mixture coefficients are, of course, presupposed. Note that this may not be the case for systems based on other primaries. In these cases, the unit plane cannot be used for constructing a chromaticity diagram, and another suitable plane must be used instead. In every case, a suitable plane can be found.

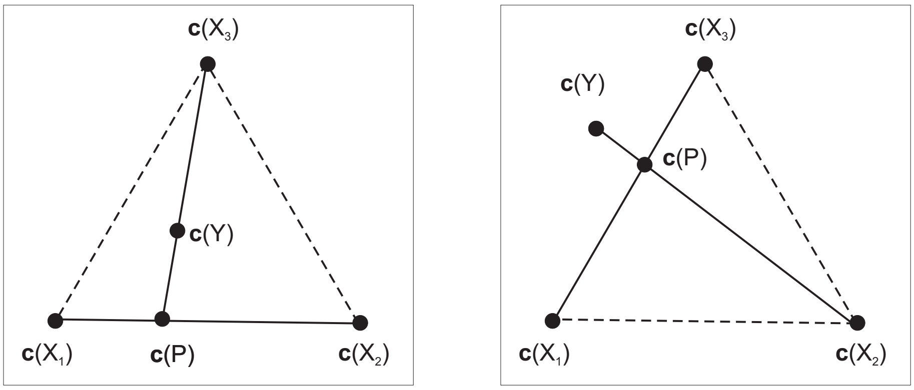

Figure 2.10: Left panel: The chromaticity $\mathbf{c}(Y)$ of a mixture $Y = x_1 \cdot X_1 + x_2 \cdot X_2 + x_3 \cdot X_3$, $x_i > 0$ of three stimuli must be located within the triangle defined by the chromaticities $\mathbf{c}(X_1), \mathbf{c}(X_2)$ and $\mathbf{c}(X_3)$. See text for further explanation. Right panel: Representation of a 'negative' colour match $Y + x_2 \cdot X_2 \sim x_1 \cdot X_1 + x_3 \cdot X_3$ in the chromaticity diagram.

$\mathbf{c}(Y)$ is a *convex combination* of the chromaticities $\mathbf{c}(X_i)$, i.e.

$$\mathbf{c}(Y) = \lambda_1 \cdot \mathbf{c}(X_1) + \lambda_2 \cdot \mathbf{c}(X_2) +, \cdots, \lambda_N \cdot \mathbf{c}(X_N) \tag{2.31}$$

with $\sum_{i=1}^{N} \lambda_i = 1$ and $0 < \lambda_i < 1$ for all i: The convex hull of a set of vectors equals the set of all convex combinations of these vectors *by definition*.[7]

Imaginary reference lights and the CIE system It should now be clear that the chromaticity of any stimulus must be located in the convex hull of the chromaticites of all monochromatic stimuli, that is in the region defined by the spectral locus and the purple line (see Figure 2.8). Clearly, this region, which we shall refer to as the 'region of *real* chromaticities', is not a triangle. If it were, as was erroneously assumed in the early days of colour science, one would expect that positive tristimulus values could be obtained for all light beams provided that one special set of reference lights were chosen, namely those represented by the corners of the triangle; for any other set of reference lights at least some stimuli would by necessity have at least one negative tristimulus value. Under this assumption, then, one particular set of reference lights would be set apart from all others, and it appears natural to think of each of these reference lights as exciting just one of the three cone types. Estimating the spectral sensitivities of the cones based on the colour matching functions would then be straightforward: Indeed, the cone sensitivies would simply be identical to the colour matching functions obtained using this special set of reference lights. However, as the region of real chromaticities is not a triangle, this simple scheme will not work, and as will be discussed in the next section, the estimation of the cone sensitivities ultimately had to be based on other principles.

The non-triangularity of the region of real chromaticities implies that no set of real reference lights whatsoever yields colour matching functions which are positive over the entire spectrum.

---

[7] We have already seen that eqn. 2.31 holds for $N = 2$, since eqn. 2.30 can also be stated as $\mathbf{c}(Y) = \lambda_1 \cdot \mathbf{c}(X_1) + \lambda_2 \cdot \mathbf{c}(X_2)$ with $\lambda_1 + \lambda_2 = 1$ and $0 < \lambda_i < 1$. Starting with this, the validity of eqn. 2.31 for any natural number N can be easily ascertained using mathematical induction. For notational convenience, I just show how to get from $N = 2$ to $N = 3$, how to get from any N to $N + 1$ should then be obvious: We have $\mathbf{c}(x_1 X_1 + x_2 X_2 + x_3 X_3) = \lambda_1 \mathbf{c}(x_1 X_1 + x_2 X_2) + \lambda_2 \mathbf{c}(X_3) = \lambda_1(\lambda_1' \cdot \mathbf{c}(X_1) + \cdot \lambda_2' \cdot \mathbf{c}(X_2)) + \lambda_2 \mathbf{c}(X_3) = \lambda_1 \lambda_1' \mathbf{c}(X_1) + \lambda_1 \lambda_2' \mathbf{c}(X_2) + \lambda_2 \mathbf{c}(X_3)$. It is clear that the mixture coefficents here, namely $\lambda_1 \lambda_1'$, $\lambda_1 \lambda_1'$ and λ_2 add to unity since their sum is $\lambda_1(\lambda_1' + \lambda_2') + \lambda_2$ and we have $\lambda_1 + \lambda_2 = \lambda_1' + \lambda_2' = 1$. It is also clear that the mixture coefficients are all between 0 and 1 since the multiplication of two numbers from this interval are also within it.

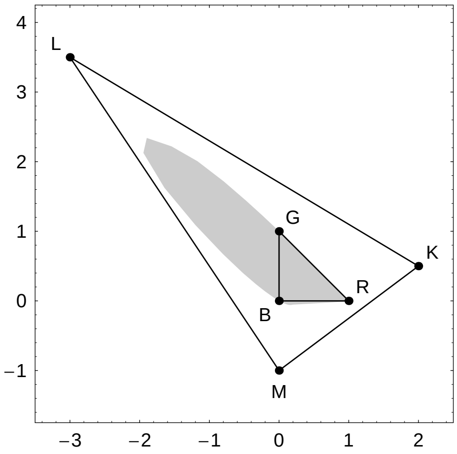

Figure 2.11: Expressed in terms of a set of imaginary reference lights such as $\mathbf{K}, \mathbf{L}, \mathbf{M}$ every physically realisable stimulus has positive tristimulus values, since every real chromaticity is within the triangle KLM. This is impossible with any set of physically realisable reference lights such as $\mathbf{R}, \mathbf{G}, \mathbf{B}$, since these must have chromaticities within the region of real chromaticities (shown here in grey): The resulting triangle can impossibly cover the entire region of real chromaticities.

Thus, irrespective of which set of reference lights is used, it is inescapable that one sometimes has to add one or two of the reference lights to the target light in order to establish a colour match, say,

$$Y + x_2 \cdot X_2 \sim x_1 \cdot X_1 + x_3 \cdot X_3. \tag{2.32}$$

For illustration, the representation of this kind of match in the chromaticity diagram is shown in the right panel of Figure 2.10. In a purely formal sense, though, it is quite possible to obtain all-positive colour matching functions by defining them in terms of so-called **imaginary reference lights**. In Figure 2.11, a triangle KLM has been chosen such that the region of real chromaticities is contained entirely within it. Now, if the points K, L, M have the rg-chromaticities $(r_K, g_K), (r_L, g_L)$ and (r_M, g_M), they may be said to represent the tristimulus vectors

$$\begin{pmatrix} R_i \\ G_i \\ B_i \end{pmatrix} = \begin{pmatrix} r_i \\ g_i \\ 1 - r_i - g_i \end{pmatrix}, \tag{2.33}$$

where $i = K, L, M$. Expressed relative to these tristimulus vectors instead of relative to the RGB primaries every stimulus has all-positive tristimulus vectors. Doing this is completely analogous to specifying stimuli relative to the RGB tristimulus vectors of another set $R'G'B'$ of real reference lights, except that the tristimulus vectors of K, L and M do not represent any physically realisable stimulus. This is why one speaks of *imaginary* reference lights.

As we have already discussed, it is essentially immaterial which set of reference lights is being used for colourimetric specification, since the resulting sets of colour matching functions describe exactly the same information. On the practical side, though, agreeing upon a specific system of colour specification for use in industry and commerce is clearly desirable. The *Commission Internationale de l'Eclairage* (or CIE) have recommended the use of a particular set of colour matching functions $\overline{x}(\lambda), \overline{y}(\lambda), \overline{z}(\lambda)$ which are defined in terms of the imaginary reference lights $\mathbf{X}, \mathbf{Y}$ and $\mathbf{Z}$. The choice of these specific reference lights was based on multifaceted considerations of practical and computational convenience. Since imaginary stimuli were chosen instead of real ones, for instance, every stimulus will have an all-positive tristimulus vector. Furthermore, the $\mathbf{Y}$ reference lights was chosen such that the second colour matching function $\overline{y}(\lambda)$ coincides with the

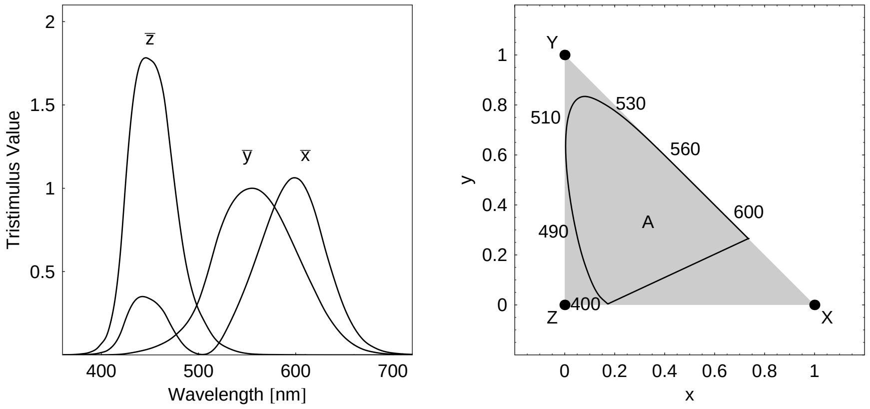

Figure 2.12: Left panel: The CIE 1931 2° colour matching functions $\overline{x}(\lambda), \overline{y}(\lambda)$ and $\overline{z}(\lambda)$. Note that they are all positive at all wavelengths. This is because they are specified in terms of *imaginary* reference lights $\mathbf{X}, \mathbf{Y}$ and $\mathbf{Z}$. Right panel: The corresponding chromaticity diagram, where $x = X/(X+Y+Z)$ and $y = Y/(X+Y+Z)$. The closed line defines the region of real chromaticities. Again due to the use of imaginary reference lights, this region is completely within the triangle XYZ spanned by the reference lights.

photopic luminous efficiency function $V(\lambda)$, which is used in the specification of luminance (see later, page 35 ff.). Other relevant considerations included obtaining a region of real chromaticities which is practical when plotting colourimetric data. The CIE 1931 colour matching functions and the corresponding xy-chromaticity diagram are shown in Figure 2.12.

The cone sensitivities Until now, we have mainly discussed the psychophysical part of trichromatic theory, which, based on the Grassmann laws of trivariance and linearity as first principles, parsimoneously describes the colour matching behaviour of human observers, but the ultimate explanation for these psychophysical facts are to be sought in the physiology of the retina. Clearly, the trivariance of colour matching is accounted for by assuming that a light stimulus entering the eye is coded by three different kinds of receptors, each responding to light stimuli in a characteristic manner, as was suggested very early in the history of colour science (e.g. Palmer, 1777; Young, 1802). The linearity laws can be accounted for by assuming that the response of these receptors L, M and S to an arbitrary light X are uniquely determined by the three integrals in equation set 2.1 on page 13, as already shown in footnote 1 on page 16.

Under this commonly accepted hypothesis, the cone excitation vector $\mathbf{e}(X)$ of a any stimulus X is linearly related to its tristimulus vector $\mathbf{t}(X)$: Assume that the reference lights $\mathbf{R}, \mathbf{G}$ and $\mathbf{B}$ have the cone excitation vectors $\mathbf{e}(\mathbf{R}) := (L_R, M_R, S_R)^t, \mathbf{e}(\mathbf{G}) := (L_G, M_G, S_G)^t$ and $\mathbf{e}(\mathbf{B}) := (L_B, M_B, S_B)^t$, respectively. Furthermore, let an arbitrary stimulus X have the tristimulus values R_X, G_X, B_X, which means that the colour match $X \sim R_X \cdot \mathbf{R} + G_X \cdot \mathbf{G} + B_X \cdot \mathbf{B}$ holds. By the linking proposition, according to which two stimuli appear equal whenever they have the same cone excitation vectors, we then have

$$\mathbf{e}(X) = \mathbf{e}(R_X \cdot \mathbf{R} + G_X \cdot \mathbf{G} + B_X \cdot \mathbf{B}), \tag{2.34}$$

which (using eqns. 2.6 and 2.7 on page 16, which are implied by equation set 2.1 on page 13) can be rewritten as

$$\mathbf{e}(X) = R_X \cdot \mathbf{e}(\mathbf{R}) + G_X \cdot \mathbf{e}(\mathbf{G}) + B_X \cdot \mathbf{e}(\mathbf{B}). \tag{2.35}$$

In matrix notation, this can be rewritten as

$$\begin{pmatrix} L_X \\ M_X \\ S_X \end{pmatrix} = \mathbf{M} \cdot \begin{pmatrix} R_X \\ G_X \\ B_X \end{pmatrix} \tag{2.36}$$

with

$$\mathbf{M} := \begin{pmatrix} L_R & L_G & L_B \\ M_R & M_G & M_B \\ S_R & S_G & S_B \end{pmatrix}. \tag{2.37}$$

Since any triplet $(r(\lambda_0), g(\lambda_0), b(\lambda_0))^t$ of values of the colour matching functions $\overline{r}(\lambda), \overline{g}(\lambda), \overline{b}(\lambda)$ at wavelength λ_0 is just the tristimulus values of a monochromatic light with wavelength λ_0, and the triplet $(l(\lambda_0), m(\lambda_0), s(\lambda_0))^t$ the corresponding triplet of values of the cone sensitivity functions $\overline{l}(\lambda), \overline{m}(\lambda), \overline{s}(\lambda)$, we have

$$\begin{pmatrix} \overline{l}(\lambda) \\ \overline{m}(\lambda) \\ \overline{S}(\lambda) \end{pmatrix} = \mathbf{M} \cdot \begin{pmatrix} \overline{r}(\lambda) \\ \overline{g}(\lambda) \\ \overline{b}(\lambda) \end{pmatrix}. \tag{2.38}$$

From this equation it is clear that each of the three spectral sensitivity curves of the cones must be linear combinations of the colour matching functions, which imposes strong restrictions on their possible shapes. It is also clear that the task of determining the spectral sensitivity functions of the cones reduces to determining the the matrix $\mathbf{M}$.

The reduction hypothesis of dichromatic colour blindness Based on the classical **reduction hypothesis** of dichromatic colour blindness, the cone sensitivity functions can be determined by comparing the colour matches of normal observers with those of colour-blind ones. According to this commonly accepted hypothesis, the three classical classes of colour-blind observers – protanopes, deuteranopes and tritanopes – lack the L-cones, the M-cones and the S-cones, respectively. Importantly, though, the remaining two types of cones are assumed to work in the same way as in normal observers. Since these observers have just two functional types of cones, they are called *dichromats*, whereas normal observers, who have three types of cones are called *trichromats*. Historically, König and Dieterici’s (1886) psychophysical determination of the cone sensitivity functions based on this hypothesis is regarded as seminal. The reduction hypothesis in itself is much older, though, as it was hinted to already by Palmer (see Mollon, 2003). In the following, we shall briefly consider how the reduction hypothesis allows the determination of the cone sensitivity functions based on the colour matches of trichromatic and dichromatic observers. Other techniques can be, have been and are being used but often involve different auxiliary assumptions. The original procedure of König and Dieterici (1886), for instance, is based on the auxiliary assumption that, in the case of dichromatic observers, the two monochromatic lights at the endpoints of the visual spectrum each stimulate just one of the two cones they posses. Some more modern procedures are described in Stockman et al. (1993) and Sharpe and Stockman (1999). The procedure described below involves no auxiliary assumptions and elucidates the basic idea.

Dichromatic observers accept symmetric colour matches made by trichromatic observers, but confuse some stimuli which appear different to a trichromatic observer. Based on the reduction hypothesis, one would of course expect that whenever two stimuli evoke identical M- and S-cone excitations but different L-cone excitations, they should appear different to a trichromatic observer but identical to a protanope, simply because he lacks the L-cones. Thus, any pair of stimuli which appear different to a trichromat but equal to a protanope differ only in L-cone excitation and is called a protanopic **confusion pair**. Similarly, deuteranopic and tritanopic confusion pairs differ only in M-cone and S-cone excitation, respectively.

An important corollary of this is that the confusion pairs of protanopic, deuteranopic and tritanopic observers each uniquely defines the 'direction' of an axis in (trichromatic) colour space corresponding to pure L-cone, M-cone, or S-cone excitation, respectively. If, for instance, two stimuli P and P' are a protanopic confusion pair, with cone excitation vectors $\mathbf{e}(P), \mathbf{e}(P')$, then $M(P) = M(P')$ and $S(P) = S(P')$, hence

$$\Delta\mathbf{e_P} := \mathbf{e}(P) - \mathbf{e}(P') = \Delta L_P \cdot \mathbf{e_L} \tag{2.39}$$

where $\Delta L_P := L(P) - L(P')$ and

$$\mathbf{e_L} := \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix}. \tag{2.40}$$

Defining $\mathbf{N} := \mathbf{M}^{-1}$, we have $\mathbf{t} = \mathbf{N} \cdot \mathbf{e}$ as a direct corollary of eqn. 2.36. Using this, it is easily shown[8] that, for the corresponding difference of tristimulus vectors $\Delta\mathbf{t_P} := \mathbf{t}(P) - \mathbf{t}(P')$, we have

$$\Delta\mathbf{t_P} = \Delta L_P \cdot \mathbf{t_L} \tag{2.41}$$

where

$$\mathbf{t_L} := \mathbf{N} \cdot \mathbf{e_L}. \tag{2.42}$$

Thus, the difference $\Delta\mathbf{t_P}$ of the tristimulus vectors of any protanopic confusion pair always has the same 'direction' in tristimulus space as $\mathbf{t_L}$, and along this direction, only the amount of $L-$cone excitation varies, while the $M-$ and $S-$cone excitations are constant. Clearly, strictly analogous statments can be made for deuteranopic and tritanopic confusion pairs. Thus, the coefficent vector of a tristimulus vector expressed in terms of three basis vectors having the directions of $\mathbf{t_P}, \mathbf{t_D}$ and $\mathbf{t_S}$, may be interpreted as its cone exitation vector. Whereas the direction of these basis vectors is uniquely determined by the confusion pairs, their lengths may be arbitrarily chosen; This merely fixes the units of measurement for L-, M- and S-cone excitation.

Trichromatic theory provides no profound rationale for choosing particular units of measurement, so they may be arbitrarily chosen. If one arbitrarily chooses the values of ΔL_P, ΔM_D and ΔS_T, i.e. the differences of L-,M-, and S-cone excitation corresponding to the protanopic, deuteranopic and tritanopic confusion pairs $(P, P'), (D, D')$ and (T, T'), as units of measurements for L, M and S, respectively, then the matrix N has the corresponding tristimulus vector differences $\Delta\mathbf{t_P}, \Delta\mathbf{t_D}$ and $\Delta\mathbf{t_T}$ as columns, i.e.

$$\mathbf{N} = \begin{pmatrix} \Delta R_p & \Delta R_d & \Delta R_t \\ \Delta G_p & \Delta G_d & \Delta G_t \\ \Delta B_p & \Delta B_d & \Delta B_t \end{pmatrix}. \tag{2.43}$$

To see this, consider that when $\Delta L_P := 1$, eqns. 2.41 and 2.42 yield $\Delta\mathbf{t_P} = \mathbf{N}\cdot\mathbf{e_L}$, and remember that $\mathbf{e_L} = (1, 0, 0)^t$ so that $\mathbf{N} \cdot \mathbf{e_L}$ is just the first column of $\mathbf{N}$. The claim is shown for the two other columns of N by applying strictly analogous reasoning to the deuteranopic and tritanopic confusion pairs.

---

[8] We have $\Delta\mathbf{t_P} := \mathbf{t}(P) - \mathbf{t}(P') = \mathbf{N} \cdot \mathbf{e}(P) - \mathbf{N} \cdot \mathbf{e}(P') = \mathbf{N}(\mathbf{e}(P) - \mathbf{e}(P')) = \mathbf{N} \cdot \Delta\mathbf{e_P} = \mathbf{N} \cdot \Delta L_P \cdot \mathbf{e_L} = \Delta L_P \cdot \mathbf{N} \cdot \mathbf{e_L}$. Defining $\mathbf{t_L} := \mathbf{N} \cdot \mathbf{e_L}$ yields eqn. 2.41.

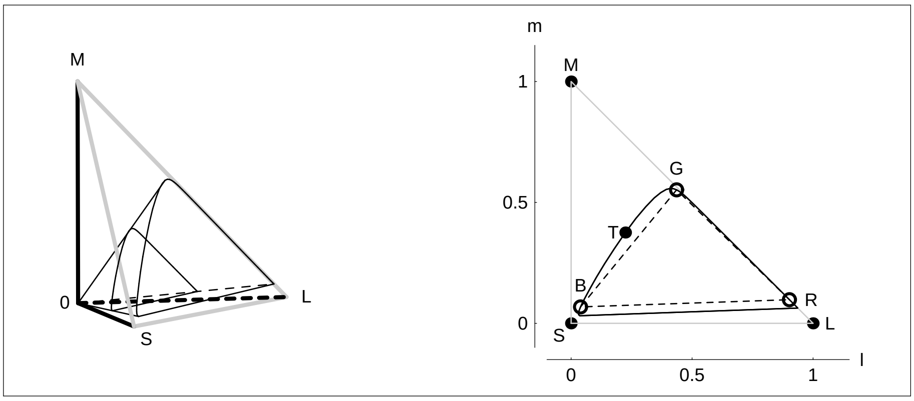

Figure 2.13: Left: The colour cone in cone excitation space, where the axes correspond to pure L-, M- and S-cone excitation, respectively. Note that the axes are outside the colour cone, hence no realisable stimulus stimulates a single cone type alone. Right: a corresponding chromaticity diagram, where $l := L/(L+M+S)$ and $m := L/(L+M+S)$. Analogously to the rg and xy chromaticity diagrams, this diagram is constructed by intersecting the colour plane with the unit plane $L+M+S=1$, and projecting this intersection orthogonally onto the $S=0$ plane. The points L, M and S represent cone excitation vectors having only one non-zero entry: L, M or S, respectively. The points R, G and B, represent the cone excitation vectors of the reference lights $\mathbf{R}, \mathbf{G}$ and $\mathbf{B}$. The stimulus T has one negative coordinate (R) with respect to these reference lights, but only positive cone excitation values.

Once the matrix $\mathbf{N}$ has been determined in this way, the cone sensitivity functions $\bar{l}(\lambda), \overline{m}(\lambda)$ and $\bar{s}(\lambda)$ can be computed based on the colour matching functions using equation 2.38 with $\mathbf{M} = \mathbf{N}^{-1}$. Each of the cone sensitivity functions thus obtained can be multiplied with a scalar in order to change the unit of measurement for each cone type. Criteria often used in order to define a unit of measurement include scaling the cone sensitivity functions such that they all peak at unity, or such that that $L(E) = M(E) = S(E) = 1$ for a particular equal energy stimulus E.

Cone excitation space and cone chromaticity diagrams The cone sensitivity functions according to the estimate of Stockman et al. (1993) have already been shown in Figure 2.1 on page 14. If one takes a look at these curves, it is evident that no monochromatic light stimulates just one of the cone types alone. In fact, every physically realisable stimulus stimulates more than one cone type. This is evident in Figure 2.13, which shows the colour cone in the cone excitation space corresponding to these sensitivity functions, as well as a corresponding chromaticity diagram. In the latter, the triangle LMS represents the all-positive octant of cone excitation space, i.e. all cone excitation vectors for which the entries L, M and S are all positive. The triangle RGB, on the other hand, represents all the cone excitation vectors that can be evoked by a physically realisable mixture of the three reference lights $\mathbf{R}, \mathbf{G}$ and $\mathbf{B}$, i.e. those that correspond to tristimulus vectors in the all-positive octant of the RGB tristimulus space. The interesting lesson to be learned from this is that though many physically realisable stimuli are outside the all-positive octant of tristimulus space, for instance those located on the part of the spectral locus between R and G, they are nevertheless all within the all-positive octant of cone excitation space.

Viewed from this perspective, it should be clear that the notions of ‘imaginary primaries’ and ’negative tristimulus values’, which until now may have appeared rather abstract and meaningless other than in a purely mathematical sense, are actually quite meaningful also in a concrete, empirical sense. From the perspective of a tristimulus space, the axes of cone excitation space can be regarded as ‘imaginary primaries’ but from the perspective of cone excitation space, they obvi-

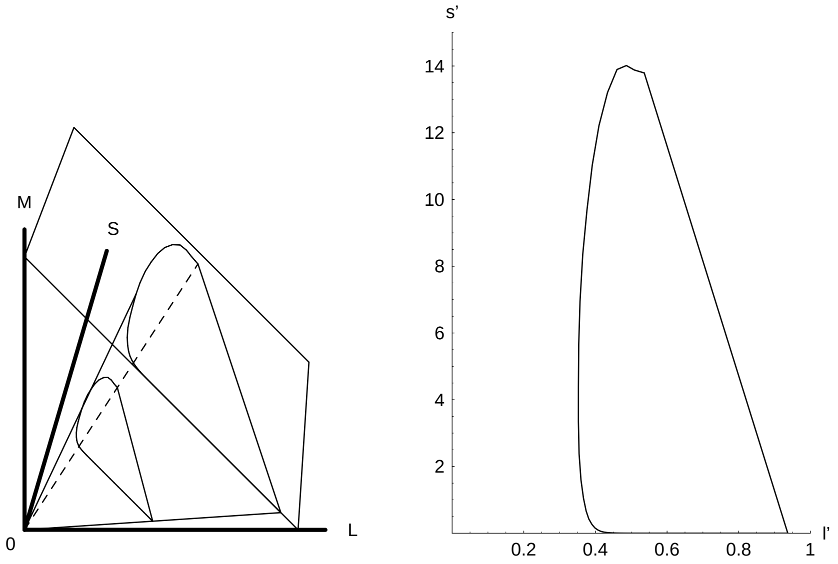

Figure 2.14: Construction of the MacLeod-Boynton (1979) chromaticity diagram. Left: The colour cone is intersected with the plane $L + M = 1$. Right: The resulting chromaticity diagram, where $l' := L/(L + M)$ and $s' := S/(L + M)$.

ously have a fairly concrete physiological interpretation. It should now also be evident that though a particular test light may have one or more negative tristimulus values, in terms of a given set of reference lights, its cone excitation vector will have only positive entries L, M and S. The test light T in the right panel of Figure 2.13, for instance, has a negative tristimulus value R, which is why its cone chromaticity plots outside the triangle RGB. As every other stimulus, though, it is located inside the triangle LMS, which means that it has positive L-,M- and S-cone excitation.

A cone chromaticity diagram that enjoys more widespread use than the one in Figure 2.13 is that devised by MacLeod and Boynton (1979). The construction of this chromaticity diagram is illustrated in Figure 2.14. In principle, a chromaticity diagram can be constructed by using any plane in colour space which intersects the colour cone completely. In the MacLeod-Boynton diagram, the plane $L + M = 1$ is used instead of the unit plane $L + M + S = 1$. Based merely on the basic facts of trichromatic theory, either of these choices is as good as the other. But when one also takes into account experimental findings concerning the so-called luminance mechanism (which will be discussed later, see page 35 ff.), which codes something grossly corresponding to the intensity-dimension of stimuli, using the plane $L + M = 1$ has distinct practical advantages. It has been found that the S-cones do not contribute to luminance. Instead, luminance corresponds to a linear sum of the M- and L-cone excitation (Eisner & MacLeod, 1980). Thus, by choosing appropriate units of measurements for M- and S-cone excitation the luminance of a stimulus is just $L + M$, and thus the plane $L + M = 1$ is a plane of equiluminance, whereas the unit plane $L + M + S$ is not. The chromaticity coordinates of the MacLeod-Boynton diagram are $l' := L/(L + M), m' := M/(L + M)$ and $s' := S/(L + M)$, whereby only m' and s' must be explicitly stated, since $m' = 1 - l'$. An immediately useful property of the MacLeod-Boynton diagram is that – in a plane of equiluminance – the values of l', m' and s' are directly proportional to L-, M-, and S-cone excitation, respectively.

The law of persistence We have now considered the major aspects of trichromatic theory, i.e. the Grassman laws and their implications as well as the physiological linking proposition. Before we turn to issues which are beyond the basic facts of trichromatic theory, a few thoughts on its proper scope may be in order. It is, for instance, misleading to think of points in colour space as representing colour impressions. Instead, a point in colour space, i.e. a tristimulus vector, is more properly thought of as representing an equivalence class of stimuli which all appear indistinguishable to the observer. Indeed every point in the colour cone – except those on the curved part of the spectral locus and in the plane of purples – represents an infinite number of different, but mutually metameric, lights. The reason why it is misleading to think of tristimulus vectors as representing colour impressions, is that the perceived colour of a stimulus depends on the temporal and spatial context in which it is viewed, as evidenced for instance by such phenomena as successive and simultaneous contrast. Thus two stimuli with the same tristimulus vectors may appear different in colour when they are viewed in different contexts. This is the reason why the viewing conditions are kept under strict experimental control in the colour matching experiments which found the basis of trichromatic theory. For instance, the two fields that are to be matched are typically viewed against a background of complete darkness. There is no necessity to this, though. According to the law of persistence ('persistence of optical equations', Kries, 1905), two stimuli which appear equal when viewed against a (common) dark background, will continue to appear equal when viewed against any other (common) background as well. An immediate consequence of this is that when the colour match $\mathbf{R} \sim R \cdot \mathbf{R} + G \cdot \mathbf{G} + B \cdot \mathbf{B}$ is valid when the two fields are viewed against a dark background, it will also be valid when the two fields are viewed against any other background, hence the tristimulus values R, G, B of a stimulus are independent of the background against which the target is viewed. Thus, the same colour space, with the same equivalence classes of stimuli is obtained irrespective of which background the matching stimuli are viewed against, whereas the perceived colour corresponding to the tristimulus vectors may change. It is important, though, that the two stimuli that are to be matched are viewed against the *same* background. If the two targets are viewed against different backgrounds, instead of a common one, a colour match generally breaks down, as in the case of simultaneous contrast.

2.2 Trichromatic theory and colour appearance

Although it is tempting to think of the elements of cone excitation space or colour space as representing colours, everything that the hard facts of trichromatic theory tells us about perceived colour is that physically different lights which map to the same point in colour space appear equal (because they evoke the same cone excitation triplet). *Which* colour impression they evoke can not be deduced from the laws of trichromatic theory (Wright, 1972). Remember that the task of a subject in a classical colour matching experiment is to make the two semi-fields appear equal in colour, he is never actually asked what colour they appear. In order to account for how colour impressions depend on the physical stimulus, then, it is necessary to go beyond trichromatic theory proper and develop further auxiliary theoretical notions.

According to one widespread idea, the perceived colour of any stimulus can be described in terms of the three perceptual variables **hue**, **saturation** and **brightness**. Indeed, in the case of an isolated and homogeneous spot of light viewed against a background of darkness (which are the standard viewing conditions for classical colour matching experiments) these perceptual variables seem to be intimately related to certain dimensions of colour space. Stimuli which vary only in intensity, and therefore have the same chromaticity, differ only in brightness, whereas hue and saturation remain (at least roughly[9]) constant. This observation provides a further motivation for

---

[9] Some deviations from this are known as the Bezold-Brücke effect, see Wyszecki and Stiles (1982, p. 420 ff.).

the use of chromaticity diagrams: They are often said to yield a two-dimensional representation of colour-space where only the brightness-dimension is lacking.

The perceptual variables hue and saturation seem to be represented by a kind of polar coordinates in the chromaticity diagram: Somewhere in the centre of the diagram there is a unique point which represents lights which appear achromatic (hueless, i.e. white, gray or black). Any straight line drawn from this point outwards represents lights which evoke colours of constant hue, whereby perceived saturation increases with the distance from the achromatic point. Accordingly, the most saturated colour impressions are evoked by the lights represented by the spectral locus and the line of purples. As one moves along the spectral locus and the line of purples, which define a closed path enclosing the achromatic point, perceived hue varies in an ordered sequence including all possible hue impressions; one obtains a full ‘color circle’.

In a three-dimensional representation of the colour cone, then, we may say that a central line extending from the origin represents achromatic colours, and that planes of constant hue extend outwards from this line (see Colour Plate. II, page 137), whereby something like the ‘height’ of a point in the cone represents brightness.

2.2.1 Hue, saturation, brightness, and Helmholtz coordinates

The above observations suggest that is could be possible to re-parametrise colour space in a way which yields a simple relation between the perceptual variables hue, saturation and brightness and variables of colour space. A classical re-parametrisation of this kind are the so-called Helmholtz coordinates dominant wavelength, purity, and luminance. We shall now consider each of the variables in turn, and discuss to which extent they capture the perceptual variables.

Dominant wavelength It is obvious from the chromaticity diagram that any stimulus with a chromaticity between the achromatic point and the spectral locus is metameric to a mixture of a monochromatic stimulus and a stimulus represented by the achromatic point. The wavelength of the monochromatic constituent in this mixture is then said to be the dominant wavelength of the stimulus. By this definition, it is clear that it is impossible to assign a dominant wavelength to stimuli with chromaticities between the achromatic point and the purple line. In these cases, it is customary to assign it a *complementary* dominant wavelength instead, which is simply the dominant wavelength of a chromaticity which is located on the opposite side of the achromatic point (see Figure 2.15).

It is essential to note that the dominant wavelength of a stimulus depends on which chromaticity is regarded as achromatic. This is not quite unproblematic, since trichromatic theory has nothing to say about which stimuli appear achromatic. For purely colourimetric purposes, it may be quite sufficient to define a certain stimulus as achromatic per definition, for instance the spectrum of a particular daylight, or an equal energy spectrum, which tend to look at least approximately achromatic. Once this definition is made, any stimulus has a well-defined dominant wavelength which can be calculated based on the colour-matching functions. If, however, dominant wavelength is intended to represent perceived hue, then the notion that lines of constant hue converge on the achromatic point requires that a perceptually defined achromatic point is used for reference. In this case, it might be necessary to define the achromatic point individually, since location of the achromatic point in the chromaticity diagram may differ, at least slightly, for different individual observers, even for who are considered to be colour normal.

The results of psychophysical experiments suggest that lines of constant hue are not quite straight (e.g. Burns, Elsner, Pokorny, & Smith, 1984; Wilson & Brocklebank, 1955). These deviations from linearity are generally referred to as the **Abney effect**. They are sufficiently small to say that to a first approximation, dominant wavelength may be considered an adequate index

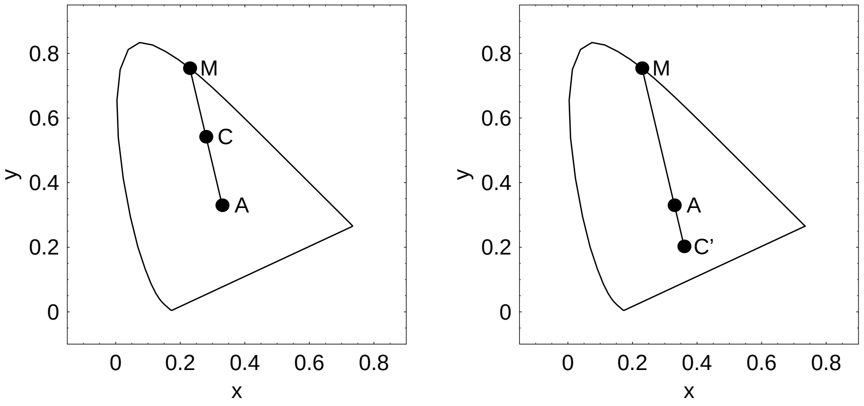

Figure 2.15: Definition of dominant wavelength. Left: For a stimulus with the chromaticity C, a line from the achromatic point A can be drawn through it until it intersects the spectral locus in the point M. The wavelength of the monochromatic stimulus with the chromaticity M is the dominant wavelength of the stimulus with the chromaticity C. Right: For stimuli with chromaticities like C', which are located between the achromatic point and the purple line, this will not work. In these cases, the line from C' through A is extended until it intersects the spectral locus in the point M. The wavelength of the monochromatic stimulus with chromaticity M is then the *complementary* dominant wavelength of the stimulus with chromaticity C'.

of perceived hue as long as it is defined relative to a perceptually defined achromatic point. Apart from this, though, it should be kept in mind that hue differences are far from proportional to differences in dominant wavelength (see, for instance, Wyszecki & Stiles, 1982, p. 570 ff).

Purity The purity of a stimulus is a ratio of distances in the chromaticity diagram indicating how far its chromaticity is displaced from the achromatic point towards the spectral locus or the purple line. If C is the chromaticity of the stimulus, A that of the achromatic point and M the point where a line from A through C intersects either the spectral locus or the purple line, then the purity of the stimulus is $|CA|/|MA|$, where $|AB|$ is the length of the line from A to B (See the left panel of Figure 2.15). Obviously, then, purity is a measure which may assume values between zero (for the achromatic stimulus) and unity (for monochromatic stimuli or stimuli from the purple line). Similar to dominant wavelength, the measure of purity depends on the choice of an achromatic stimulus, but unlike dominant wavelength, which is independent of the choice of a particular chromaticity diagram, calculating the purity of a stimulus in terms of one chromaticity diagram may yield quite different results than calculating it in terms of another. Further qualification is therefore necessary in order to keep different measures of purity apart. If purity is calculated based on the CIE 1931 xy-diagram, the term **excitation purity** is used. The term **colourimetric purity** is used when purity is calculated based on a chromaticity diagram which is constructed using a plane of equiluminance in colour space (the CIE xy-diagram is not), such as for instance the MacLeod-Boynton (1979) chromaticy diagram[10]. In Fig. 2.16, loci of constant exitation purity

---

[10] In the original definition, the colourimetric purity of a tristimulus vector (or, stimulus) $\mathbf{X}$ is defined as follows: Any tristimulus vector $\mathbf{X}$ can be written as a sum $\mathbf{M} + \mathbf{A}$ of a tristimulus vector $\mathbf{M}$ from the spectral locus (or the purple line) and an 'achromatic' tristimulus vector $\mathbf{A}$. The colourimetric purity p_c of $\mathbf{X}$ is then the ratio of the luminance $\ell(\mathbf{M})$ of the 'monochromatic constituent' and the total luminance $\ell(\mathbf{M} + \mathbf{A}) = \ell(\mathbf{M}) + \ell(\mathbf{A})$ of the mixture, i.e. if $\mathbf{X} = \mathbf{M} + \mathbf{A}$, then $p_c(\mathbf{X}) := \ell(\mathbf{M})/(\ell(\mathbf{M}) + \ell(\mathbf{A}))$. To see that this definition is equivalent to the one given in terms of an 'equiluminant' chromaticity diagram, consider the following reasoning: An 'equiluminant' chromaticity diagram is obtained by defining the chromaticity coordinates x_i of an arbitrary tristimulus vector $\mathbf{Y}$ as $y_i := Y_i/\ell(\mathbf{Y})$

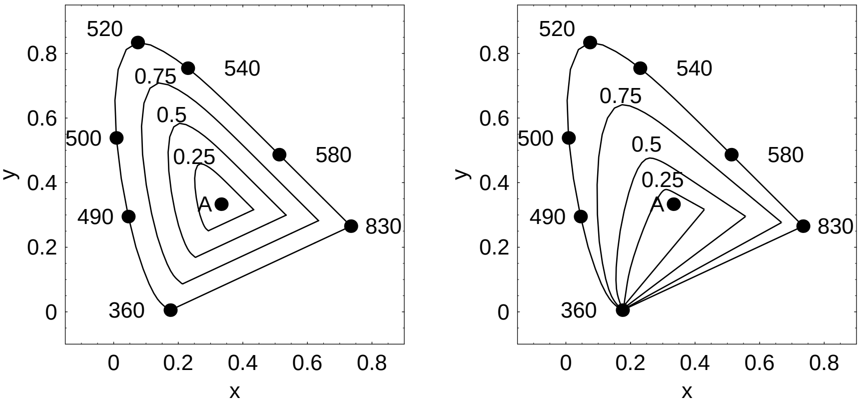

Figure 2.16: Left: Loci of constant excitation purity in the CIE xy-diagram for the purities $0.25, 0.5$ and 0.75. Right: Corresponding loci of constant *colourimetric* purity.

(left) and constant colourimetric purity (right) are shown for comparison in the CIE xy-diagram.

It is well known that monochromatic lights of different wavelengths, which by definition all have a purity of 1, may vary quite considerably in perceived saturation (e.g. Sinden, 1923). Accordingly, it is clear that no measure of purity can be used as an index of perceived saturation which is independent of dominant wavelength. For any fixed dominant wavelength, though, it is probably unproblematic to assume that perceived saturation is a monotonically increasing function of purity.

Luminance As the foundation of the technological discipline of colourimetry, trichromatic theory provides an answer to the question about which stimuli look equal in colour (when they are viewed under equal viewing conditions), namely those which have the same tristimulus vectors. A question with similar theoretical and practical interest, treated within the discipline of **photometry** is: Which stimuli appear equal in **brightness**? Just as colourimetry provides three functions of wavelength which can be used to predict which spectra appear equal in colour, namely the colour matching functions, the aim of photometry is to establish a single function of wavelength which can be used to predict which spectra appear equally bright, i.e. have the same 'luminous efficiency'.

Once a monochromatic stimulus M_{λ_0} of a fixed wavelength λ_0 and intensity m_{λ_0} has been chosen as a reference, the relative luminous efficiency of any other monochromatic stimulus M_λ may be determined, by measuring the intensity m_λ at which it appears to have the same brightness

---

for all $i = 1, 2, 3$. Thus, we have $X_i = x_i \cdot \ell(X), M_i = m_i \cdot \ell(M)$ and $A_i = a_i \cdot \ell(A)$ for all i. Since we have $X_i = M_i + A_i$, this means that $x_i \cdot \ell(X) = m_i \cdot \ell(M) + a_i \cdot \ell(A)$, or, rearranging and using $\ell(X) = \ell(M) + \ell(A)$,

$$x_i = \frac{\ell(M)}{\ell(M) + \ell(A)} \cdot m_i + \frac{\ell(A)}{\ell(M) + \ell(A)} \cdot a_i. \tag{2.44}$$

Defining $\alpha := \frac{\ell(M)}{\ell(M)+\ell(A)}$, this can be rewritten as $x_i = \alpha \cdot m_i + (1 - \alpha) \cdot a_i$ or $x_i = a_i + \alpha \cdot (m_i - a_i)$, hence $x_i - a_i = \alpha \cdot (m_i - a_i)$, which in turn means that

$$\frac{x_i - a_i}{m_i - a_i} = \alpha = \frac{\ell(M)}{\ell(M) + \ell(A)}. \tag{2.45}$$

The left side of this equation corresponds to the definition of colourimetric purity as the ratio of line lengths in the equiluminant chromaticity diagram, the right side is the definition in terms of the luminance ratio.

as M_{λ_0}. The ratio m_{λ_0}/m_λ may then be understood as the luminous efficiency of M_λ relative to that of the reference lights M_{λ_0}. The function $B(\lambda) := m_{\lambda_0}/m_\lambda$ may then be called a luminous efficiency function ('brightness function'). For this brightness function to be practically useful, one would like to be able to compute the brightness ℓ of any spectrum $x(\lambda)$ by multiplying each of its monochromatic constituents with the luminous efficiency for that wavelength and then integrating over the visible spectrum:

$$\ell(x) := \int_\lambda B(\lambda)x(\lambda)\,d\lambda. \tag{2.46}$$

In this equation, ℓ is by definition a linear function of the stimulus, i.e

$$\ell(x(\lambda) + y(\lambda)) = \ell(x(\lambda)) + \ell(y(\lambda)) \tag{2.47}$$

and

$$\ell(t \cdot y(\lambda)) = t \cdot \ell(x(\lambda)) \tag{2.48}$$

for any stimuli $x(\lambda), y(\lambda)$ and intensity factors t. This linearity law is commonly referred to as **Abney's Law** or, highlighting the analogy to the corresponding linearity laws of colour mixture, the **4th Grassmann law**.

Unfortunately, though, it is quite clear that Abney's Law is not valid for equal brightness judgements.[11] As was demonstrated by Abney and Festing (1886), however, an experimental technique called **heterochromatic**[12] **flicker photometry** yields a measure of luminous efficiency which obeys the linearity law. In heterochromatic flicker photometry, the two stimuli which are to be compared are presented in temporal succession at the same location of the visual field. Since the two stimuli are repeatedly interchanged, the observer perceives a periodic change in colour and brightness commonly referred to as flicker. The task of the observer is to adjust the intensity of one of the stimuli such that the perceived flicker is minimised, and at that point the two stimuli are said to have the same luminous efficiency.

The fact that flicker photometry yields a measure of luminous efficiency which obeys Abney's Law is not its only advantage over direct heterochromatic brightness matching. Judging whether two stimuli of different colour appear to be equally bright is notoriously difficult, leading to large variability in the data obtained with different observers. Worse still, measurements obtained with a single observer may also be quite variable. Flicker photometry, on the other hand, yields quite reliable measurements. Furthermore, it has been found that a number of other psychophysical techniques yield measurements which are consistent with those obtained with flicker photometry, suggesting that it is tapping a visual mechanism of more general significance.

Ironically, then, one may say that flicker photometry yields a measure of 'brightness' which has all the properties one would like a measure of brightness to possess, apart from the fact that it does not actually represent perceived brightness. Luminance is a term with many meanings, but the most common usage is probably 'whatever flicker photometry measures'. In order to provide an objective measure of luminance, the CIE has defined a standard luminous efficiency function $V(\lambda)$ for the so-called standard observer (see Fig. 2.17). Using this function, the luminance ℓ of a stimulus $x(\lambda)$ is computed as

$$\ell(x) := k \int_\lambda V(\lambda)x(\lambda)\,d\lambda, \tag{2.49}$$

where k is just a constant defining the unit of measurement.

---

[11] Except under some very special circumstances (cf. Lennie, Pokorny, & Smith, 1993)

[12] Heterochromatic means 'having different chromaticities'. Obviously, judging when two stimuli with *equal* chromaticity is equally bright is completely uninteresting, since this would be the case whenever they have the same tristimulus values.

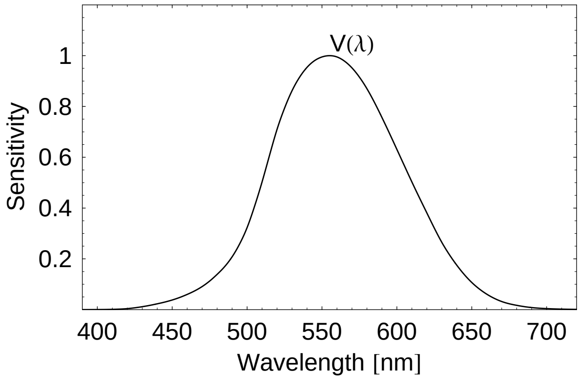

Figure 2.17: The CIE (1924) 2 degree photopic luminosity curve $V(\lambda)$.

Trichromatic theory tells us that the visual mechanism which is responsible for $V(\lambda)$, as well as any other visual mechanism operating under daylight conditions, must receive its input from the three types of cones. The fact that luminance is a linear function of spectra suggests that it is coded by a mechanism which sums cone excitations linearly, that is, the luminance ℓ of any stimulus can be written as a linear sum of its cone excitation values L, M and S, i.e.

$$\ell = l \cdot L + m \cdot M + s \cdot S, \tag{2.50}$$

with fixed weights l, m and s.[13] For any fixed value of luminance ℓ this is just the equation of a plane in cone excitation space, hence stimuli of constant luminance are represented by planes in cone excitation space. Loci of constant brightness, on the other hand, are *not* planes in colour space, but instead curved surfaces. This is because heterochromatic brightness matches violate Abney's Law. The general finding, commonly referred to as the **Helmholtz-Kohlrausch effect**, is that perceived brightness increases with purity for stimuli of constant luminance (Kohlrausch, 1935).

Returning to our original question regarding the relation between the variable brightness and variables of colour space, we may say the following. Luminance, which is a simple linear sum of the basis vectors of colour space, is useful for many practical applications, and is of profound theoretical significance, but it does not represent perceived brightness. It remains a conceivable possibility, though, to find a non-linear function of the variables of colour space that represents perceived brightness. This function would have to be rather complicated though, and take the Helmholtz-Kohlrausch effect into account. The fact that the size of the Helmholtz-Kohlrausch effect depends on the dominant wavelength of the stimulus complicates matters further. A perhaps more fundamental problem, though, is the observation that heterochromatic brightness matches are subjectively difficult to perform and yield unreliable measurements, suggesting that judging the brightness of heterochromatic stimuli is not a very natural task.

2.2.2 Complementarity of colours

A further important aspect of perceived colour which appears to be described (though by no means explained) in a grossly adequate manner by variables of trichromatic theory is that of complementary colours. In terms of trichromatic theory, two stimuli are said to be complementary whenever they can be additively mixed to yield a stimulus which appears achromatic. Expressed in terms of

---

[13] As already noted, psychophysical evidence suggests that $s = 0$, i.e. that the S-cones do not contribute to luminance (Eisner & MacLeod, 1980).

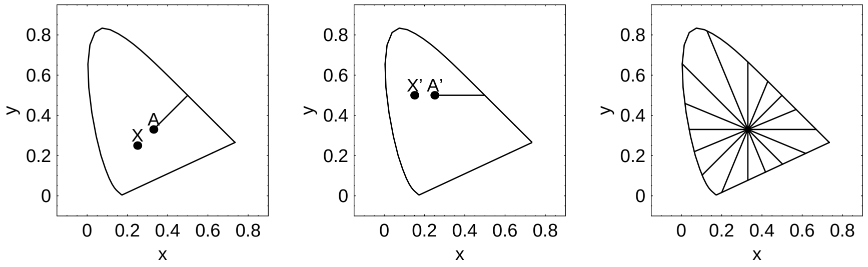

Figure 2.18: Left: The complementaries of the chromaticity X are the chromaticities on the line from the achromatic point A to the spectral locus. Since complementarity is defined relative to the achromatic point, the complementaries of X are located on a line of constant hue (the right-hand panel shows lines of constant hue). Middle: If complementarity were defined relative to another point A' in the chromaticity diagram, the complementaries of a given chromaticity, for instance X', would in general *not* correspond to a line of constant hue.

the chromaticity diagram, this is the case whenever the line segment connecting the chromaticities of the two stimuli passes through the achromatic point. Apart from the choice and definition of the achromatic point, this definition is perfectly rigid and couched completely in terms of the hard facts of trichromatic theory. However, the choice of the achromatic point as a point of reference for defining complementarity can only be justified by observations concerning colour appearance, which go beyond trichromatic theory proper. The notion that lines of constant hue converge on the achromatic point can be used to motivate this particular choice: If complementarity is defined with respect to the achromatic point, then the set of complementaries corresponding to a given chromaticity is a line of constant hue. If, on the other hand, complementarity were defined with respect to any other chromaticity, the set of complementaries of a given chromaticity would in general correspond to a line in chromaticity space corresponding to different hues (see Figure 2.18).

Any two stimuli which are complementary according to this definition are called **additive complementaries**. The notion of complementary colours, though, is much older than trichromatic theory, and stems above all from observations made in connection with the perceptual phenomena of (negative) afterimages and simultaneous colour contrast. When a coloured patch is fixated for some time and then substituted by a neutral background, an afterimage can be perceived (see Colour Plate III on page 138). The colour of the afterimage seems to depend in a characteristic manner on the colour of the patch initially fixated. If one adapts to a yellow patch for instance, the afterimage is usually violet. Conversely, if one adapts to a violet patch the afterimage is yellow. Yellow and violet are therefore called **afterimage complementaries**. Every hue is thought to have a characteristic afterimage complementary in much the same manner as every stimulus has an additive complementary line of constant dominant wavelength. Analogous observations can be made with simultaneous contrast. When a patch which appears achromatic when viewed against a surround of darkness is embedded in a coloured surround, the patch appears tinged in the hue which is the **contrast complementary** of the surround colour.

A critical observation which makes the definition of additive complementaries appear very useful and theoretically interesting is that colours which are afterimage complementaries or contrast complementaries appear to correspond, at least roughly, to stimuli which are additive complementaries in the sense that the colour impression of an afterimage is matched by that evoked by a light which is the additive complementary of the stimulus used to induce the afterimage.[14]

---

[14] I use the term ‘match’ in a somewhat cavalier manner here. Actually it is very difficult, perhaps even impossible,

Though it is generally accepted that afterimage complementaries correspond roughly to additive complementaries, it is debatable how exact this correspondence actually is. As discussed by Wilson and Brocklebank (1955), many researchers have noted that afterimage complementaries sometimes deviate from additive complementaries. The afterimage complementary of yellow is for instance violet, whereas the additive complementary is blue. In their own experiments, though, they found that these deviations could be explained by taking the Abney-shift into account, i.e the fact that lines of constant hue are sometimes slightly curved.

We now conclude our summary of how colour appearance can be described in terms of concepts related to trichromatic theory. As we have already pointed out, the subject of colour appearance is best regarded as a set of hypotheses extraneous to trichromatic theory proper, since a) they are based on a much weaker empirical basis than the Grassmann Laws and the corresponding physiological three-receptor hypothesis and b) their validity may be contested without ever coming into conflict with these more well-established facts.

Whereas colour appearance is beyond trichromatic theory proper, it is the principal subject of Hering's (1920) **opponent colours theory**, which we shall discuss briefly in the following. At the time it was proposed it was both intended and conceived of as an alternative to trichromatic theory. In the course of time though, conceptual clarifications of some earlier misconceptions have made it clear that opponent colours theory is by no means at odds with trichromatic theory proper. The prevailing modern view is that trichromatic theory describes the inital coding of colour in the retina, wheras opponent colour theory describes the processing and recoding of this colour signal at higher levels of the visual system.

2.3 Opponent colours theory

A question which has been discussed for centuries is how many basic colours there are, i.e. how many colours one would need in order to be able to produce any other colour as a mixture of the basic colours. At first consideration, trichromatic theory might seem to provide a definite answer to this. This is, however, only true if the term 'colour' is intended to mean 'light beam', or, alternatively, 'cone excitation vector'. However, as a famous quote of Newton reminds us, "the Rays to speak properly are not coloured"[15], and obviously, one might add, cone excitation vectors aren't either. Ewald Hering (1920), the founding father of **opponent colours theory**, was very aware of this, and couched his answer in terms of colours *per se*, i.e. *qua* colour impressions. According to this theory, there are six basic colour impressions, namely the four **unique hues** red, green, yellow and blue, as well as the two basic achromatic colour impressions black and white.

Hering identified the four unique hues based on the phenomenological observation that they do not involve other hue impressions. Whereas an orange, for instance, may be said to appear both reddish and yellowish, a pure yellow can not be described as a perceptual mixture of other colour impressions[16]. Hence, yellow is a unique hue and orange is not. The same argument applies for the three other unique hues blue, red and green. Every other hue can be described as a perceptual

---

to make a real stimulus appear exactly equal to an afterimage.

[15] "And if at any time I speak of Light and Rays as coloured or endued with Colours, I would be understood to speak not philosophically and properly, but grossly, and accordingly to such Conceptions as vulgar People in seeing all these Experiments would be apt to frame. For the Rays to speak properly are not coloured. In them there is nothing else than a certain Power and Disposition to stir up a Sensation of this or that Colour." (Newton, 1704/1952, Book One, Part II, Prop. II, p. 124 ff)

[16] It is of essential importance that we speak of a *perceptual* mixture here. Couched in terms of a *physical* mixture of lights, this statement is *not* true; It is quite possible to mix a light which appears red with a light which appears green in order to obtain a light which appears pure yellow. Historically, failure to distinguish properly between perceptual colour mixture and the physical mixture of lights was one of the obstacles which had to be overcome in order to appreciate that opponent colours theory is quite compatible with trichromatic theory.

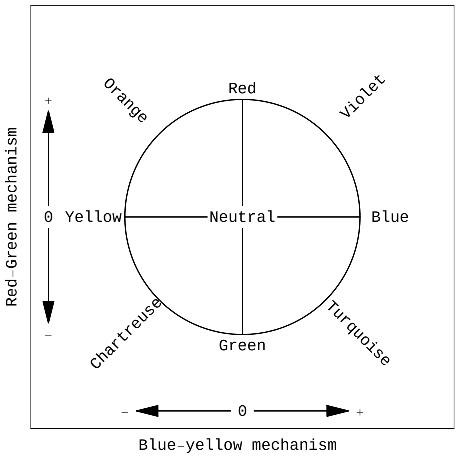

Figure 2.19: How combinations of the two opponent mechanisms ‘blue-yellow’ and ‘red-green’ are thought to give rise to different colour impressions. The diagram can be thought of as a kind of chromaticity diagram, where the achromatic point is in the centre. If the blue-yellow mechanism is in equilibrium, the colours that can be perceived are represented by the central vertical line in the diagram, and include pure greens of all saturations, pure reds of all saturations, and the neutral colour (achromatic color). Which of these colours are perceived is determined by the activity of the red-green mechanism. Note that when both mechanisms are in equilibrium, an achromatic color will be perceived. When only one of the mechanism is in equilibrium, a unique hue is perceived.

mixture of a pair of unique hues. Violet, for instance, can be described as a mixture of red and blue, turquise as a mixture of blue and green, and so on. This is illustrated in Color Plate IV on page 138, where the outer cirle schematically indicates the amounts of two unique hues contained in any hue of the colour circle.

A basic feature of Hering’s theory, which is also evident in this illustration, is that the four unique hues are arranged in two pairs of mutually exclusive, or opponent, colours. No colour impression can be both bluish and yellowish, hence blue and yellow are opponent hues. In the same way, red and green are opponent hues. A pair of non-opponent unique hues, such as blue and green can be (perceptually) mixed to produce a colour impression which contains traces of both, but a pair of opponent hues neutralise or cancel each other. This is the reason why a a unique blue and a unique yellow light, when physically mixed in the right proportions, produce a stimulus which appears achromatic, and not bluish yellow, whereas a physical mixture of two lights with non-opponent hues, say, unique green and a unique yellow, appears both greenish and yellowish.

In order to explain these characteristics of perceptual colour mixture, Hering envisioned two physiological mechanisms, each driven by a pair of antagonistic processes. The blue-yellow mechanism is excited by stimuli with a blue ‘valence’ B, and inhibited by stimuli with a yellow ‘valence’ Y, or, possibly, vice versa; In any event, the two kinds of colour valence influence the mechanism in an antagonistic fashion. If the yellow and blue valences are equally strong, the mechanism is in equilibrium and may be said to yield a zero output, since then $B - Y = 0$. If the blue valence is stronger than the yellow valence, the mechansim yields a positive output, since $B - Y > 0$, in the converse case it yields a negative output, since $B - Y < 0$. The same reasoning applies, the necessary changes being made, to the red-green mechanism.

As illustrated in Figure 2.19, the only colours that can be seen when the blue-yellow mechanism is in equilibrium are unique reds or unique greens, and – if the red-green mechanism happens to also be in equilibrium – neutral (achromatic) colours. Analogously, either a unique blue,

Figure 2.20: Simultaneous brightness contrast: The central squares are physically equal, but the one embedded in the dark surround appears brighter.

an achromatic colour, or a unique yellow can be seen when the blue-yellow mechanism is in equilibrium, depending on the activity of the blue-yellow mechanism. All the other intermediate, non-unique colours result when neither of the mechanisms are in equilibrium; orange for instance, will be perceived when the blue-yellow mechanism yields a negative output and the red-green mechanism a positive output. This scheme also allows for a simple hypothesis of hue perception: Colours of constant hue, but different saturation levels, all have the same ratio of excitations of the blue-yellow and red-green mechanism, respectively.

In addition to the blue-yellow and red-green opponent mechanism, Hering also postulated a black-white mechanism. In contrast to the former, though, this mechanism is not opponent in the sense that it has a neutral point which yields a colour impression which is neither white nor black; white and black yields perceptual mixtures which are some shade of gray, and any shade of gray can be said to resemble both black and white to some degree. In spite of this difference between the two hue mechanisms and the achromatic black-white mechanism, the latter is also referred to as an opponent mechanism since Hering also conceived of this mechanism as a combination of two antagonistic processes. This assumption was primarily based on observations concerning temporal and spatial influences on achromatic colour appearance, i.e. adaptation and simultaneous lightness contrast. Hering pointed to the fact that the sensation of black is not evoked by the absence of light alone, but requires either previous stimulation with a lighter stimulus (adaptation) or the presence of lighter stimuli at other locations in the visual field (simultaneous contrast). This observation can easily be made by turning out all light sources in a room with no windows or other openings through which light may enter. Immediately after turning out the lights, one may have the impression of pitch black, but after some time the sensation will not be pitch black anymore, but instead some (dark) *grey*. Since this sensation of grey appears in the absence of any physical stimulation, it is commonly referred to as ‘intrinsic grey’.[17] It is also sometimes called ‘neural grey’, in accordance with the assumption that it is due to some kind of spontaneous neural activity. The reason why we, in spite of this, actually quite frequently perceive objects as pitch black in everyday life, is that there are generally other objects present in the visual field which reflect more light. How the neighbouring regions which reflect more light may make a given stimulus appear more blackish is demonstrated by the well-known phenomenon of simultaneous lightness contrast (see Figure 2.20). In Hering’s theory, direct stimulation of a point on the retina with light induces whiteness, whereas previous or neighbouring stimulation with light induces blackness, and these two processes interact in an antagonistic manner to determine perceived lightness.

---

[17] German: ‘Eigengrau’ or ‘Eigenlicht’.

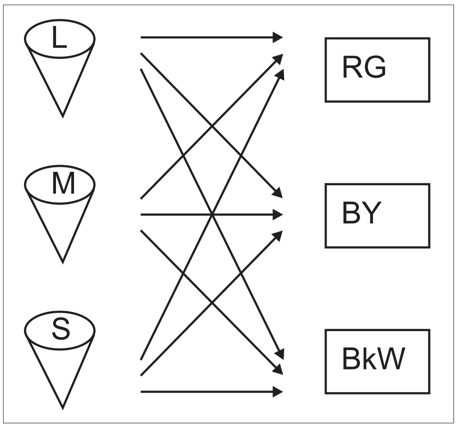

Figure 2.21: The basic idea behind zone theory. The opponent mechanisms (right) combine inputs from the cones (left). The outputs of the cones may influence the opponent mechanism in an excitatory as well as an inhibitory manner.

Zone theory In Hering's original theory, the red-green, blue-yellow and black-white mechanisms were conceived of as receptive elements of the retina, which obviously is in conflict with trichromatic theory. Quite early, though, several authors suggested that Hering's opponent mechanisms may be best sought at higher levels of the visual system. This idea, together with the growing realisation that the three-dimensionality of colour space implied by trichromatic theory is not incompatible with the excistence of four basic colours (or six, if black and white are included), set the stage for the development of so-called **zone-theories** of colour vision (Kries, 1905). Since the opponent mechanisms must directly or indirectly (e.g. via an intermediate stage) be fed by the cones, this idea implies that the response of each of the opponent mechanisms should be a particular function of the cone-excitations. Theoretical work suggested that these functions could even be quite simple. Schrödinger (1920b) and Judd (1949), for instance, showed that hypothetical mechanisms summing (and subtracting) the outputs of the three types of cones could yield the kind of response characteristics one would expect the opponent mechanism to possess (See Figure 2.21).

The basic logic behind this is quite simple. Schrödinger noted that monochromatic stimuli which evoke the impression of unique blue and unique yellow are, within the precision of measurement, additive complementaries. That is, their respective chromaticities are collinear with the achromatic point, or, put differently, the cone excitation vectors of the two unique hues and any achromatic stimulus are all located in the same plane through the origin of cone excitation space. Such a mechanism would be a good candidate for the red-green opponent mechanism, since the theory postulates that unique yellow, unique blue and achromatic colours are the equilibrium colours of the red-green mechanism, i.e. they should be perceived whenever the red-green mechanism is in equilibrium (produces a zero output). For cone excitation vectors outside of the equilibrium plane, the mechanism would produce a positive output on one side, and a negative one on the other, signalling either redness and greenness, respectively. Schrödinger also suggested that the blue-yellow mechanism was linear, defining another null plane in cone excitation space containing the equlibrium colours unique green and unique red (as well as the achromatic colours, of course).

Planes in colour space are defined by an equation of the form

$$l \cdot L + m \cdot M + s \cdot S = 0 \tag{2.51}$$

where l, m and s are fixed real coefficients defining the plane, and L, M and S are cone excitations. Accordingly, a mechanism which sums the cone excitations L, M and S with the weighting factors l, m, s would produce a zero output for cone excitation vectors in this plane. The red-green (RG) opponent mechanism and the blue-yellow (BY) can thus be defined by the coefficients $l_{RG}, m_{RG}, s_{RG}, l_{BY}, m_{BY}$ and s_{BY} in the equations

$$RG = l_{RG} \cdot L + m_{RG} \cdot M + s_{RG} \cdot S \tag{2.52}$$

and

$$BY = l_{BY} \cdot L + m_{BY} \cdot M + s_{BY} \cdot S, \tag{2.53}$$

respectively. Hering's black-white (BkW) mechanism is often identified with the luminance mechanism, which is also defined by a linear sum of the cone excitation values, i.e

$$BkW = l_{BkW} \cdot L + m_{BkW} \cdot M + s_{BkW} \cdot S. \tag{2.54}$$

The upshot of this is that the opponent mechanisms can define a new set of basis vectors in cone excitation space, and when vectors in cone excitation space are expressed relative to these basis vectors, their entries represent the activity of the three opponent mechanisms. Rewriting the above in matrix notation, we have

$$\begin{pmatrix} RG \\ BY \\ BkW \end{pmatrix} = \mathbf{M} \cdot \begin{pmatrix} L \\ M \\ S \end{pmatrix} \tag{2.55}$$

with

$$\mathbf{M} := \begin{pmatrix} l_{RG} & m_{RG} & s_{RG} \\ l_{BY} & m_{BY} & s_{BY} \\ l_{BkW} & m_{BkW} & s_{BkW} \end{pmatrix}. \tag{2.56}$$

Thus, if a stimulus X has the tristimulus vector $\mathbf{e}(X)$, then the three entries of the vector $\mathbf{o}(X) := \mathbf{M} \cdot \mathbf{e}(X)$ represent the activity of the three opponent mechanisms.

The elements of $\mathbf{M}$ defining the RG mechanisms can, in principle, be determined once the cone-excitation vectors of two different colours which are in red-green equilibrium (say, unique yellow and achromatic) are known. The null plane of the RG mechanism is then uniquely defined by the three points given by these two vectors and the origin. The elements of the normal vector of this plane are the coefficents l_{RG}, m_{RG}, s_{RG} of eqn. 2.52. However, if the vector $n_{RG} := (l_{RG}, m_{RG}, s_{RG})^t$ is normal to the null plane, then any vector $c \cdot n_{RG}$, where c an arbitrary scalar factor is also normal to this null plane. Accordingly, the null plane determines the coefficients l_{RG}, m_{RG}, s_{RG} defining the RG mechanism only up to scalar multiplication with a common factor. The same holds analogously for the BY mechanism. This means that knowledge of the null planes of the two mechanisms, and their respective normal vectors gives us

$$RG = c_{RG} \cdot (l_{RG} \cdot L + m_{RG} \cdot M + s_{RG} \cdot S) \tag{2.57}$$

and

$$BY = c_{BY} \cdot (l_{BY} \cdot L + m_{BY} \cdot M + s_{BY} \cdot S), \tag{2.58}$$

where c_{RG} and c_{BY} are the only unknowns. In opponent colours theory, the explicit goal is to describe colour appearance quantitatively in terms of the responses of the opponent mechanisms.

This can be used to fix the relative scaling of the opponent colour mechanisms, i.e. the ratio c_{RG}/c_{BY}: So-called binary hues, such as orange, violet, turquoise and chartreuse are each conceived of as having equal amounts of two unique hues. A 'true orange', for instance, is assumed to be 50% red and 50% yellow. Accordingly, it makes sense to choose the ratio of c_{RG} to c_{BY} such that $RG = BY$ for a stimulus which appears a true orange.

Once the matrix $\mathbf{M}$ is known, one may use it to compute spectral sensitivity curves of the opponent mechanism based on the cone sensitivity functions. In Colour Plate V sensitivity curves of the opponent mechanisms based on an estimate of this matrix originally proposed by Judd is shown. The spectral sensitivities of the opponent colour mechanisms are often called **valence** curves, since they are thought to reflect the 'amounts' of the unique hue sensations evoked by a spectral stimulus. This assumption is central to opponent colours theory, and makes it a general theory of **colour appearance**. As we have seen, the opponent mechanisms can be thought of as emerging through a simple basis transformation of cone-excitation space. As we have already seen, such basis transformations are part and parcel of trichromatic theory. But opponent colours theory goes beyond trichromatic theory in claiming that this particular basis is **natural for describing perceived color**[18].

The popularity of opponent colours theory today is without doubt largely due to the seminal work of Dorothea Jameson and Leo M. Hurvich. In a series of influential papers, they suggested that several phenomena of colour vision can be parsimoneously accounted for within the framework of opponent colours theory, including purity discrimination, wavelength discrimination and the Bezold-Brücke shift (Jameson & Hurvich, 1955; Hurvich & Jameson, 1955).

The notion that the valence curves of the opponent mechanism yield a direct representation of perceived colour, suggest that they can be used to predict the colour appearance of any light stimulus (disregarding the influence of adaptation, of course). The so-called hue coefficients defined by Hurvich and Jameson (1955) suggest themselves as useful for this purpose. According to theory, any hue can be thought of as a perceptual mixture of two unique ones. Any violet for instance, can be thought of as a mixture of unique blue and unique red. How close the violet is to either of these unique hues can be quantified by the ratio of its blue valence to the sum of the blue and the red valence, i.e. $B/(B + R)$. Obviously, if this coefficient is above 0.5 there it has more blue valence than red valence, if it is less, the converse is true. Clearly, the alternative measure $R/(B + R)$ yields the same information, since $B/(B + R) + R/(B + R) = 1$. Generally, one may define the hue coefficients by the expressions

$$K_{RG} := |RG|/(|RG| + |BY|) \tag{2.59}$$

and

$$K_{BY} := |BY|/(|RG| + |BY|). \tag{2.60}$$

Depending on the sign or RG, K_{RG} may be thought of as giving the proportion of redness or greenness, respectively, and the sign of BY makes K_{RG} a measure of either blue or yellow content. The spectral curves of these hue coefficients are shown in Colour Plate VI on page 139.[19] In theory, these curves can be used to predict which colours observers actually perceive when viewing a given stimulus. If the observer views a monochromatic stimulus with a wavelength of 600 nm for instance, the curves in Color Plate VI predict that he should perceive it as 50% red and 50% yellow, i.e. a true orange. Colour naming experiments in which the observers are asked to state their 'perceived hue coefficients' explicitly (e.g., for a violet stimulus, 'What is the percentage of

---

[18] Hering (1920) spoke of a 'natural colour system'.

[19] Note that $K_{RG} = 1 - K_{BY}$, as follows immediately from equations 2.59 and 2.60.

blueness in this colour?') do indeed yield results which are grossly in line with this prediction (Werner & Wooten, 1979).

Without doubt due to the fact that opponent colours theory deals explicitly with colour appearance it has received much attention by linguists and anthropologists, who are interested in such questions as how colour is represented by language or whether colour terms are culturally determined or reflect specific features of our visual system which are shared across cultural borders.

The physiological basis of opponent colours theory Though opponent colours theory is, in the first instance, a phenomenological theory of colour appearance, Hering also ventured hypotheses regarding its physiological basis. From the observation that no colours are simultanously perceived as red and green or blue and yellow, he drew the conclusion "that a physiological process having redness and greenness, or, alternatively, yellowness and blueness, as its perceptual correlate is either not conceivable at all, or but possible under quite special and unusual circumstances" (Hering, 1920, p. 49, my translation)[20].

Later on, a plethora of physiological findings have documented cells in the visual pathway which responds positively to some wavelengths and negatively to others. Cells which show such a response characteristic, generally referred to as spectral opponency, suggest themselves as the physiological substrate of Hering's opponent mechanisms. An early finding of this kind is that of Svaetichin (1956), who recorded from cells in the goldfish retina. Sveatichin originally believed that he was recording from cone receptors, but further investigations indicated that he had probably been recording from horizontal cells, which combine the outputs of different cones (Kaiser & Boynton, 1996). Spectrally opponent cells were later also found in the LGN of the macaque monkey by De Valois, Abramov, and Jacobs (1966). They presented monochromatic stimuli from the visual spectrum to the animals and measured how the firing rate of LGN cells changed with wavelength.

As suggested in the schematic plots of Figure 2.22, a typical finding was that the cells responded with a firing rate below the spontaneous firing rate[21] at some wavelengths and with an increased firing rate at others. Some of the cells responded with excitation to long wavelengths and inhibition to short wavelengths, whereas others were excited by short wavelengths and inhibited by long wavelengths. For both kinds of cells, the point in the spectrum where the sign of the response changed was quite variable. However, for both kinds of cells, the distribution of the crossover points was clearly not Gaussian, and there was some indication that the distributions were bimodal, suggesting that both kinds of cells could by categorised into two subclasses according to whether the crossover point was at a low wavelength or at a high one. According to this classification, the four kinds of cells were labelled with the terms +R-G, -R+G, +B-Y and -B+Y. As a gross illustration, the typical response pattern of each of these types of cells are shown schematically in Figure 2.22. The former two types have their crossover points at a high wavelength (approximately where unique yellow is) and the latter have their crossover points at a lower wavelength (approximately where unique green is).

De Valois et al. (1966) also found cells in the LGN which were not spectrally opponent, i.e. they responend with increased firing rate to all wavelengths. The mean spectral response rate of these non-opponent cells was grossly similar to the spectral sensitivity of the psychophysically determined luminance function $V(\lambda)$. These cells suggest themselves as a candidate neural subtrate for the black-white mechanism.

Generally, such neurophysiological findings clearly support one aspect of Hering's opponent

---

[20] In this connection, it may be interesting to note that some later investigations suggest that under 'special and unusual circumstances' two opponent colours can indeed be perceived simultaneously (Crane & Piantanida, 1983; Billock, Gleason, & Tsou, 2001).

[21] Many cells fire at a certain rate even in the absence of any stimulation. This is the spontaneous firing rate.

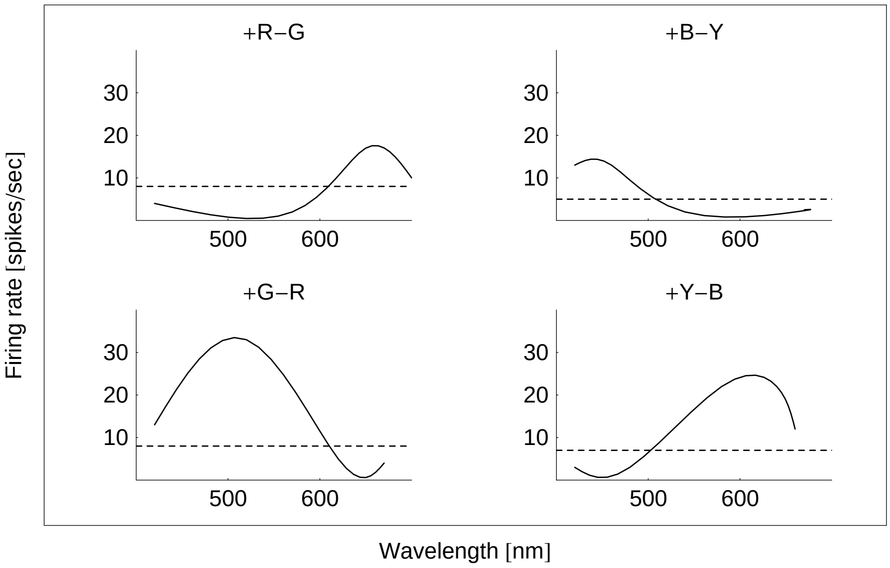

Figure 2.22: Schematic illustration of typical results from De Valois, Abraham and Jacob's (1966) study. The curves here are drawn by hand and only give a rough indication of their original results. Each panel represents the firing rate of one neuron in response to different monochromatic lights from the visual spectrum. The curves have positive and negative lobes which correspond grossly to the positive and negative lobes of the psychophysically defined valence curves of opponent colours theory. The categorisation of different neurons are based on this similarity. Note that two kinds of opponent neurons correspond to each opponent mechanism (e.g. $+R - G$ and $+G - R$ for the red-green mechanism), and that the range of possible negative firing rates (firing rates below the spontaneous firing rate of the neuron) is more limited than the range of possible positive firing rates.

colours theory, namely the assumption of an antagonistic, or **opponent** response, but remain somewhat unclear concering the issue of the **unique hues**. Whereas many studies have documented spectral opponency in ganglion cells of the retina and cells of the LGN, their chromatic properties are rather diverse. This diversity is not what one would intuitively expect based on Hering's theory of unique hues. Instead, the theory of unique hues would suggest the existence of spectrally opponent cells which clearly fall into two distinct classes corresponding to the red-green and the yellow-blue opponent mechanism, respectively. In the De Valois et al's (1966) experiment, for instance, one would expect the loci of the crossover points (zero-crossing in Figure 2.22) to be very similar within each class of cell, and ideally this crossover point should correspond closely to the wavelengths of the unique hues. But the crossover points obtained in De Valois et al's study vary so much that the classification into BY and RG cells is far from unproblematic, as pointed out by the authors themselves (p. 969).

In a more recent study, though, Derrington, Krauskopf, and Lennie (1984) found that cells in the parvocellular layer of the macaque LGN fell into two distinct groups with more modest variation within each of the groups. However, the null-planes of these chromatic mechanisms were distinctly different from those one would expect from opponent colours theory. In this respect they agree with the results of the psychophysical experiments of Krauskopf, Williams, and Heeley (1982), who found two axes in chromaticity space to be independently adaptable, suggesting the existence of two neural mechanisms. The axes found in this study also clearly differ from the axes one would expect based on the unique colours of opponent colours theory (see also Mollon & Cavonius, 1987).

In summary, then, we may say that while there is a plethora of evidence for neural mechanisms exhibiting spectral opponency, there is little clear evidence for mechanisms which respond to those axes in colour space which correspond to the unique hues, and the neural mechanisms behind the unique hues of opponent theory have yet to be identified (Valberg, 2001).

The failure to identify a clear neural substrate of unique colours, however, does obviously not have any implications concerning the validity of the psychophysical theory. Assuming that the psychophysical theory is in itself correct, it simply means that its neural substrate has yet to be found. Opponent colours theory plays an important role in many psychophysical investigations on colour vision, both as a theoretical framework and as a convenient tool. We shall therefore now consider some relevant psychophysical findings.

Linearity of the opponent mechanisms When Schrödinger (1920b) and Judd (1949) showed how opponent colours theory could be related to trichromatic theory, they developed models which assumed that the opponent mechanisms were linear combinations of the cone inputs. With the empirical evidence available at the time, this was probably not an unreasonable assumption. Whether this is in fact true or not, however, is of crucial theoretical and practical significance (Krantz, 1975a). If linearity can be presupposed, it would mean that the response functions of opponent colour theory inherit the properties of the colour matching curves or the cone sensivities which makes them so useful for 'objective' colour specification. Due to the Grassmann Laws, the colour matching curves, which are only given for monochromatic stimuli, can be used to calculate the tristimulus values of any light stimulus. Similarly, linearity of the opponent mechanisms would mean that the opponent response of any light stimuli can be computed based on the opponent response curves for monochromatic lights.

The available empirical evidence suggests, however, that linearity can not be presupposed. In rigorous psychophysical experiments Larimer, Krantz, and Cicerone (1974, 1975) found that though the red-green mechanism appeared to be linear, the blue-yellow mechanism is clearly not. A non-linearity of the blue-yellow mechanism is also suggested by the findings of Werner and Wooten (1979) and Webster, Miyahara, Malkoc, and Raker (2000). In the data of the latter study, there is also some indication that the red-green mechanism may also be nonlinear.

The results of Webster et al's (2000) study are also of special interest since they question the assumption that opponent mechanisms which are simple reweightings of the cone inputs can account for the loci of the opponent hues. If this were the case, one would expect interindividual variations in the loci of oppoent hues to be due to individual differences in receptor sensitivities or the weighting factors. Accordingly, across observers, the locus of one unique hue should correlate with the loci of other unique hues. The data of Webster et al. (2000), however, are clearly at odds with this expectation.

Unique hues as a tool for studying adaptation and context effects The psychological primacy of unique colours postulated by opponent colours theory makes them appear very useful for psychophysical studies of adaptation and simultaneous contrast. As demonstrated by the phenomenon of simultaneous contrast, for instance, the perceived colour of a given patch of light may change radically when the context in which it is embedded is changed. So called asymmetric matching is one standard method for quantifying the influence of the context. In this procedure the subjects view two target patches a and b embedded in different surrounds A and B, respectively, and are asked to change the tristimulus values of patch a such that it appears equal in colour to b. The difference in the tristimulus values of a and b may then be taken as a measure of the context effect. This method has some drawbacks, however. If both stimuli are presented simultaneously, then one may not safely assume that the perceived colour of a is only influenced by A since B is also present in the visual field and may also have an influence. If the stimuli are presented sequentially, on the other hand, the observer has to rely on memory to establish a match, which introduces another possible source of error.

For these and similar reasons one often uses the technique of grey settings instead, in which the observers are asked to adjust the chromaticity of the target patch until it appears achromatic. The difference of the chosen chromaticities for different surrounds may then be taken as a measure of the context effect. In this technique no external comparison stimulus is needed, since the observer judges the patch to be grey according to a fixed internal criterion. The feasiblity of this method rests on the fact that in any plane of constant luminance, there is only one chromaticity which appears achromatic, whereas there are many chromaticities which can be said to have a given chromatic colour. Asking the observers to make the target appear violet, for instance, makes the task very difficult and vague since there are many chromaticities which may be said to look violet in some way. A drawback of this method is, however, that the context effect can only be measured for stimuli which appear grey, and it cannot be ascertained whether the influence of the surround is comparable for other colours.

Opponent colours theory, however, suggest that this problem can be ameliorated, since it contends that besides grey, the unique colours are also psychologically special and easy to identify. Studying how the loci of the unique colours are influenced by context therefore suggests itself as a further method which has basically almost the same virtues as achromatic settings but allows one to measure the effect for a wider range of target colours. Accordingly, unique hue settings have been employed in a number of studies on context effects in colour vision (e.g. Walraven, 1976; Shevell, 1978; Mausfeld & Niederée, 1993). Opponent colours theory suggests that it should be easier for an observer to adjust the colour to be ‘a good yellow’ than, say, ‘a good violet’, since in the former case he can rely on the criterium that a unique yellow should be ‘neither reddish nor greenish’, whereas no comparable criterium can be offered for non-unique hues. In contrast with this expectation, though, Malkoc, Kay, and Webster (2001) reported that unique hue settings were no more reliable than settings for binary hues. Their conclusion that the obtained “pattern of results does not support a perceptual organization in which the dimensions of red-green, blue-yellow, and bright-dark have superordinate status” poses a serious problem for opponent colours theory.

A further related problem with unique hue settings is that they are actually much more variable than one would perhaps expect based on many experimental reports when considered in isolation. A recent meta-analysis of Kuehni (2004) documents and discusses the large ranges of unique hue settings obtained in a number of studies. The impressive size of this variability becomes clear in his statement that “one observer’s unique blue can be another’s unique green and vice versa, and the same for yellow and green.”

Chapter 3

Models of simultaneous contrast

> The assumed validity of the von Kries formulation may have led investigators to believe that the change in color appearance resulting from a chromatic background is a well-understood phenomenon. This is false.
>
> - STEVEN K. SHEVELL (1978)

As we have seen, trichromatic theory provides the means for predicting which light beams will look identical in colour. In itself, it does not provide any means for predicting what colour a given stimulus will appear, but the notion that perceived hue, saturation and lightness correspond to dominant wavelength, purity and luminance, respectively can be said to offer at least a gross indication of the relation between points in colour space and colour appearance. However, trichromatic theory only deals with how the local stimulus influences colour perception under a fixed state of adaptation, it does not deal with how colour perception is influenced by temporal and spatial context. As is well known, these influences can be quite dramatic, and are ubiquitous in real-life colour perception.

In order to develop a more complete theory of colour vision, these important influences must obviously also be accounted for. Although a great number of context effects in colour vision are known, and there is ample reason to believe that they cannot all be explained by a common mechanism, the simultaneous contrast effect shown in Colour Plate I is, for different reasons, often considered to be a paradigmatic case of spatial context effects. On the one hand, this is just a matter of terminology. Le Grand (1957), for instance, uses the term 'simultaneous contrast' collectively to refer to *all* spatial context effects in colour vision. On the other hand, the phenomenon demonstrated in Color Plate I is often thought to be the consequence of a perceptual mechanism which plays a fundamental role in most of everyday colour perception. Helmholtz (1911) and Hering (1920), for instance, both argued that simultaneous contrast, defined more narrowly as 'the effect observable in displays like Colour Plate I', as well as the biologically important phenomenon of colour constancy, are due to a common perceptual mechanism.

Most modern theories of simultaneous contrast can be said to be developments, refinements or combinations of the theories of Helmholtz and Hering, respectively (Kingdom, 1997). Following Mach (1866), Hering favoured an explanation in terms of interactions between neural elements of the retina, commonly referred to as **lateral inhibition** in the modern literature. Helmholtz, on the other hand, advocated a theory based on his general concept of perception as the product of **unconscious inferences**. These two theories differ in many interesting respects, and much empirical work has been devoted to evaluating the relative merit of the two theories (Kingdom, 2003). At present, though, we shall not delve into this difficult and multi-faceted issue, but constrain ourselves to noting that good arguments can be made in favour of both theories[1]. Our primary concern here shall instead be quantitative models of simultaneous contrast, which, though often formulated within a certain theorical framework are often more eclectic in nature.

3.1 The von Kries model

A classical model of the influence of context on perceived colour, commonly referred to as the **von Kries coefficient scheme** (Kries, 1905), can be considered a simple extension of trichromatic theory. According to this model, the role of the context is to change the overall sensitivity of each of the cone types L,M,S. If this is the case, the cone excitation values evoked by a stimulus cannot in general be computed in the conventional manner, i.e. as

$$L_x = \int l(\lambda)x(\lambda)\, d\lambda \tag{3.1}$$

$$M_x = \int m(\lambda)x(\lambda)\, d\lambda \tag{3.2}$$

$$S_x = \int s(\lambda)x(\lambda)\, d\lambda, \tag{3.3}$$

since one would have to use corrected cone sensitivities $l'(\lambda), m'(\lambda)$ and $s'(\lambda)$, instead of the standard cone sensitivities $l(\lambda), m(\lambda)$ and $s(\lambda)$. It can readily be assumed that the corrected cone sensitivities must be rescaled versions of the original ones, since if this were not the case, stimuli which are metameric in one state of adaptation would fail to be metameric in other states of adaptation, as is not the case, according to the law of persistence (see page 32). Accordingly, we have $l'(\lambda) = \rho_l \cdot l(\lambda), m'(\lambda) = \rho_m \cdot m(\lambda)$ and $s'(\lambda) = \rho_s \cdot s(\lambda)$, where ρ_l, ρ_m and ρ_s are positive scaling factors, and the corrected cone excitation values of a stimulus $x(\lambda)$ can be computed as

$$L'_x = \int \rho_l \cdot l(\lambda)x(\lambda)\, d\lambda \tag{3.4}$$

$$M'_x = \int \rho_m \cdot m(\lambda)x(\lambda)\, d\lambda \tag{3.5}$$

$$S'_x = \int \rho_s \cdot s(\lambda)x(\lambda)\, d\lambda. \tag{3.6}$$

Since the scaling factors $\rho_i,\ i = l, m, s$ do not depend on wavelength, they can be written in front of the integrals instead of within them. Then the remaining integrals are identical to those representing the standard cone excitation values in equations 3.1 through 3.3, hence

$$L'_x = \rho_l \cdot L_x \tag{3.7}$$

$$M'_x = \rho_m \cdot M_x \tag{3.8}$$

$$S'_x = \rho_s \cdot S_x, \tag{3.9}$$

or, in matrix notation,

$$\mathbf{e}'_\mathbf{x} = \mathbf{D} \cdot \mathbf{e}_\mathbf{x} \tag{3.10}$$

where

$$\mathbf{e}'_\mathbf{x} := \begin{pmatrix} L'_x \\ M'_x \\ S'_x \end{pmatrix}, \mathbf{e}_\mathbf{x} := \begin{pmatrix} L_x \\ M_x \\ S_x \end{pmatrix}, \text{and } \mathbf{D} := \begin{pmatrix} \rho_l & 0 & 0 \\ 0 & \rho_m & 0 \\ 0 & 0 & \rho_s \end{pmatrix}. \tag{3.11}$$

---

[1] A very readable and thorough treatment of Helmholtz's and Hering's respective theories is given by Gelb (1929).

Since the sensitivity changes are thought to depend on the temporal or spatial context in which the target stimulus is viewed, the von Kries coefficients ρ_i should be understood as functions $\rho_i(C)$ of the context C. Accordingly, the diagonal matrix $\mathbf{D}$ as well as the von Kries-transformed cone excitation vector $\mathbf{e}'_\mathbf{x}$ are also functions of C. To underscore this, we write equation 3.10 in the form

$$\mathbf{e}'_\mathbf{x}(C) = \mathbf{D}(C) \cdot \mathbf{e}_\mathbf{x}. \tag{3.12}$$

The above may be understood as the generic form of the von Kries model. More specific submodels are obtained by stating how the von Kries coefficents depend on the characteristics of the context, i.e. by stating the functions $\rho_i(C)$ explicitly. Many different, more or less specific proposals have been made in the literature, but common to most formulations is the convention of normalising the functions $\rho_i(C)$ to unity for a standard viewing context S. This procedure is appropriate since we are only interested in how the sensitivities change with changes in the context, and convenient, since it implies that in the standard viewing context, the corrected cone excitation values are identical to the standard cone excitation values.

Even without making any explicit assumptions regarding the functions $\rho_i(C)$, though, it is quite simple to test whether the generic von Kries model describes the influence of context on colour appearance adequately. As follows immediately from equation 3.10, the corrected cone excitation vector of a stimulus viewed in a context A is a linear transform of its corrected cone excitation vector when viewed in another context B, i.e.

$$\mathbf{e}'_\mathbf{x}(A) = \mathbf{D}(B, A) \cdot \mathbf{e}'_\mathbf{x}(B) \tag{3.13}$$

with

$$\mathbf{D}(B, A) := \mathbf{D}(B) \cdot \mathbf{D^{-1}}(A) = \begin{pmatrix} \frac{\rho_l(B)}{\rho_l(A)} & 0 & 0 \\ 0 & \frac{\rho_m(B)}{\rho_m(A)} & 0 \\ 0 & 0 & \frac{\rho_s(B)}{\rho_s(A)} \end{pmatrix}. \tag{3.14}$$

A simple corollary of this is that a set of targets which have identical von Kries corrected cone chromaticities in terms of one context, will still have mutually identical von Kries corrected cone chromaticities in terms of any other context. To see this, remember that all cone excitation vectors of a given chromaticity can be written as

$$c \cdot \begin{pmatrix} L \\ M \\ S \end{pmatrix} \tag{3.15}$$

where the vector is fixed and c is a positive nonzero number. Multiplying a vector of this form with the transformation matrix $\mathbf{D}(B, A)$ yields

$$c \cdot \begin{pmatrix} \frac{\rho_l(B)}{\rho_l(A)} \cdot L \\ \frac{\rho_m(B)}{\rho_m(A)} \cdot M \\ \frac{\rho_s(B)}{\rho_s(A)} \cdot S \end{pmatrix}. \tag{3.16}$$

Again, the vector is fixed and $c > 0$, hence all vectors of this form have the same chromaticity.

Now, since we know that in the standard viewing condition of a single target stimulus presented in a context of complete darkness all stimuli with a given chromaticity differ only in perceived brightness and are equal in perceived saturation and hue, the von Kries model predicts that they should continue to be equal in saturation and hue for any other viewing context. This prediction has been tested and found to be at odds with the empirical findings in a number of studies. Jameson and Hurvich (1959), for instance, presented monochromatic stimuli embedded in a coloured surround to their observers and asked them to report the perceived saturation and hue of the target patch in numerical terms. In general, the results changed markedly when the luminance of the target patch was changed. Since these changes occurred with a constant chromaticity of the target patch and a constant surround, this result cannot be accounted for by the von Kries model.

Violations of this prediction of the von Kries model have also been reported in a number of other studies. Werner and Walraven (1982) asked subjects to adjust the chromaticity of a target patch embedded in a coloured surround until it appeared achromatic, and found that the chromaticities of the settings changed quite substantially with the luminance level of the target patch. This result is clearly at odds with the predictions of the von Kries model, especially when one considers that it is well documented that the chromaticity of the subjectively achromatic point is independent of the luminance when the stimulus is viewed in a context of complete darkness (Walraven & Werner, 1991). A similar result has been reported by Chichilnisky and Wandell (1996), and also studies of simultaneous contrast using unique hue settings are known to yield results which cannot be accounted for by the von Kries model (e.g. Walraven, 1976; Shevell, 1978).

A further possibility to test the validity of the von Kries hypothesis is to rely on so-called asymmetric colour matching. The observer views two target stimuli a and b embedded in different surrounds A and B and is asked to adjust the tristimulus values of one of the target stimuli until both appear equal in colour. If settings are made for several different target stimuli and a fixed pair of surrounds, the von Kries hypothesis predicts that the cone excitation values of the settings for the target b should be a linear function of those for the target a. This prediction is clearly violated in the experiments of Smith and Pokorny (1996) and Miyahara et al. (2001).

All of these studies investigated the effect of uniform coloured surrounds on the perceived colour of a target patch, i.e. simultaneous colour contrast in the narrower sense, so with respect to this situation at least, we may safely conclude that the von Kries model is inadequate. It should be noted, though, that the von Kries model was not originally proposed to account for simultaneous contrast. Instead, von Kries (1905) intended it to account for the influence of ‘chromatic adaptation’ or ‘successive contrast’, i.e. the influence of pre-exposure to a coloured stimulus on the perceived colour of a stimulus presented at a later time at the same location of the visual field. Much empirical data has been collected in this setting also. The motivation behind many early studies of this phenomenon was however not to test the validity of the von Kries model, but to determine the spectral sensitivities of the cones: It can be shown that if the von Kries model applies, then data from asymmetric matching experiments can be used to determine the matrix with which the colour matching curves must be multiplied in order to obtain the cone fundamentals (see Wyszecki & Stiles, 1982, p. 431). However, this undertaking turned out to be quite unsuccesful. In order to account for the empirical findings, it was even suggested that there could be more than three types of cones (MacAdam, 1956). However, in view of modern evidence this appears unreasonable, and the failure to establish the cone fundamentals by this method can also be attributed to invalidity of the von Kries hypothesis. Indeed, later non-linear models of chromatic adaptation have been proposed, which of course are incompatible with the von Kries model (MacAdam, 1961).

In summary, it is quite clear that the von Kries model fails to account for many psychophysical findings, and this failure can be understood as the starting point for the development of other more refined models, which we shall consider later (sections 3.2 and 3.3). Nevertheless, the von Kries model continues to play a central role in the colour vision literature. There are several reasons for this. Obviously, the failure of the von Kries model to account for psychophysical data does not necessarily imply that sensitivity changes in the cones do not play a role in simultaneous or successive contrast. Even if von Kries-type sensitivity regulation were at work, one could obtain more complicated psychophysical data if additional mechanisms also influence the percept. Indeed, the models of simultaneous contrast we shall consider later all incorporate von Kries adaptation as a contributing submechanism. Furthermore, under some experimental conditions, the von Kries model actually yields a reasonably good description of the data (Speigle & Brainard, 1999), so it may be valid for some situations.

Furthermore, based on purely computational deliberations it can be argued that the kind of transformation of receptor inputs suggested by the von Kries hypothesis can, under a number of plausible assumptions, actually be quite good for achieving colour constancy (Ives, 1912; West & Brill, 1982; Finlayson, Drew, & Funt, 1994). In this context, the von Kries model can be conceived of as an abstract model of colour transformations, and it is not necessary to assume that the transformation is implemented as cone sensitivity changes.

3.2 The two-process model

The failure of the von Kries model suggests that sensitivity regulation of the cone photoreceptors cannot yield a complete account of the influence of context on perceived colour. In order to address this problem, Jameson and Hurvich (1959) proposed a more elaborate model which, in addition to the multiplicative transformation of cone signals posited by the von Kries hypothesis, also incorporates an additive transformation of neural signals at the opponent level. According to this model, the cone excitation vector $\mathbf{e_x}$ of the target patch is first subjected to a von Kries transformation $\mathbf{D}(C)$ yielding the von Kries-corrected cone excitation vector $\mathbf{e}'_\mathbf{x}(C) = \mathbf{D}(C) \cdot \mathbf{e_x}$. This signal is then subjected to opponent recoding. Since the opponent mechanisms are assumed to be linear combinations of the cone outputs, the triplet of resulting opponent signals $\mathbf{o}'_\mathbf{x}(C)$ is obtained by multiplying the von Kries-corrected cone excitation vector with a 3-by-3 matrix $\mathbf{M}$, i.e. $\mathbf{o}'_\mathbf{x}(C) = \mathbf{M} \cdot \mathbf{e}'_\mathbf{x}(C)$. After the opponent recoding, the context C contributes a constant opponent signal represented by the vector $\mathbf{i_o}(C)$. Accordingly, perceived colour is determined by the opponent signal represented by the vector

$$\mathbf{o}''_\mathbf{x}(C) := \mathbf{o}'_\mathbf{x}(C) + \mathbf{i_o}(C), \tag{3.17}$$

or, writing it out in full,

$$\mathbf{o}''_\mathbf{x}(C) = \mathbf{M} \cdot \mathbf{D}(C) \cdot \mathbf{e_x} + \mathbf{i_o}(C). \tag{3.18}$$

The model owes its name to the fact that the context C is assumed to excert its influence on two levels of neural processing, namely at the cone level through the von Kries transform $\mathbf{D}(C)$ and at the opponent level through the additive signal $\mathbf{i_o}(C)$. In terms of opponent colours theory, the additive term $\mathbf{i_o}(C)$ can be understood as a shift in the neutral point of the opponent mechanisms. Furthermore, since the excitations of the opponent mechanisms are supposed to yield a straightforward index of perceived colour, the additive component can be said to represent a fixed amount of colour induced into the target patch by the context.

In purely formal terms it is equally possible, though, to formulate the model without reference to opponent mechanisms. According to the model, two target stimuli x and y viewed in the contexts C and D, respectively, should appear equal in colour whenever $\mathbf{o}''_\mathbf{x}(C) = \mathbf{o}''_\mathbf{y}(D)$. Obviously, since the matrix $\mathbf{M}$ is invertible, this will be the case whenever $\mathbf{M}^{-1}\mathbf{o}''_\mathbf{x}(C) = \mathbf{M}^{-1}\mathbf{o}''_\mathbf{y}(D)$, hence saying that perceived colour depends on the expression $\mathbf{o}''_\mathbf{x}(C)$ is formally equivalent to saying that it depends on the expression $\mathbf{M}^{-1}\mathbf{o}''_\mathbf{x}(C)$. Thus, we may equally well say that perceived colour is represented by the expression

$$\mathbf{e}''_\mathbf{x}(C) := \mathbf{M^{-1}}\mathbf{o}''_\mathbf{x}(C) = \mathbf{D}(C) \cdot \mathbf{e_x} + \mathbf{M^{-1}} \cdot \mathbf{i_o}(C), \tag{3.19}$$

or, defining $\mathbf{i_e}(C) := \mathbf{M^{-1}} \cdot \mathbf{i_o}(C)$,

$$\mathbf{e''_x}(C) = \mathbf{D}(C) \cdot \mathbf{e_x} + \mathbf{i_e}(C). \tag{3.20}$$

Remembering that $\mathbf{D}(C) \cdot \mathbf{e_x}$ is just the von Kries-transformed cone excitation vector $\mathbf{e'_x}$ we have

$$\mathbf{e''_x}(C) = \mathbf{e'_x} + \mathbf{i_e}(C). \tag{3.21}$$

In formal terms, then, the only difference between the von Kries model and the two-process model is that the latter allows for the additive term $\mathbf{i_e}(C)$. In general, it is assumed that the additive term is inversely related to the von Kries-corrected cone excitation vector of the context stimulus (Jameson & Huvich, 1972), i.e. if $\mathbf{e_c}$ is the cone excitation vector of the surround, then $\mathbf{i_e}(C) = -k \cdot \mathbf{e'_c}$ where k is a positive constant. Accordingly, the additive term is often referred to as subtractive. As ballpark figures, we may note that Jameson and Hurvich (1961) give estimates of k for different spatial arrangements of the target and the context stimulus, ranging from 0.14 for the case where the target patch is spatially separated from the inducing context stimulus to 0.56 for the case where the target stimulus and the context stimulus share a common side.

One advantage of the two-process model over the von Kries model is that, due to the inclusion of the subtractive term, it can account for the perception of pitch black. As we have already noted, the absence of light per se does not evoke the perception of a good black, but instead of intrinsic grey. Only by adding a brighter temporal or spatial context, pitch black can be perceived. Now, stimuli which appear achromatic when viewed in a standard dark context have cone excitation vectors of the form $c \cdot (1,1,1)^t$ where c is a non-negative number representing the intensity of the stimulus and the perceived brightness of the stimulus decreases monotonically with decreasing values of c. Extrapolating from this observation, one would expect a quantitative model to represent an achromatic stimulus which appears even blacker than the absence of light by a triplet of negative numbers. There is no way, however, that one can subject the cone excitation vector $(L, M, S)^t$ of a real light to a von Kries transform and obtain a vector $(\rho_l \cdot L, \rho_m \cdot M, \rho_s \cdot S)^t$ of von Kries transformed cone excitation values with negative components since L, M, S and the gain factors ρ_i are all positive. Adding a triplet of negative numbers, though, obviously does the trick, and this is allowed for by the additive term in the two-process model.

3.3 The contrast-coding model

When referring to the kind of simultaneous contrast which occurs when a target stimulus is embedded in a uniform coloured surround, the effect of the additive term in the two-process theory can be described by saying that a fraction of the von Kries-corrected cone excitation vector of the surround is subtracted from the von Kries-corrected cone excitation vector of the target: Since $\mathbf{i_e}(C) = -k \cdot \mathbf{e'_c}$, the model Eqn. 3.21 can be written as

$$\mathbf{e''_x}(C) = \mathbf{e'_x} + \mathbf{i_e}(C) = \mathbf{e'_x} - k\mathbf{e'_c}. \tag{3.22}$$

Exactly how large this fraction k might be is left open by the two-process theory. In this section, we shall discuss models of simultaneous contrast which give a definite and theoretically very interesting answer to this question. According to these models, the *entire* surround vector is subtracted, as opposed to just a part of it, i.e. $k = 1$. Inserting this into equation 3.22, we obtain

$$\mathbf{e''_x}(C) = \mathbf{e'_x} - \mathbf{e'_c} = D(C) \cdot \mathbf{e_x} - D(C) \cdot \mathbf{e_c} = D(C) \cdot (\mathbf{e_x} - \mathbf{e_c}), \tag{3.23}$$

which makes it clear that $\mathbf{e''_x}(C)$, which is intended to be an index of perceived colour, is a function, represented by the von Kries transform $D(C)$, of the *difference* between the standard cone excitation vector of the target and that of the surround. Clearly, this kind of model suggests that the cone excitation vector of the surround is an equally fundamental determinant of perceived colour as the cone excitation vector of the target stimulus itself. This might appear counterintuitive, but there are in fact many lines of empirical evidence supporting this hypothesis, which I shall refer to as the **contrast-coding model**. As discussed by Whittle (2003), the general idea behind this model has been put forward a number of times and can actually be said to be quite old. In modern work on simultaneous contrast, though, the idea has perhaps been put forward most forcefully in the respective contributions of Whittle and Walraven (e.g. Whittle & Challands, 1969; Whittle, 1992, 1994b, 2003; Walraven, 1976, 1979; Werner & Walraven, 1982). We shall start the discussion of the empirical evidence for contrast-coding by considering Walraven's (1976) classic experiment.

3.3.1 Walraven's experiment

In this experiment, the observers made unique yellow settings for target stimuli embedded either in a red surround or a surround of complete darkness. Both targets and surrounds where composed of monochromatic stimuli which evoke zero or close to zero S cone excitation. Hence the stimuli can be completely characterised in terms of their L and M cone excitation values and be represented in a two-dimensional L vs M cone excitation space. When a surround of complete darkness was used, the unique yellow settings at different luminance levels of the target T fell on a line through the origin of L vs M cone excitation space given by the equation $L_T = \alpha \cdot M_T$. This is exactly what one would expect based on the widespread assumption that a constant 'colour', which varies only in brightness, is coded by a given ratio of cone excitation values, or equivalently, a fixed chromaticity. In this case, the chromaticity is given by the ratio α.

When the target was viewed in a red surround, the unique settings for different luminance levels of the target did however not result in a constant ratio of L and M cone excitation anymore. Instead, Walraven found that the ratio was constant if the cone excitation values of the surround S *were first subtracted from those of the surround*, i.e. the ratio $\beta = (L_T - L_S)/(M_T - M_S)$ was constant for stimuli which appeared unique yellow in a given surround. Although the values of β where different for different luminance levels of the red surround, this general relation was found to be valid for all surrounds. It is also interesting to note that it is also valid for the dark surround since in that case $L_S = M_S = 0$, hence $L_T = \beta \cdot M_T$, as was observed.

Clearly, this finding suggests that the difference between the cone excitation vector of the target and that of the surround plays a more fundamental role in determining perceived colour than the cone excitation vector of the target itself. The different values of β for different surrounds can be accounted for by assuming that this difference vector is subjected to a surround-dependent von Kries transformation before the ratio is computed. If we assume that the perceived colour, in this case unique yellow, is determined by the ratio

$$\gamma := \frac{\rho_L(S) \cdot (L_T - L_S)}{\rho_M(S) \cdot (M_T - M_S)}, \tag{3.24}$$

for any surround, then the different values of

$$\beta(S) = \frac{(L_T - L_S)}{(M_T - M_S)} = \frac{\rho_M(S)}{\rho_L(S)} \cdot \gamma \tag{3.25}$$

for different surrounds can be accounted for by assuming that they evoke different ratios of the von Kries coefficients $\rho_M(S)$ and $\rho_L(S)$, which is quite reasonable since the latter are generally assumed to depend on the surround. Generalising from the findings of Walraven, where S-cone excitation did not play a role, to the case of triplets of cone excitation values, this leads to the

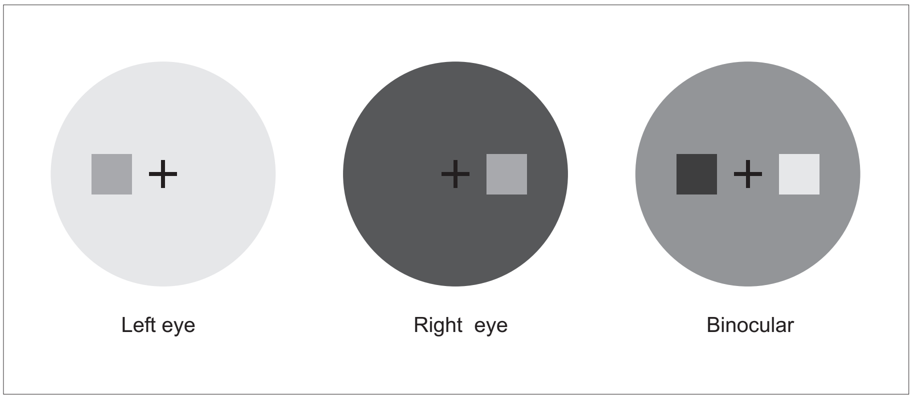

Figure 3.1: The experimental technique of haploscopically superimposed displays (after Whittle, 1994b). The left configuration is presented to the left eye, the one in the middle to the right eye. The right panel gives an impression of what the observer then perceives. The square targets in the left and middle panels are physically identical, yet appear somewhat different due to simultaneous contrast. This effect is greatly enhanced in the binocularly fused image. The reader may try to verify this by free-fusing the left and middle panels. A coloured version of this is shown in Color Plate VII on page 140.

hypothesis that perceived colour is determined by the expression

$$\mathbf{w_t}(S) := \mathbf{D}(S) \cdot (\mathbf{e_t} - \mathbf{e_s}) \tag{3.26}$$

where $\mathbf{e_t}, \mathbf{e_s}$ are the cone excitation vectors of the target and the surround, respectively, and $\mathbf{D}(S)$ is a surround-dependent diagonal matrix of von Kries coefficients.

3.3.2 Whittle's experiment

Whereas Walraven (1976) obtained evidence for the contrast-coding model with the method of unique hue settings, and later also with the technique of achromatic settings (Werner & Walraven, 1982), Whittle and Challands (1969) had earlier drawn very similar conclusions using a quite different experimental technique of asymmetric colour matching. In this experimental paradigm, which was first described by Hering (1890) as an excellent means for demonstrating simultaneous colour contrast, a mirror system is used to present two different images to the left and the right eye of the observer. As illustrated in Figure 3.1, each eye sees a target patch (square) embedded in a larger surround (disk) with a fixation cross in the middle. In the left eye, the target square is left to the centre of the surround, whereas in the right eye, it is on the other side. When the observer fuses the two images binocularly, he perceives a single surround with a single fixation cross and two target squares, as illustrated on the right side of the figure. Usually, the perceived colour of the fused surround is then some average of the colour of the two surrounds when they are viewed 'as is', i.e. without fusing them. This phenomenon is commonly referred to as binocular summation. The phenomenon of immediate interest to us, though, is the colour appearance of the target patches. In Figure 3.1 the two target patches in the left and middle panels have the same luminance, yet appear somewhat different since they are embedded in different surrounds. When the two stimuli are binocularly fused, this effect of simultaneous brightness contrast is greatly enhanced, as suggested in the right panel. The reader may try to verify this for himself by fusing the left and middle images in the figure. This effect cannot only be observed with achromatic stimuli as in Figure 3.1, but also with coloured stimuli, as shown in Colour Plate VII on page 140.

Using this kind of stimulus presentation, commonly referred to as a **haploscopically superimposed displays** (HSD), Whittle and Challands (1969) measured the size of the simultaneous brightness contrast effect for achromatic stimuli using asymmetric brightness matching. In their experiments, a target patch F of fixed luminance $\mathcal{L}_F$ was presented in a surround of complete darkness to the left eye, and the observer was asked to adjust the luminance $\mathcal{L}_T$ of a second target patch T, which was presented to the right eye and embedded in a surround C of luminance $\mathcal{L}_C$, until it appeared equally bright as the fixed target. For a given luminance level $\mathcal{L}_F$ of the fixed patch, matches were made for several different luminance levels of the surround C.

The results obtained in this study clearly support the idea of contrast coding. Any combination of target and surround luminances in the right eye which were matched to a given fixed stimulus F in the left eye yielded approximately the same value of the expression

$$b_F := \frac{\mathcal{L}_T - \mathcal{L}_C}{\mathcal{L}_C + \mathcal{L}_0}, \tag{3.27}$$

where $\mathcal{L}_0$ is a constant. If the perceived brightness of the fixed patch, which is presented to the left eye, is not influenced by the stimulus presented to the right eye, any combination of target and surround luminances which yields the same value of b_F should have the same perceived brightness, and accordingly, b_F can be interpreted as an index of perceived brightness. Defining $\rho(C) := 1/(\mathcal{L}_C + \mathcal{L}_0)$ we have

$$b_F := \rho(C) \cdot (\mathcal{L}_T - \mathcal{L}_C), \tag{3.28}$$

which makes it clear that, according to this logic, perceived brightness depends on the difference between the luminance of the target and the surround scaled by a surround-dependent gain factor $\rho(C)$. Note the analogy to equation 3.26, which can also be written as

$$\begin{aligned} w_T^L(C) &= \rho_L(C) \cdot (L_T - L_C) && (3.29)\\ w_T^M(C) &= \rho_M(C) \cdot (M_T - M_C) && (3.30)\\ w_T^S(C) &= \rho_S(C) \cdot (S_T - S_C) && (3.31) \end{aligned}$$

for cone excitation values L, M and S.

An immediate and important corrollary of the contrast-coding model in equation 3.28 is that an increment should never be matched to a decrement. If a target patch T embedded in a surround C is a luminance increment, i.e. $\mathcal{L}_T > \mathcal{L}_C$, then it should never be matched to a target patch T' embedded in a surround C' which is a luminance decrement, i.e $\mathcal{L}_{T'} < \mathcal{L}_{C'}$. Indeed, in accordance with the model such increment-decrement matches never occur in experiments using haploscopically superimposed displays (Whittle, 1994b). This implies that one can produce arbitrarily large simultaneous contrast effects just by using sufficiently different surrounds. By way of example, if a target T with a luminance of 11 cd/m^2 is embedded in a surround C of luminance 10 cd/m^2, then whatever the luminance of the other surround C' is, say 1100 cd/m^2, the subject will have to choose a luminance greater that this for the target patch T' in order to establish a match. Obviously, then, at the point where the target patches appear equal in brightness, their luminances differ by a factor of more than 100. The same logic can be applied to the contrast coding model for cone-excitation vectors. In this case, the model in equation 3.26 predicts that a target which is an $X-$cone decrement (where X is L, M or S) can never be matched to a target which is an $X-$cone decrement. The dramatic effects of the surround on the perceived colour of the target patch suggested by this scheme remind one of the well-known quote attributed to the french painter Eugène Delacroix (1798-1863): “Give me the mud of the streets and I will turn it into the lucious flesh of a woman ... if you allow me to surround it as I please.”

3.3.3 Other evidence for contrast-coding

Experimental evidence for contrast-coding has been obtained by using a number of different experimental techniques, each with their own characteristic advantages and disadvantages. The results of some kinds of experiments are in clear agreement with the contrast-coding hypothesis, whereas the results of other kinds of experiments are less unanimous. Significantly, experiments in which the haploscopically superimposed displays of Whittle and Challands (1969) have been used yield results which are generally in very good agreement with the contrast-coding hypothesis (Chichilnisky & Wandell, 1995; Shepherd, 1997, 1999; Niederée & Mausfeld, 1997; Richter, 2002). Studies similar to that of Walraven (1976), though, in which the stimuli are presented in ‘direct view’, i.e. without the ‘tricks’ of the haploscopically superimposed displays, have produced less unequivocal evidence for the contrast-coding hypothesis. Shortly after the publication of Walraven’s study (1976), Shevell (1978) presented evidence from a similar study which were at odds with the contrast-coding hypothesis. Walraven (1979) contested Shevell’s criticism, though, and a rather vigorous controversy ensued (Shevell, 1980; Drum, 1981; Adelson, 1981; Werner & Walraven, 1982; Davies et al., 1983; Nerger, Piantanida, & Larimer, 1993).

In the comparatively simple experimental situation studied by Walraven (1976) and Shevell (1978) where the stimulus is seen in direct view, then, the contrast-coding hypothesis becomes problematic. On the one hand, these problems are purely empirical; Shevell’s findings simply do not agree with the predictions of the model. On the other hand, though, it can be argued that, in this situation, the contrast-coding model simply cannot work based on purely logical deliberations. The problem is that the contrast-coding model makes absurd predictions for the case of zero contrast between target and surround, as noted both by Mausfeld and Niederée (1993) and Whittle (1994a).

3.3.4 The zero-contrast problem

To appreciate the nature of this problem, consider two target stimuli a and b, embedded in the uniform coloured surrounds A and B, respectively. Assume further that each target has the same cone excitation vector as the surround in which it is embedded, i.e. $\mathbf{e_a} = \mathbf{e_A}$ and $\mathbf{e_b} = \mathbf{e_B}$. Obviously, in this case there is no contrast, and accordingly no visible border between target and surround. Clearly, then, the target patch evokes the same colour impression as the surround, since the target and the surround constitute a single uniform surface (see Colour Plate VIII on page 140). We may say that at zero contrast, the target ‘inherits’ the perceived colour of the surround. Now, if the two surrounds A and B are of different colours, then the target patches a and b must also appear to have different colours. This is, however, at odds with the contrast-coding model, which predicts that the two target patches should appear to have the same colour: According to the model (equation 3.26), the expressions

$$\mathbf{w_a}(A) \quad = \quad \mathbf{D}(A) \cdot (\mathbf{e_a} - \mathbf{e_A}) \tag{3.32}$$

and

$$\mathbf{w_b}(B) \quad = \quad \mathbf{D}(B) \cdot (\mathbf{e_b} - \mathbf{e_B}) \tag{3.33}$$

should be the respective indices of perceived colour for the target a in the surround A and for the target b in the surround B. At zero target-surround contrast they must both be equal to the zero vector since then $\mathbf{e_a} - \mathbf{e_A} = \mathbf{0}$ and $\mathbf{e_b} - \mathbf{e_B} = \mathbf{0}$. Whether the von Kries transforms $\mathbf{D}(A)$ and $\mathbf{D}(B)$ differ, does not matter, since they are both linear mappings and thus map the zero vector onto the zero vector.

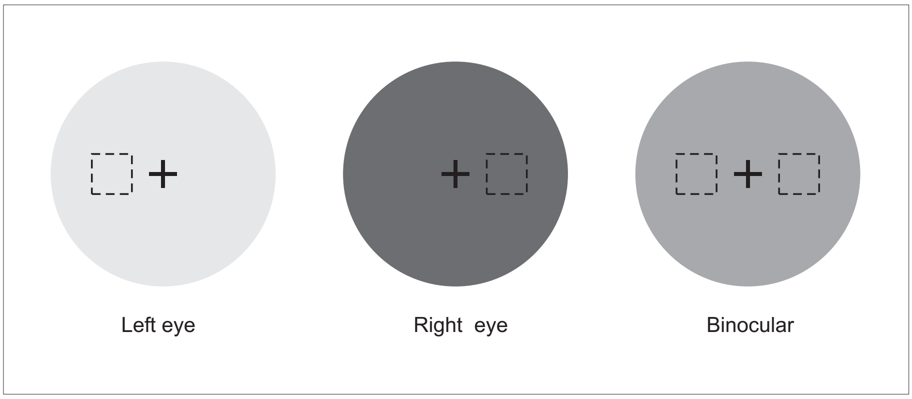

Figure 3.2: The zero-contrast problem does not apply to haploscopically superimposed displays. If zero contrast targets are presented in both eyes (left and middle), both the single-eye stimuli and the binocularly fused percept (right) consist of a large uniform surface. At the locations corresponding to the targets, the observer obviously has identical colour impressions. The dashed contours are drawn here only to indicate the locations of the zero-contrast targets, which are of course indistinguishable from the surrounds, both in the single-eye views and the binocularly fused view. In this illustration, achromatic stimuli are shown, but obviously the same reasoning applies to coloured stimuli.

The problem of zero contrast does not only apply to the specific model discussed here, where the target-surround difference vector is subject to subsequent von Kries scaling. Clearly, by substituting the von Kries transform in the contrast coding model by any function which maps the null vector onto itself, a variant of the model is obtained which is also subject to the zero contrast problem. Several such variants have been discussed in the literature (e.g. Mausfeld & Niederée, 1993; Chichilnisky & Wandell, 1995; Niederée & Mausfeld, 1997; Richter, 2002). Quite generally, I shall use the term contrast-coding model for any model of simultaneous contrast which is subject to the zero-contrast problem.

3.3.5 The scope of the zero-contrast problem

The zero-contrast problem makes it clear that contrast-coding models make an absurd prediction when applied to a fairly simple situation, and therefore cannot realistically serve as a complete model of everyday colour perception. However, under the special viewing conditions employed in many studies of simultaneous contrast, this prediction is actually not absurd at all. Rather, it agrees with the facts for quite trivial reasons. In the case of asymmetric colour matching experiments using haploscopically superimposed displays, for instance, a zero-contrast target presented in one of the surrounds will always appear equal to a zero-contrast target presented in the other surrounds. To see this, consider that when both targets are at zero contrast, the stimulus presented to each eye is just a large uniform surround, and the surrounds from both eyes are binoculary fused into a single large surround (see Figure 3.2). At the two locations of the visual field where the targets would be visible if they were distinguishable from the surround, one obviously has identical colour impressions.

The haploscopically superimposed display is not the only experimental technique in which the otherwise absurd zero-contrast prediction of the contrast coding models is trivially true. Wuerger (1996) and later also Beer and MacLeod (2001) performed asymmetric colour matching experiments using an experimental ‘trick’ which makes the two physically different surrounds appear equal in colour. The ‘trick’ consists in blurring the edges of the surrounds, as shown in Figure

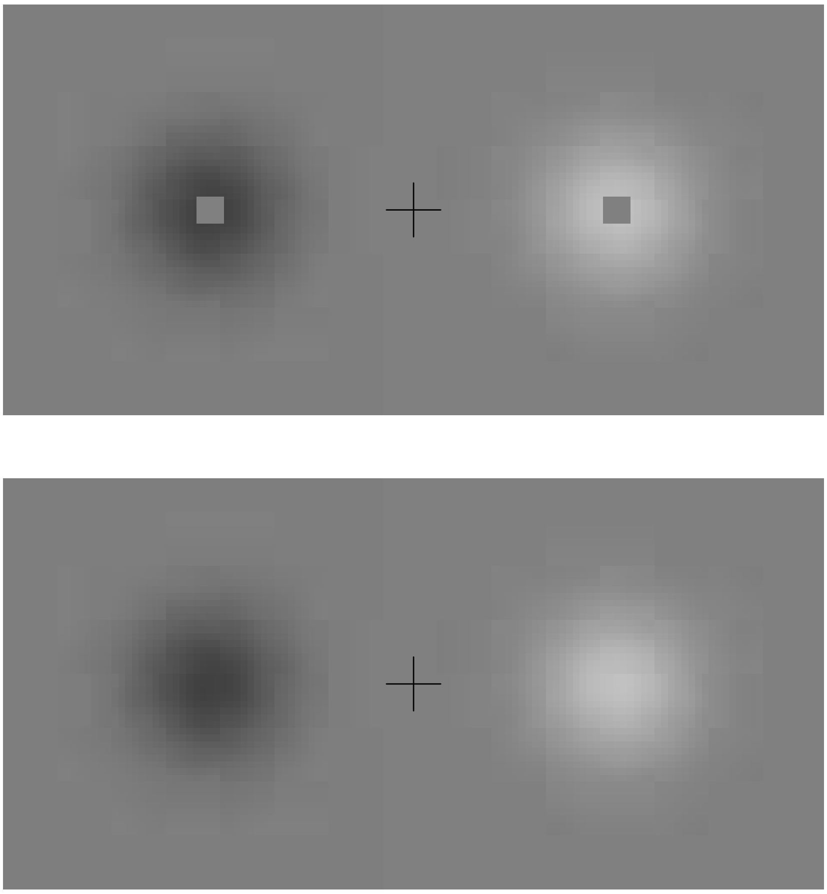

Figure 3.3: Illustration of the blurring-technique used by Wurger (1966), and Beer and MacLeod (2001). Top: The targets (squares) are presented in differently coloured (here differently bright) surrounds, the outer borders of which are shaded gradually towards a common mid-grey background. When one maintains strict fixation of the central cross for some time, the outer borders of the local surrounds become invisible, and one has the impression of two targets embedded in a common, mid-gray surround. Thus even though the immediate surrounds of the targets are different, they appear equal. This general effect can also be appreciated in the lower panel, where the targets are omitted. In the actual experiments, the target stimuli are flashed onto the surrounds for a brief period of time. This appears to be necessary because under strict fixation also the targets have a tendency to become invisible, making a judgement impossible. Flashing the targets makes them visible also under strict fixation. The targets in the top display are physically identical, but appear slightly different due to simultaneous contrast. The reader may try to verify that the effect is enhanced when the outer borders of the immediate surrounds become invisible.

3.3. In the top panel of this figure two target stimuli (squares) are embedded in a light and a dark surround with blurred outer edges. If one fixates the cross in the middle for a little while, the two physically different surrounds and the common mid-grey background will all appear to be equally bright. The target squares then appear to be embedded in the same mid-grey surround, although their immediate surrounds are physically quite different. Using coloured stimuli instead of the achromatic ones in Figure 3.3, Wuerger (1996) exploited this effect and asked her subjects to adjust the colour of one of the target patches until it appeared equal to the other. Her results, as well as those of Beer and MacLeod (2001) were comparable to those obtained with the haploscopically superimposed displays: In order to establish a match, the observers essentially equated the differences between the cone excitation vector of the target patch and its immediate surround, and the observers obviously will do so even if the contrast between the target and the immediate surround is zero: If both targets have zero contrast, i.e. are indistinguishable from the immediate surround one has the situation shown in the bottom panel of Figure 3.3. Whenever the central cross is fixated such that the immediate surrounds appear equal, one just perceives a single uniform mid-grey background, hence the zero-contrast prediction is trivially true.

3.3.6 Contrast-coding and everyday colour perception

As we have seen, experiments using haploscopically superimposed displays or the edge-blurring technique of Wuerger (1996) generally yield very strong simultaneous contrast effects consistent with contrast-coding. Experiments using less artificial viewing conditions, though, for instance asymmetric matching experiments where the two centre-surround stimuli are seen side-by-side in direct view, generally yield smaller simultaneous contrast effects, which are difficult to account for by contrast-coding (Whittle, 1994a). Now, since the latter kind of experiment obviously is closer to the conditions of everyday colour perception, one might wonder what the theoretical significance of the contrast-coding observed in the haploscopically superimposed displays might be. According to Whittle (1994b, 1994a, 2003), contrast-coding is a feature of retinal mechanisms. The fact that clear evidence for contrast-coding is seldom obtained under more ordinary viewing conditions can be explained by assuming that the effects of contrast-coding on the retinal level are overshadowed or counteracted by other mechanisms at higher levels of the visual system (Shapley & Reid, 1985). Assuming this to be the case, it can be argued that the advantage of the HSD-technique over more ordinary viewing conditions is that it eliminates or balances out the contribution of higher mechanisms and therefore makes it possible to study the behavior of retinal mechanisms in a more direct way. Theoretically, these higher mechanisms are thought to perform some kind of ‘integration’ of the difference signals provided by the retinal mechanisms (Whittle & Challands, 1969; Whittle, 1994a; Arend, 1973). To appreciate what is meant by integration consider the following example: Two target stimuli a and b are embedded in the surrounds A and B as suggested in Figure 3.4 (top). For simplicity, we assume all surfaces to be achromatic so that we only have to consider the single dimension of luminance instead of three-dimensional cone excitation vectors. Assume that the two centre-surround stimuli have the luminance profile shown in the lower part of the figure. Suppose now, that the retina extracts the difference signals $\Delta aA := a - A, \Delta bB := b - B$ and $\Delta AB := A - B$ from the borders in the stimulus. Only these difference signals, as opposed to the absolute signals a, b, A and B are conveyed to higher levels of the visual system. How could the higher centres of the brain use this difference signals to decide whether the targets a and b have the same luminance? Clearly, they could compute the luminance difference Δab of the targets a and b by taking all the difference signals into account, i.e.

$$\begin{aligned} \Delta ab &= \Delta aA + \Delta AB - \Delta bB \\ &= (a - A) + (A - B) - (b - B) \\ &= a - b, \end{aligned} \tag{3.34}$$

as should be evident from the figure, and the target patches a and b should of course have the same luminance whenever Δab is zero. Inserting this into the above equation yields the solution $a = b$, which means that the target stimuli should be regarded as equal whenever their absolute luminances are equal. If the brain performs this kind of computation, then, there would be no simultaneous contrast effect even if the retina only codes difference signals; physically identical

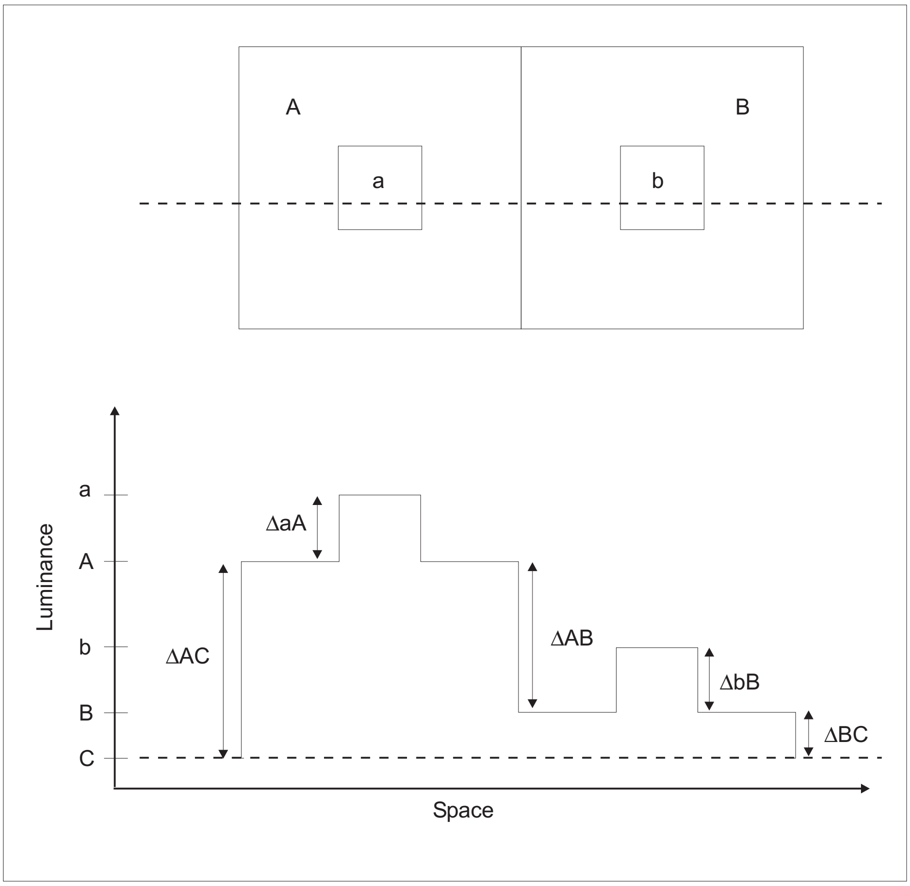

Figure 3.4: The contrast-coding-plus-integration perspective applied to two juxtaposed centre-surround stimuli (top). The lower panel shows the luminance profile along a horizontal path (upper dashed line). See text for further explanations.

targets should appear equal even if they are embedded in different surrounds. If, however, the brain for some reason fails to take the difference signal ΔAB into account, one would expect quite a strong simultaneous contrast effect. Neglecting the term ΔAB in eqn. 3.34 yields

$$\Delta ab = \Delta aA - \Delta bB. \tag{3.35}$$

As above, the targets should be regarded as equal whenever Δab is zero, but now this should be the case whenever $\Delta aA = \Delta bB$, or equivalently, $a - b = A - B$. This means that the target patches should appear equal whenever their luminances differ by an amount just as large as the luminance differences of the surrounds; if the surrounds are very different, we would expect a large effect.

We can now understand what is meant by integration; it simply means taking more difference signals into account than those at the borders of the target stimulus itself. In the above simple example, integration means taking not just the local difference signals ΔaA and ΔbB, but also the non-local difference signal ΔAB into account. This notion of integration can be used to explain the differences between the results obtained with the HSD and the edge-blurring technique on the one hand and those obtained with side-by-side stimuli presented in direct view on the other hand. In the HSD, as well as the blurring technique, the difference between the two surrounds A and B is not perceived, i.e. the brain assumes ΔAB to be zero although it actually is not. Thus, integration breaks down, and we obtain large simultaneous contrast effects: In essence, target-surround differences are equated. In side-by-side stimuli, on the other hand, smaller simultaneous contrast effects are generally found (Whittle, 1994a). This can be explained by assuming that the difference ΔAB is taken into account, which is plausible, since in side-by-side displays, it is readily perceived by the observer. We must assume, however, that integration is only partial. If it were complete, we would expect no simultaneous contrast effect at all, which is seldom, if ever, the case, even in side-by-side displays.

This perspective, according to which the retina codes difference signals and the brain integrates them (Whittle & Challands, 1969; Arend, 1973) has several interesting theoretical implications. For one thing, the special viewing conditions of the HSD and the blurring technique, which otherwise might be regarded as unnatural and artificial, can be regarded as cleverly devised experimental tools for studying the workings of retinal mechanisms, since they eliminate, or balance out, the influence of the integration performed at higher levels of the visual system. Furthermore, it implies that in order to account for everyday colour perception, both a model of contrast-coding and a model of integration is needed; contrast-coding alone will not suffice. This also means that the zero-contrast problem cannot be used as an argument against contrast-coding models; the fact that two zero-contrast targets in different surrounds do not appear equal, as predicted by contrast-coding alone, can simply be attributed to a mechanism of integration.

3.3.7 Evidence for contrast-coding under ordinary viewing conditions

As already mentioned, a substantial bulk of the evidence for contrast-coding stems from studies using the HSD or the blurring-technique. Based on the contrast-coding-plus-integration model of colour perception delineated above, one would expect to obtain data consistent with a simple contrast-coding model under these special viewing conditions. Under other, more ordinary viewing conditions, though, one would *not* expect to obtain results consistent with a simple contrast-coding model, since in this case one would also expect mechanisms of integration to influence the results. The findings of Walraven (1976), however, seem to be at odds with this expectation. In his study, the observers made unique yellow settings for a target stimulus embedded in a uniformly coloured surround. Since both the target and the surround were presented in direct view, there is no obvious reason why integration should fail to play a role. Nevertheless, his results could be well described by a pure contrast-coding model. This finding poses a number of challenging questions. Although the HSD experiments and Walraven's experiment are both generally cited as evidence for contrast-coding, they suggest rather different theoretical scenarios. Whereas the HSD experiments suggest that contrast-coding appears under special viewing conditions where integration is eliminated, Walraven's experiment suggests that contrast-coding occurs under ordinary viewing conditions. In the former case, the contrast-coding model can be understood as describing just one stage of the total process which determines perceived colour under ordinary viewing conditions, in the latter case it appears that the contrast-coding model describes the total process. Furthermore, whereas the zero-contrast prediction of the contrast-coding model is trivially true in the HSD experiments, it is patently absurd in the direct-view situation studied by Walraven. Hitherto, we have only formulated the zero-contrast prediction for asymmetric colour matching experiments, stating that two zero-contrast targets should appear equal even if they are embedded in different surround. We did not state what colour the zero-contrast targets should appear, only that they should be equal. Based on opponent colours theory, it is, however, possible to derive a more specific expectation which can be directly related to Walraven's experimental situation, namely that zero-contrast targets should always appear achromatic even if the surround is coloured: According to the contrast-coding model, the perceived colour of the target is determined by the expression

$$\begin{aligned} \mathbf{w_t}(C) &:= \mathbf{D}(C) \cdot (\mathbf{e_t} - \mathbf{e_c}) \\ &= \mathbf{e}'_\mathbf{t} - \mathbf{e}'_\mathbf{c} \end{aligned} \tag{3.36}$$

where $\mathbf{D}(C)$ is a von Kries transform, and $\mathbf{e_t}, \mathbf{e_c}$ are the cone excitation vectors of the target and the surround, respectively, and $\mathbf{e'_t}, \mathbf{e'_c}$ the corresponding von Kries-transformed cone excitation vectors. This model can equally well be expressed in terms of the variables of opponent colours theory instead of the cone excitation vectors of trichromatic theory. To this end, the von Kries-transformed cone excitation vectors must simply be multiplied with the matrix $\mathbf{M}$. Thus, in the opponent formulation, perceived colour should be determined by the expression

$$\begin{aligned} \mathbf{o''_x}(C) &:= \mathbf{M} \cdot \mathbf{e'_t}(C) - \mathbf{M} \cdot \mathbf{e'_c}(C) && (3.37)\\ &= \mathbf{M} \cdot \mathbf{w_t}(C). && (3.38) \end{aligned}$$

Now, at zero contrast, $\mathbf{w_t}(C)$ is the zero vector, and multiplying it by any matrix also yields the zero vector. Hence, at zero contrast, the three opponent mechanisms, represented by the vector $\mathbf{o''_x}(C)$ must all be zero, i.e. in equilibrium. Since opponent colours theory assumes that when the red-green and the blue-yellow mechanisms are both zero, the colour impressions should be achromatic, we would expect that any zero-contrast target should appear achromatic. Clearly this is an absurd prediction: A zero-contrast target in a coloured surround obviously appears in the same colour as the surround. If the surround is red, for instance, the target also appears red. Accordingly, the conclusion that Walraven's (1976) model cannot be taken literally appears inevitable.

Beyond this conceptual problem, there is also direct empirical evidence challenging the validity of Walraven's contrast-coding model. In an experiment very similar to that of Walraven, Shevell (1978) found that the contrast-coding model must be modified in order to account for the data. According to Shevell's model, perceived colour is determined by the expression

$$\mathbf{s_t}(C) := \mathbf{D}(C) \cdot (\mathbf{e_t} - k \cdot \mathbf{e_c}), \qquad (3.39)$$

where k is a positive constant less than unity. Note that this equation is just the two-process model, of which the respective models of Walraven and Shevell are special cases; According to Walraven's model, $k = 1$ whereas according to that of Shevell, k is less than that. For $k = 1$ we may say that the cone excitation vector of the surround is completely subtracted out, or 'discounted', for k less than unity we may say that it is only partially discounted, hence Walraven's model is called a 'full discounting model', and Shevell's is called a 'partial discounting model'. Eqn. 3.39 can be rewritten as

$$\mathbf{s_t}(C) := \mathbf{D}(C) \cdot (\mathbf{e_t} - \mathbf{e_c} + (1 - k) \cdot \mathbf{e_c}), \qquad (3.40)$$

hence the only difference between the models of Walraven and Shevell is that the latter includes the term $(1 - k) \cdot \mathbf{e_c}$, whereas the former does not. This term, which is generally referred to as the 'additive effect', is generally found to be fairly small, but it is of profound theoretical significance; One may say that Shevell's model is immune to the zero-contrast problem by virtue of including this term. Conversely, Walraven's model is subject to the zero-contrast problem because it does not include the additive term.

Whether it is, in fact, necessary to include this additive term in order to account for the empirical facts has been the subject of a long-standing and vigorous controversy (Walraven, 1979; Shevell, 1980; Drum, 1981; Adelson, 1981; Werner & Walraven, 1982; Davies et al., 1983; Nerger et al., 1993). Curiously, it proved very difficult to settle the issue. Some studies suggest that the additive term is necessary, others suggest that a pure contrast-coding model suffices. One difficulty stems from the fact that the additive effect found in the studies of Shevell (1978) and others is – in absolute terms – rather small, i.e. k is close to unity. Accordingly, the two models make very similar predictions, and very precise measurements are necessary in order to distinguish between

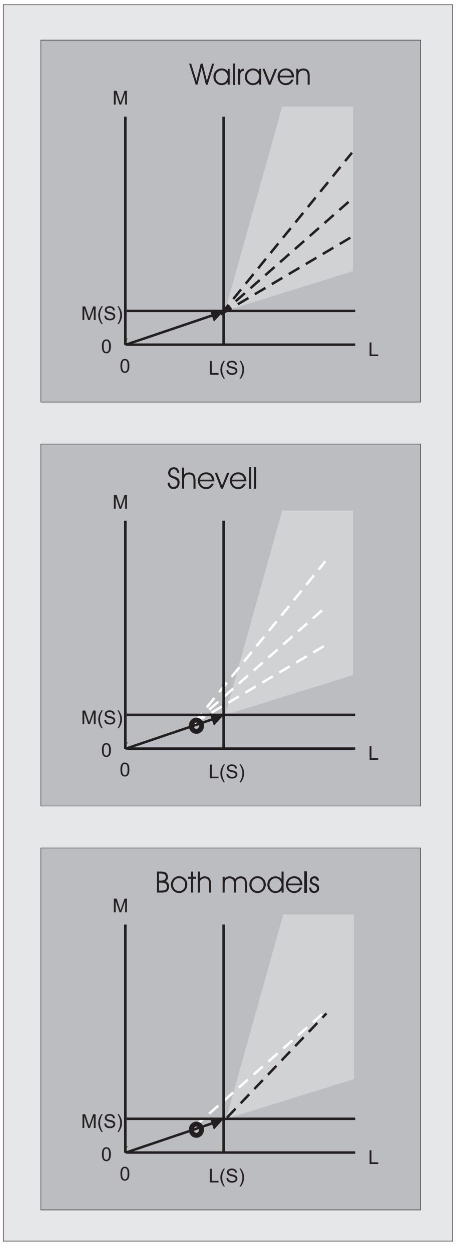

Figure 3.5: Comparison of the predictions of Walraven's (1976) and Shevell's (1978) models, here in a two-dimensional (dichromatic) cone excitation space. The horizontal axis represents the L-cone excitation of the stimuli, the vertical axis M-cone excitation. The arrow from the origin to the point $(L(S), M(S))$ represents the cone excitation vector of the surround. According to Walraven's model (top), unique yellow settings at different luminance levels should be located on a straight line through that point. The slope of this line depends on the von Kries coefficients, as suggested by the different alternatives shown here (dashed lines). For a given surround, though, the von Kries coefficients are fixed, and the data should be described by one of the lines. According to Shevell's model, the data should be described by a straight line through the point $k \cdot (L(S), M(S))$. Estimates of k suggest that it is close to unity (in the order of 0.8), hence the point shown in the middle and lower panels should be grossly realistic, and one may expect that the data should fall on one of the white dashed lines, which goes through this point. Since the slope of the prediction line is unspecified by both models, very similar predictions can be made based on the two models. This is illustrated in the lower panel. The most diagnostic cases are those where the white dashed line deviates the most from the black white line, i.e. stimuli which are rather similar to the surround. This makes it difficult to distinguish between the models based on empirical data. In the original experiments of Walraven and Shevell, this problem is excerberated by the technical restriction that only some of the potentially diagnostic measurements could be made, roughly those within the more brightly shaded region.

the alternative models. Furthermore, in Walraven's original experiment, as well as in the majority of the investigations addressing the issue, the target patch was a light beam superimposed on a background light constituting the surround. Accordingly, the target patch was always a physical increment with respect to the surround. When only such incremental target stimuli are investigated, the predictions of the two models differ most markedly for target stimuli which have low contrast with respect to the surround, as schematically illustrated in Figure 3.5.

Unfortunately, it seems that exactly under these conditions, which would be most diagnostic, subjects experience problems when trying to make their settings. In Walraven's (1979, p. 1062) words, "it is extremely difficult to judge the color of the test field (subjects typically may be at loss to report what they see). There are indications of a vague, hueless impression (an achromatic interval?) at the transition from seeing green [...] to seeing red." Also, in a study by Drum (1981), where target stimuli of very low contrast (close to threshold detectability against the surround) were used, three of the five observers made settings indicative of an additive effect, but no data where obtained from the two remaining subjects because they were "*unable to do the experiment*. In repeated attempts to measure the red-green equilibrium point, they reported that the hue of the test flash was always either greenish or so weak as to be uninterpretable" (p. 960, his italics). That the test stimulus should always look greenish is exactly what one would expect in the absence of an additive effect, so taken together, the results are equivocal: Three subjects made settings in accordance with the predictions of Shevell's model, whereas the other two behaved as one would expect from Walraven's model.

Possible problems with the perceptual task One possible explanation for the equivocal results of these experiments is that subjects are not actually able to perform the task they are asked to carry out. In most of the experiments, the observers are asked to adjust the colour of the target stimulus until it appears unique yellow, i.e. neither reddish nor greenish, or, alternatively, unique grey. According to Hering's (1920) opponent colours theory (see section 2.3), no colour should appear both reddish and greenish, nor both bluish and yellowish, so this should be a well-defined criterion. However, some investigators have reported that target stimuli embedded in a uniform surround may – especially at low target-surround contrast – evoke a dual colour impression reminiscent of perceptual transparency (Helmholtz, 1911; Masin & Idone, 1981; Brenner & Cornelissen, 1991; Brown & MacLeod, 1997; Mausfeld, 1998; Niederée, 1998). By definition, perceptual transparency involves the perception of two colours at the same location of the visual field, namely the colour of the transparent medium and the colour of the object perceived behind and 'through' it (Fuchs, 1923). In Colour Plate IX on page 141 this is illustrated in the kind of four-region stimulus typically used in studies of perceptual transparency (Metelli, 1970; D'Zmura et al., 1997; Faul, 1997; Faul & Ekroll, 2002). If such dual colour impressions are also evoked by simple centre-surround stimuli, one may expect that the task of adjusting the colour of the target patch until it appears neither reddish nor greenish might be difficult, since the target may appear as a reddish patch behind a greenish transparent layer, or vice versa, such that the total colour impression in the region of the target is both reddish and greenish. That a target stimulus embedded in a uniform coloured surround may indeed appear both reddish and greenish at the same time has been reported by several researchers (Helmholtz, 1911; Brenner & Cornelissen, 1991; Mausfeld, 1998). Such phenomena suggest themselves as a possible explanation for the equivocal results obtained in the studies contributing to the Walraven-Shevell controversy. If the target patch appears both reddish and greenish, it will obviously not be an easy task for the subjects to make a setting which appears neither reddish nor greenish, and it can be expected that they may resort to some kind of less well-defined compromise.

In the next chapter, I shall present the results of a series of experiments which strongly suggest that such dual colour impressions are indeed evoked by simple centre-surround stimuli. Using a novel and 'indirect' technique for determining the neutral point for targets embedded in a coloured surround we find that it coincides with the chromaticity of the surround, in accordance with the seemingly absurd zero-contrast prediction of Walraven's contrast-coding model. Although this result is actually predicted by the contrast-coding model, it must be regarded as paradoxical under the assumption that perceived colour can be adequately represented by a triplet of numbers, as all of the quantitative models hitherto discussed, including Walraven's, implicitly assume. If this assumption is given up, though, it is possible to develop a rational explanation of the present findings in terms of perceptual transparency. Different aspects of the experimental results are well captured by the models of Walraven and Shevell, respectively, but the total pattern of results suggests that the common basic premise of both models, namely that the colour impressions evoked by simple centre-surround stimuli only vary along three dimensions, is untenable.

Part II

Experiments

Chapter 4

The convergence paradox

A common method for measuring simultaneous colour contrast is to let subjects adjust the chromaticity of the target stimulus until it appears achromatic, or devoid of hue (e.g. Helson & Michels, 1948; Werner & Walraven, 1982). This experimental technique rests on the commonly accepted notion that achromatic colour impressions are special. Colour Plate II on page 137 illustrates how colour impressions are thought to correspond to points in colour space: Planes of constant hue intersect in a central line which represents all the achromatic colour impressions. In a plane of constant luminance, which corresponds to a planar cross-section of this space, lines of constant hue converge on a singular achromatic point, just as the spokes of a wheel meet at its hub. Thus, if an observer is asked to adjust the chromaticity of the target such that it appears achromatic, he simply has to find this point, which should be a straighforward and well-defined task.

In my own experience as a subject in experiments of this kind, though, I have sometimes felt it difficult to make exact settings in the sense that I was certain that there remained no trace of this or that hue. Since, as discussed in the previous section, a high accuracy of the settings can be essential for distinguishing between the predictions of different models, I tried to devise of a method which would make it easier for the observer to make precise settings. One idea which was explored in pilot experiments was to use more than a single target patch embedded in the same surround. If four target patches are used, for instance, one could ‘encircle’ or ‘straddle’ the achromatic point by making settings where the four patches have just noticable traces of the four opponent hues blue, yellow, red and green. To this end, the chromaticities of the four target patches were linked to each other as four small difference vectors from a common mean chromaticity, as illustrated in Figure 4.1. The observer could adjust both the mean chromaticity of the targets as well as the common length of the difference vectors, which extended in four different directions of chromaticity space, i.e. roughly ‘redwards’, ‘greenwards’, ‘yellowwards’ and ‘bluewards’ from the mean chromaticity. By iteratively adjusting the mean chromaticity and the length of the difference vectors, the observer should be able to find a point where the four targets are nearly achromatic, but different enough to say with confidence that one is slightly too green, the other slightly too red, the third slightly too yellow and the fourth slightly too bluish to be truly achromatic. At this setting, it appears natural to assume that the achromatic point should be located somewhere between the chromaticities of the four patches, presumably in the middle.

The informal pilot experiments with this method yielded highly interesting but also very unsettling results. With the new method, estimates of the achromatic point were obtained which coincided very accurately with the chromaticity of the coloured surround. Since the targets and the surround were equiluminant, this corresponds exactly to the zero-contrast prediction of the contrast-coding model: If a target has the same luminance and chromaticity as the surround it has zero-contrast and the model predicts that it should appear achromatic. Yet, as we have already discussed, this simply is not true because in this case the target is indistinguishable from the surround

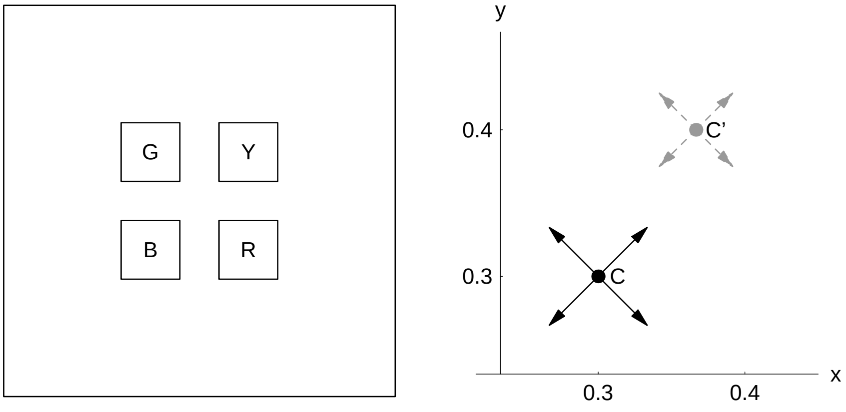

Figure 4.1: Basic logic of the pilot experiments. The observer viewed four targets embedded in a common surround, as suggested in the left panel. The chromaticities of the targets deviated in four different directions from a mean chromaticity C. The observer could adjust this mean chromaticity, as well as the common length of the offsets. The right panel shows two different possible settings, with different mean chromaticities and lengths of the difference vectors. The tips of the arrows represent the chromaticities of the targets.

and must therefore appear in the same colour as the surround, which was coloured. Thus, the new method yielded an estimate of the achromatic point which cannot possibly be obtained using the classical method of grey settings. Accordingly, when the results obtained with the new method were compared with the results obtained using classical grey settings, they turned out to be inconsistent. These inconsistencies turned out to be more than just minor differences; the total pattern of results suggested that the grey point obtained using classical achromatic settings was actually located *outside* of a circle in chromaticity space containing a full hue circle. This is clearly at odds with the common intuition that any circle in chromaticity space spanning all hues, however small its radius might be, should surround the achromatic point.

Clearly, these findings could be due to some kind of experimental artifact. If they are not, however, they must be regarded as real anomalies challenging a basic premise behind classical models of simultaneous colour contrast. In the following I report the results of a series of more formal and elaborate experiments which were designed to investigate this more closely. To anticipate, the results strongly suggest that we are dealing with a real effect which is difficult to understand from a classical perspective, but can be rationally accounted for if one assumes that simple centre-surround stimuli evoke impressions reminiscent of perceptual transparency. These experiments were performed and analysed in close collaboration with Franz Faul, Reinhard Niederée and Eike Richter, whose substantial contributions I gratefully acknowledge.

4.1 Basic logic of the experiments

Colour Plate II on page 137 illustrates schematically how the colour impressions of targets viewed in a context of complete darkness are thought to correspond to points in tristimulus space: Planes of constant hue intersect in a central axis which represents achromatic colour impressions of different brightness. In a plane of constant luminance (right panel) lines of constant hue converge on the achromatic point. The colour represented by this ‘neutral point’ should therefore have both the property of appearing achromatic and being the convergence point for lines of constant hue. We

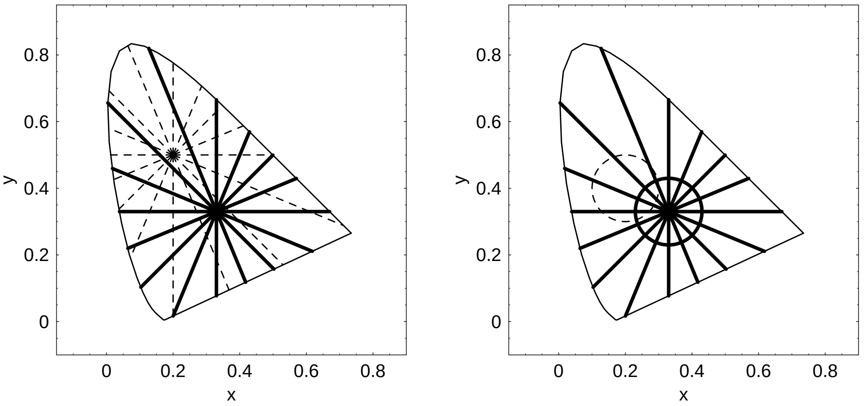

Figure 4.2: Left: Illustration of the convergence criterion. The heavy lines represents lines of constant hue, which converge on the achromatic point. If straight lines are drawn from any other point in the chromaticity diagram, they cannot all be of constant hue: Most of the dotted lines intersect many different constant hue lines. Right: Illustration of the enclosure criterion. Each of the straight lines represents a different hue. In order to intersect all of these lines, a closed path in the chromaticity diagram must enclose the achromatic point. This is the case for the heavy circle, but not for the dashed circle.

shall refer to the point at which lines of constant hue converge as the neutral point according to the **convergence criterion**. Evidently, this point should be uniquely defined: A set of straight lines extending from any other point in the chromaticity diagram can not all be of constant hue (see Figure 4.2, left). The achromatic point is also thought to possess a further property; any closed path around it in chromaticity space should contain all possible hues (**enclosure criterion**). As illustrated in the righ panel of Figure 4.2, this should only be true for the achromatic point.

These commonly accepted notions about structure in colour space and its relation to perceptual variables ultimately have their roots in an experimental paradigm that studies isolated light spots surrounded by complete darkness. In this simple case any light stimulus corresponds uniquely to a point in tristimulus space, and the neutral point in a plane of equiluminance should possess all the three properties of a) appearing achromatic, b) fulfulling the convergence criterion, and c) fulfilling the enclosure criterion. It is generally assumed, though, that the equivalence of these criteria also holds for the general case in which the target stimulus is viewed in a surround of any colour. This generalisation is a trivial corollary of classical models of simultaneous contrast. According to most models, including those discussed in the previous chapter, introducing a coloured surround simply relocates the point in tristimulus space representing a given colour impression. Conversely, the effect of the surround is just to relocate attributes of colour appearance within tristimulus space. Therefore if a given colour impression has the attributes a, b and c in the presence of a dark surround, it should also possess them in the presence of a coloured surround. Accordingly, the achromatic point should fulfill the convergence and enclosure criteria also in the presence of a coloured surround. In our experiments, we investigated whether this is in fact the case by determining the neutral point for a given surround using all three criteria and comparing the results.

4.2 General methods

The stimuli, which consisted of one or more square target patches embedded in a uniform rectanglar surround were presented on a CRT screen (resolution, 1.024 × 768 pixel; 75-Hz frame

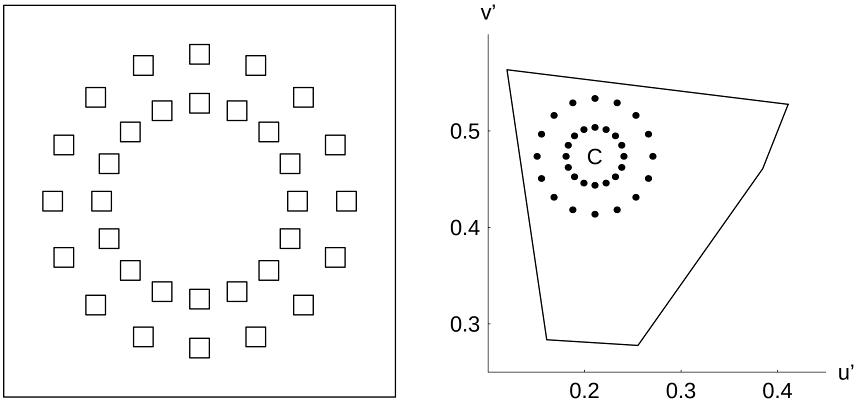

Figure 4.3: Left: Illustration of the stimulus used in the second part of experiment 1. The observers viewed two sets of 16 targets arranged in two concentric circles and embedded in a common surround (The relative sizes of the targets and the circles are correctly reproduced, but the figure is not drawn to scale, and the surround was larger). Right: The chromaticities of the targets were distributed on two corresponding circles in (u',v') chromaticity space. For clarity, the radii of the circles are here twice as large as they were in the actual experiment. The polygon represents the gamut of the monitor at the luminance used in the experiment.

rate) and viewed in a completely dark chamber. The monitor was controlled by a graphics card with a colour depth of 8 bits per red-green-blue channel. This setup was calibrated by means of a colourimeter (LMT Lichtmesstechnik, Berlin; C1210), following a standard procedure described by Brainard (1989).

4.3 Experiment 1

Stimuli and procedure In this experiment we determined the neutral point for targets presented in a coloured surround using three different experimental techniques. All the targets and the surrounds used were equiluminant at 10 cd/m^2, and data were collected for three differently coloured surrounds with the CIE 1976 (u',v') coordinates (0.2,0.49), (0.3,0.46) and (0.225,0.38), appearing approximately violet, red and yellow, respectively. In the first part of the experiment, conventional grey settings were used: The subjects were asked to adjust the chromaticity of a single square target patch (width, 0.66° visual angle) embedded in a coloured rectangular surround (width, 14.25°, height 10°) until it appeared devoid of any hue. The adjustments were made by pressing the four arrow keys of a keyboard, with each pair of arrow keys changing the chromaticity of the target along one of the two axes of chromaticity space.

In the second part of the experiment, the neutral point was instead determined according the convergence criterion. As illustrated in Figure 4.3, two sets of 16 square target patches (width, 0.33° each) were arranged on two circles ("spatial circles", radii, 1.67° and 2.5°) and displayed in the same surrounds as those used in the first part of the experiments. The chromaticities of these patches were equally spaced on two concentric circles ("chromaticity circles") in the CIE uniform chromaticity space (Wyszecki & Stiles, 1982). The radius of the inner chromaticity circle was set at a relatively small value (0.015 Euclidean distance in the (u',v') plane). The outer chromaticity circle had the same centre as the inner one, its radius being twice as large (see Figure 4.3, right). The chromaticity of each patch on the screen was set to that point of the chromaticity circles that corresponds to its location on the spatial circles. The subjects made their settings by adjusting the centre of the two chromaticity circles. Thus, the two chromaticity circles were rigidly translated in chromaticity space, keeping the relative positions of all target chromaticities unchanged. The task of the subjects was to make a setting which made each pair of target patches with the same angular position on the screen appear to have the same hue. The centre of the chromaticity circles at this setting was regarded as an estimate of the neutral point according to the convergence criterion.

In the third part of the experiment, the enclosure criterion was used. The stimulus display was identical to that used in the second part of the experiment except that the outer spatial circle of targets were omitted. In this part of the experiment, the subjects could also adjust the radius of the chromaticity circle; the task was to find a setting for the centre that enabled the perception of a full hue circle[1] while keeping the radius as small as possible. Four experienced psychophysical observers, including myself, participated in the experiment. All were colour normal as ascertained by the Ishihara Tests for Colour-Blindness (Ishihara, 1967). Every measurement was repeated 8 times.

Results Typical results from all three parts of the experiment are shown in Figure 4.4, the results of the three other observers are shown in Figure 4.5. For all three surrounds, it can be seen that the cluster of the achromatic settings is clearly different from the very similar clusters obtained using the convergence criterion and the enclosure criterion. Furthermore, the two latter clusters coincide almost perfectly with the surround chromaticity. Thus, using the convergence criterion and the enclosure criterion one obtains an estimate of the neutral point which corresponds to the chromaticity of the surround, in agreement with the problematic zero-contrast prediction of the Walraven's (1976) contrast-coding model. The achromatic settings, in contrast, yield an estimate of the neutral point which is clearly different from the chromaticity of the surround. The deviations are generally approximately in the direction of the chromaticity of equal energy white, as one would expect based on Shevell's (1978) two-process model; The deviations from the surround chromaticity may be understood as representing the 'additive effect' postulated in his model.

Taken together, our results appear paradoxical in several regards which will not have escaped the mindful reader's notice, and we shall discuss these issues thoroughly later. Before we do so, however, it is appropriate to ask whether the results are in some way due to artifact. A problematic aspect of the present experiment is the different complexity of the stimuli used in the different parts of the experiments. In the first part of the experiment only a single target patch was used, whereas in the second and third parts effectively 32 and 16 target stimuli were visible to the observer. The following experiment was designed to address this issue.

4.4 Experiment 2

Stimuli and procedure In this experiment, subjects matched the hue of only two target simuli (width, 0.66°; centre-to-centre distance, 1.38°) embedded in the same uniform surround (width, 14.25°; height, 10°). The chromaticity of one of these patches was fixed to 1 of 16 equidistant points on a circle in uv chromaticity space. The chromaticity of the second target patch could be adjusted continuously along a smaller concentric circle in chromaticity space in the search for the hue match. The center of both chromaticity circles was located at the point halfway between the background chromaticity and the achromatic point, which was determined experimentally by using the same procedure as in the first part of experiment 1. The radii of these chromaticity circles were 1.5 times and 3 times the distance between these two points, respectively (see Figure

---

[1] As a criterion for the perception of a full hue cirle, subjects were instructed to ascertain that shades of all four unique hues red, green, blue and yellow were visible.

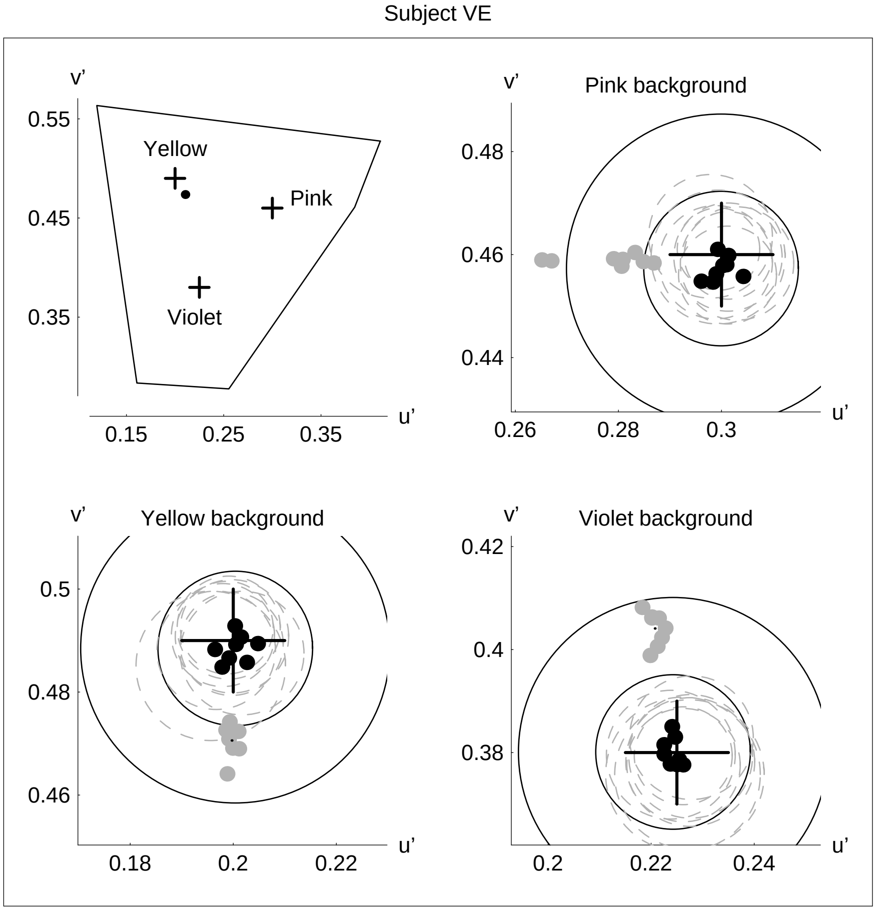

Figure 4.4: Typical results from experiment 1 (subject VE), plotted in the CIELUV (u',v') chromaticity space. Upper Left: The crosses represent the three surround chromaticities investigated. The black point represents the chromaticity of equal energy white. The surrounding polygon is the monitor gamut at the luminance used in our experiments (10 cd/m^2). The other panels represent scaled-up portions of this space. (Surround chromaticity crosses are scaled correspondingly in size.) The grey dots represent the achromatic settings made in the first part of the experiment, and the black dots represent the settings made according to the convergence criterion (second part of the experiment). Around the mean of these settings (not shown), two black solid circles are drawn with the fixed radii used. Each of the dotted circles represents one single setting according to the enclosure criterion (third part of the experiment).

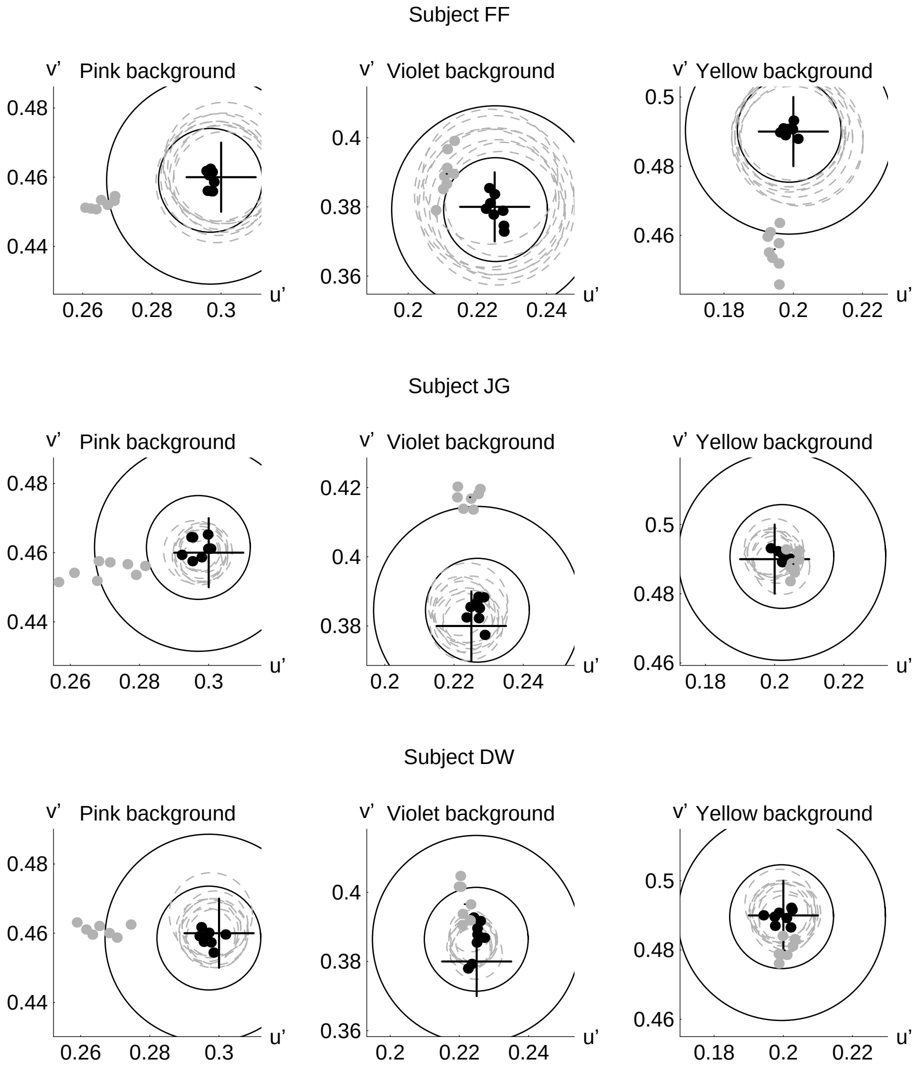

Figure 4.5: Results of experiment 1 for the subjects FF, JG and DW. Axes and symbols as in Figure 4.4.

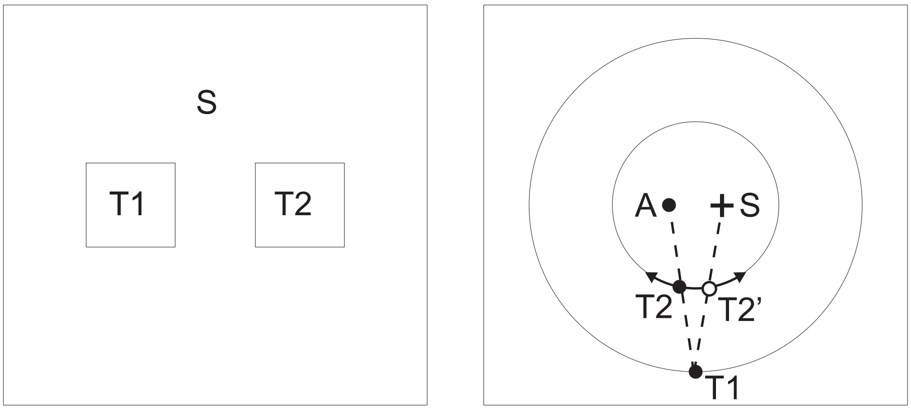

Figure 4.6: The basic logic of experiment 2. Left: The subjects viewed two targets $T1$ and $T2$ embedded in a common surround S. Right: The chromaticity of the achromatic point A was determined in a preliminary experiment. Using this information, the chromaticities of the targets $T1$ and $T2$ were sampled from two concentric circles in chromaticity space surrounding both the achromatic point A and the chromaticity of the surround S. In a given trial, the chromaticity of the target $T1$ was fixed, and the subject was asked to adjust the chromaticity of $T2$ along the inner circle until $T1$ and $T2$ appeared to have the same hue. Depending on whether lines of constant hue converge on the achromatic point or the chromaticity of the surround, one would expect the settings $T2$ and $T2'$, respectively.

4.6). Because the radii of the chromaticity circles depend on the grey settings of the subject, it often happened that the outer chromaticity circle extended beyond the monitor gamut. In these cases, the surround colours, which initially were set to the same values as in experiment 1, were made slightly less saturated and the entire procedure was repeated. The convergence criterion implies that straight lines drawn through the standard patch chromaticity and the chromaticity of the hue match intersect at the neutral point of chromaticity space. Because both the achromatic point and the surround chromaticity were located symmetrically with respect to the center of the inner chromaticity circle and well within it (see Figure 4.6), it was a priori equally possible for the lines of constant hue to converge on either of these points.

Results The results of observers GW and VE, for whom data were collected for all three surround colours are shown in Figures 4.7 and 4.8, respectively.

The results of additional subjects who only performed the experiment with one or two of the surrounds are shown in Figures 4.9 and 4.10. To estimate the neutral point for a subject and a given surround, we searched for the chromaticity N that minimised the mean angular deviation between the line from the chromaticity of the fixed target to that of the subjects settings and the corresponding lines from the fixed chromaticity to the neutral point N. As can be seen in all of the plots, the estimated neutral point lies very close to the surround chromaticity and much farther away from the mean achromatic setting. Thus, the basic results of experiment 1 can be confirmed in this more rigorous test.

4.5 Experiment 3

A central assumption underlying our use of the convergence criterion in the above experiments is that lines of constant hue are approximately straight. If, however, contrary to this assumption, the lines are strongly curved, for instance as depicted in Figure 4.11 (left), the results of experiment

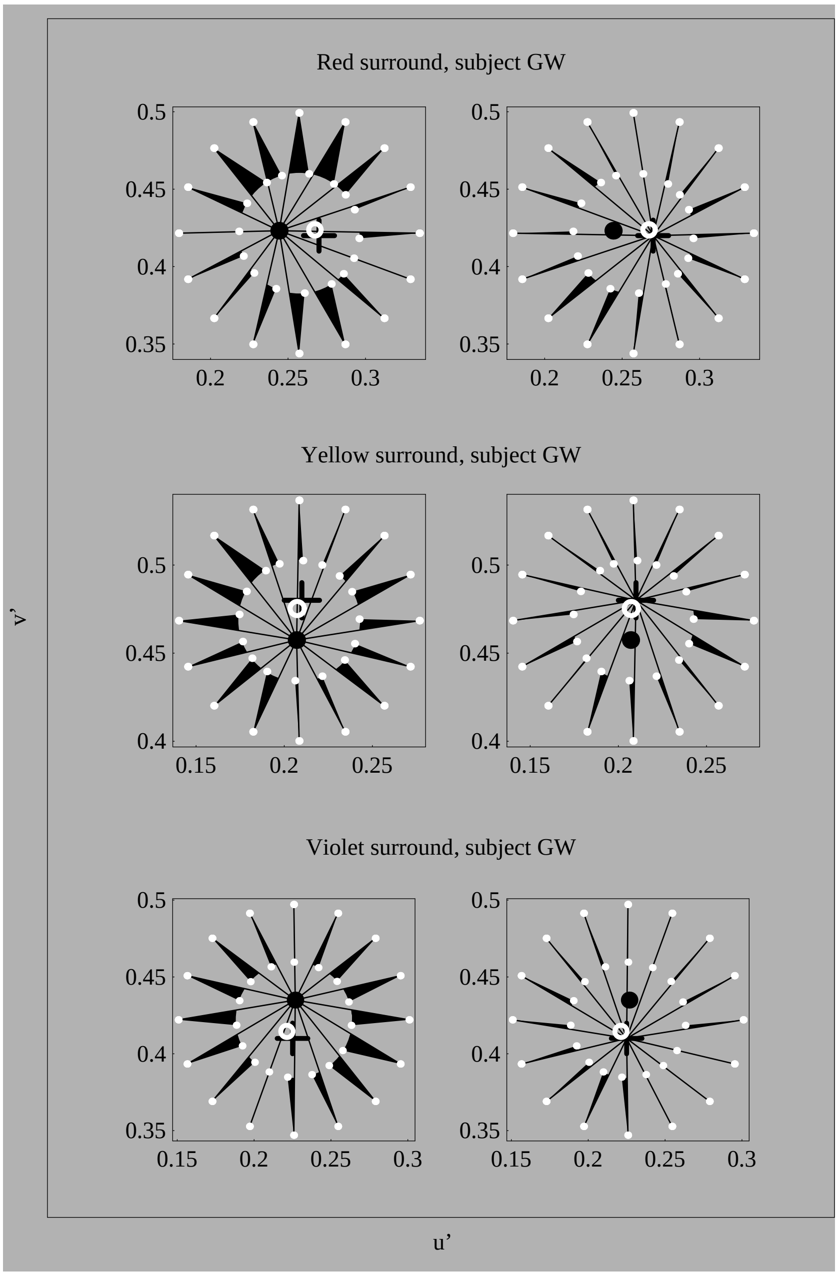

Figure 4.7: Results from experiment 2 plotted in (u’,v’) chromaticity coordinates (subject GW). Row 1 is for the red surround, row 2 is for the yellow surround and row 3 is for the blue surround. In the left column, the data are plotted together with lines predicted by the hypothesis that lines of constant hue converge on the achromatic point (black filled circle). Shorter lines are drawn from the chromaticities of the fixed patches (small white dots on the outer circle) to the corresponding mean settings (small white dots on the inner circle), and the angular deviation from the predicted lines is represented by the black wedges. The white open symbol represents the best-fitting point of convergence, and the black cross represents the surround chromaticities. In the right column, the same data are plotted, except that the lines predicted by the hypothesis that lines of constant hue converge on the surround chromaticities are given. The black wedges representing the deviations from the prediction are clearly much smaller in this case. Each data point is the mean of three repeated settings.

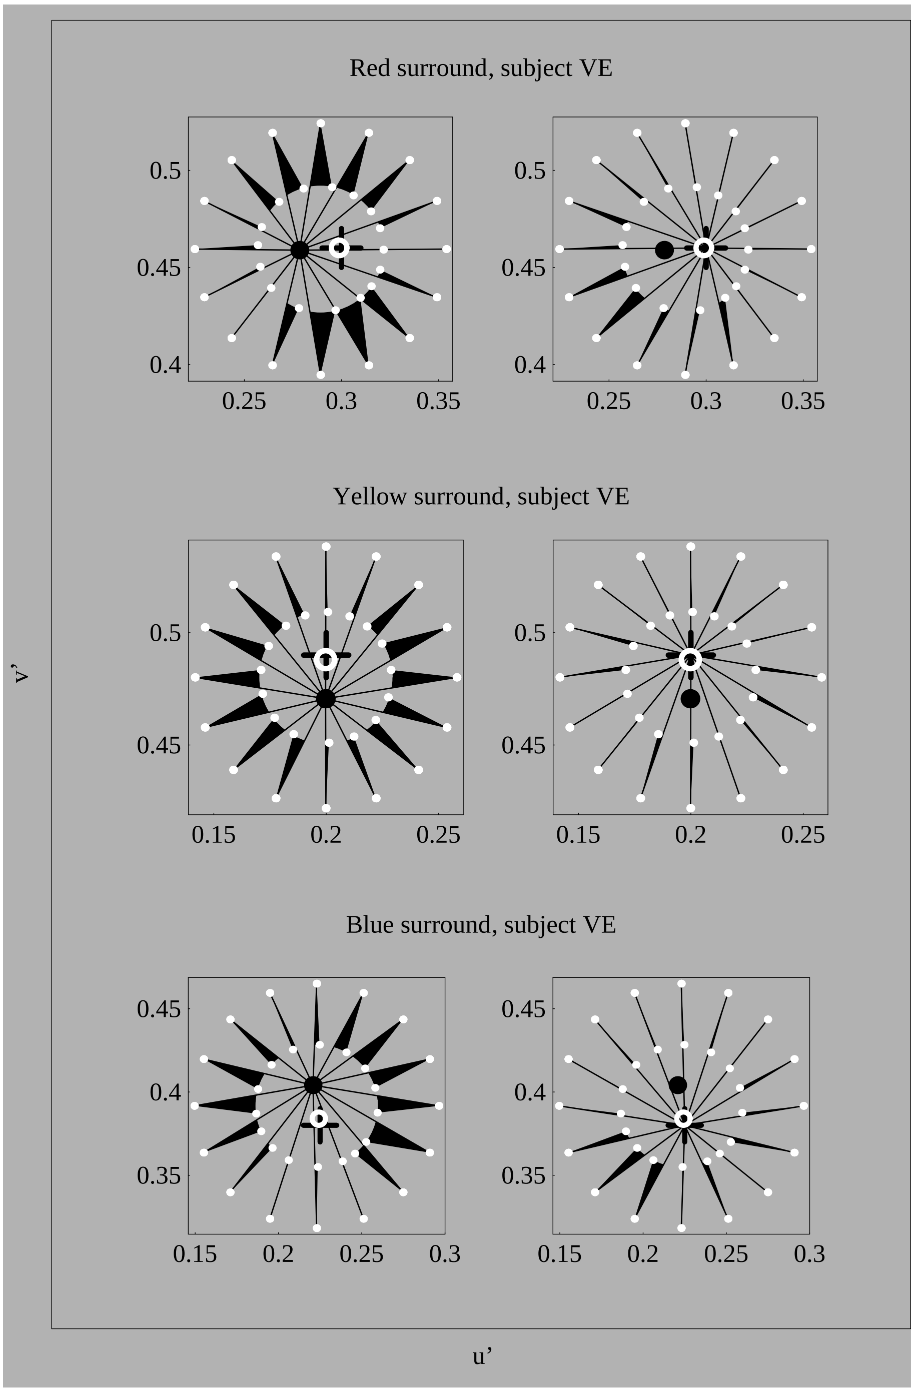

Figure 4.8: Results from experiment 2, subject VE. See caption of Figure 4.7 for explanations.

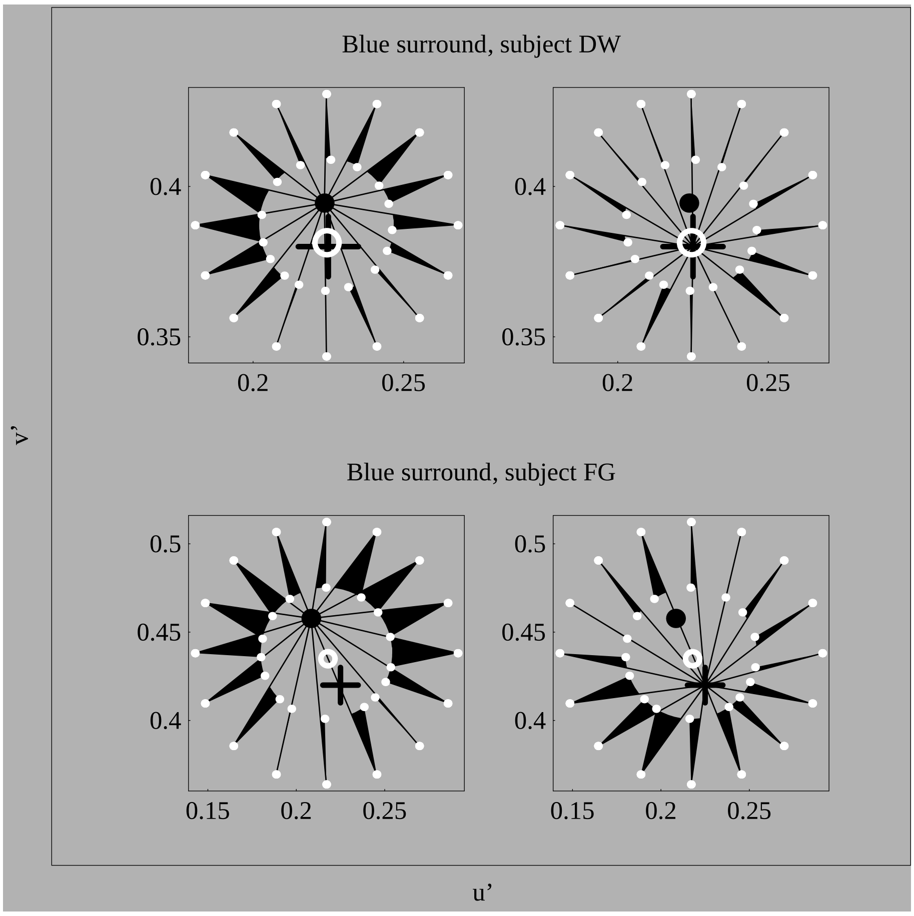

Figure 4.9: Results from experiment 2, blue surround, subjects DW and FG. See caption of Figure 4.7 for explanations.

2 may be compatible with the standard assumption that lines of constant hue converge on the achromatic point. Although only relatively small deviations from linearity have been reported (Burns et al., 1984), the possibility that large deviations may occur under the present experimental conditions cannot be ruled out *a priori*. Therefore, we conducted a third experiment in which we addressed this possibility explicitly.

Stimuli and procedure The results of experiment 2 show that if lines of constant hue were curved and thus still may converge on the achromatic point, some of them should have the strongest curvature (see Figure 4.11, left), and we focussed on these cases. We used a method with fixed stimuli instead of the adjustment procedure employed in experiment 2. The subjects viewed two target patches (width: 1.27°, centre-to-centre distance: 2.64°) in a common surround (width: 27.6°, height: 19.2°). The surrounds, which may be roughly described as yellow, pink and violet, had the uv-coordinates (0.2, 0.53), (0.3, 0.46) and (0.22,0.35), respectively. In a 4-alternative forced choice procedure, the subjects were asked to indicate whether the left (test) patch a) had the same hue as the right (standard) patch, b) had a hue different from that of the standard, c) was invisible against the surround, or d) appeared achromatic. They were instructed to use a strict criterion for a hue match.

For each surround, the chromaticities for the test patches were chosen from a triangular region

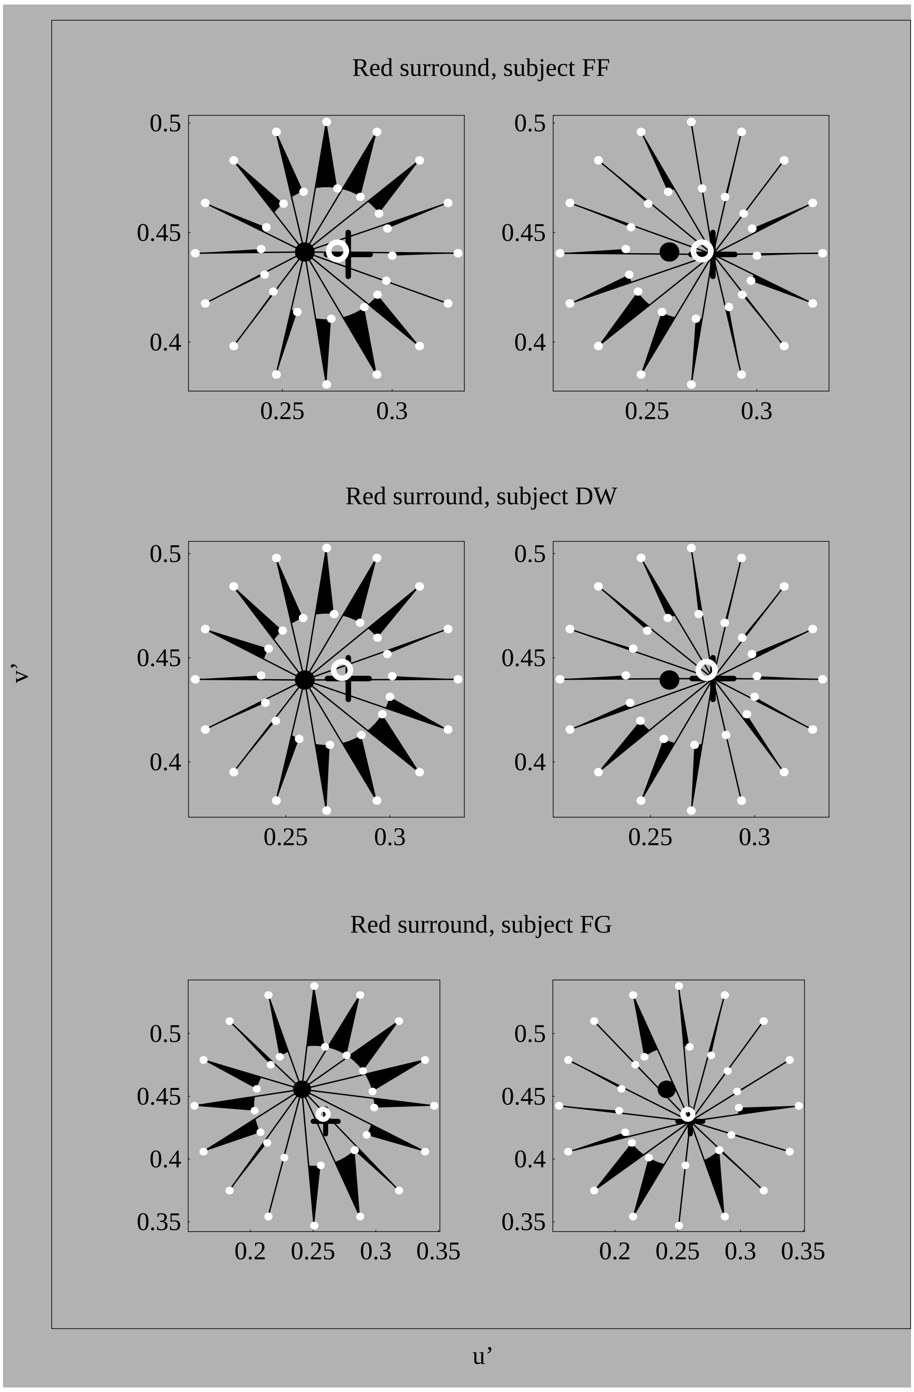

Figure 4.10: Results from experiment 2, red surround, subjects FF, DW and FG. See caption of Figure 4.7 for explanations.

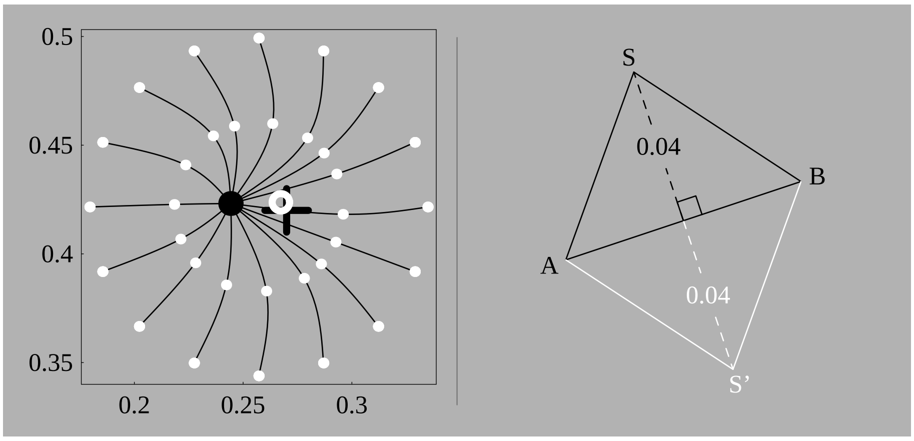

Figure 4.11: Left: A possible alternative interpretation of the data in experiment 2 (here for the red background; compare with Figure 4.7, row 1) in terms of curved lines of constant hue. Right: Illustration of the choice of target chromaticities in experiment 3 (see text for details). The black and white lines correspond to two different conditions used for each surround.

in uv-space spanned by the chromaticities A, B and S of the achromatic point, the surround, and the standard patch, respectively. All different chromaticities inside this triangle realisable at 8-bit colour resolution were used for the test patches. The achromatic point was - just as in experiment 2 - determined for each subject in a preliminary experiment. The individually determined mean A over 8 achromatic settings for a surround together with the surround chromaticity B defines a line AB in chromaticity space. For the standard patch we used a chromaticity S which (in uv chromaticity space) had a Euclidian distance of 0.04 units from this line, and was equally far away from A and B (see Figure 4.11, right panel).

Results The results for two subjects are plotted in figures 4.12 and 4.13 (the results of an additional subject were similar to those of subject MK and are not shown). It can be clearly seen that the equal hue lines for subject MK, which extend all the way down to threshold detectability of the target against the surround, are approximately straight and converge on the surround chromaticity. This equal hue line and the achromatic locus are separated by an extended region which is neither classified as achromatic nor as a hue match. The equal hue lines of observer BS look like “amputated” versions of those of subject MK. After the experiment subject BS reported that he had experienced a pronounced impression of transparency when the target was similar to the background and that he, in view of the instruction to apply a strict criterion, had not classified these cases as hue matches. Similar reports of perceptual transparency were also made by the other subjects (see discussion). The results of experiment 3 clearly rule out non-linearities of equal hue loci as a possible explanation for the dissociation of the achromatic point and the convergence point for lines of equal hues observed in experiment 1 and 2. Furthermore, since in this experiment judgements of achromaticity and hue equality were made under the exact same conditions, any artefact due to slight differences in context can also be ruled out.

4.6 Discussion

The results of all three experiments support the conclusion that there is a dissociation between the perceptual criterion of achromaticity and the structural definition of the neutral point in chro-

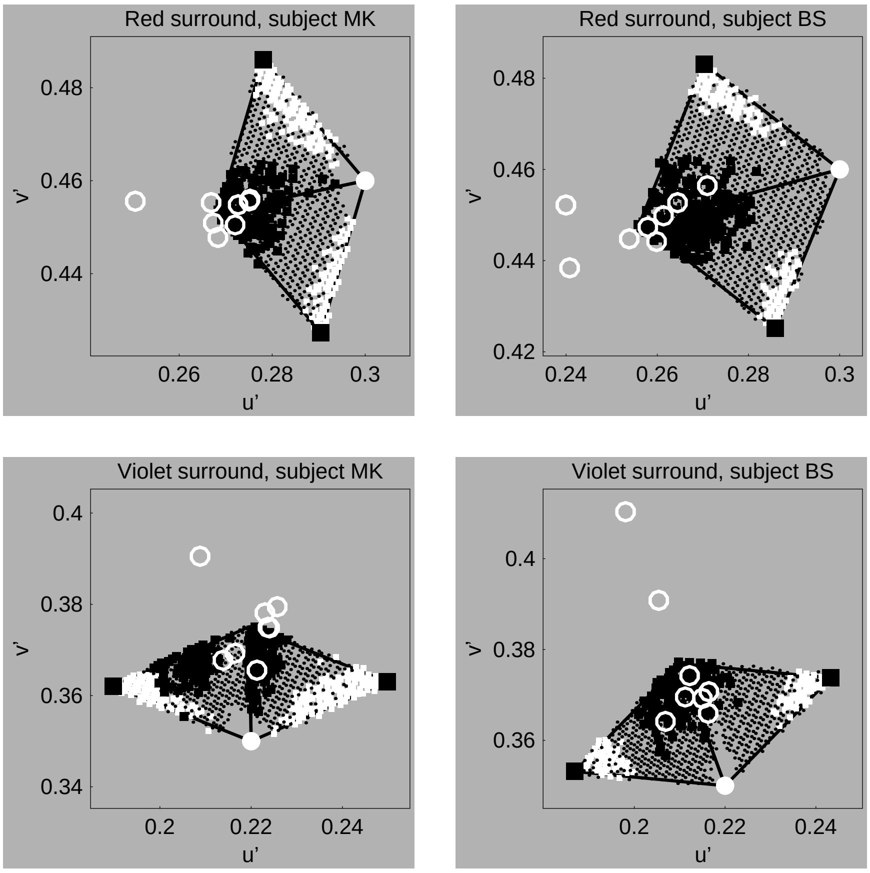

Figure 4.12: Typical results from experiment 3 in (u',v')-chromaticity coordinates, for the red and the violet surround. Each panel shows the results for one subject and one surround chromaticity. The filled white circle represents the surround chromaticity. The open white circles represent single grey settings from the preliminary experiment. For each surround, two standard target chromaticities were used, represented by the large black squares. Each point within the triangle defined by the surround chromaticity, the mean achromatic setting and the standard test patch represents one of the test chromaticities which were compared with that test patch. Points plotted as white squares represent test patches which appear equal in hue to the standard patch, those plotted as small black squares represent test patches which appear achromatic, and those plotted as small black dots represent test chromaticities that are judged neither to be achromatic nor a hue match. Chromaticity points for those test targets which appeared indistinguishable from the surround are omitted in the plot, yielding an empty region close to the white filled circle.

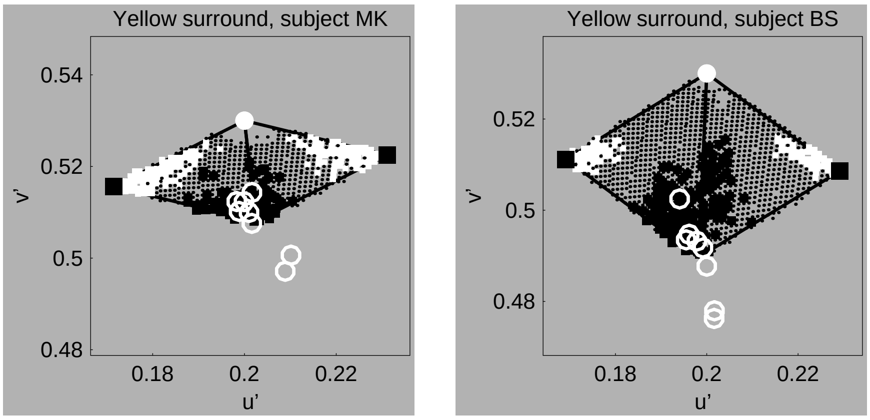

Figure 4.13: Typical results from experiment 3, for the yellow surround. See caption of Figure 4.12 for explanation.

maticity space in terms of the convergence criterion. Lines of equal hue converge on the surround chromaticity, and not on the achromatic point. Therefore, contrary to common intuitions, the chromaticity of the background should be regarded as a more natural centre of chromaticity space than the achromatic point.

Interestingly, our estimates of the neutral point obtained with achromatic settings are consistent with Shevell's (1978) two-process model, whereas our estimates of the neutral point using the convergence criterion and the enclosure criterion are consistent with Walraven's (1976) and Whittle's (1994b, 2003) contrast-coding models. Taken together, though, the present findings must appear paradoxical if one tries to understand them in terms of classical notions of colour space with three dimensions such as hue, saturation and brightness. In the following, we shall briefly consider some conceptual issues raised by our findings, as well as how they can be resolved.

To begin with, the finding that lines of constant hue converge on the chromaticity of the surround appears to be a contradiction in terms. To see this, consider first the unproblematic case of lines of constant hue converging on an achromatic colour. Any colour on a given line of constant hue can then be considered to be a mixture of two colours. All colours on a line of constant red hue, for instance, can be considered to be a mixture of a saturated red colour and white (or, depending on brightness, any achromatic colour), as illustrated in Colour Plate X on page 142. Thus, along every line of constant hue, the colour impressions which are represented contain increasing amounts of whiteness as saturation is decreased. The admixture of white does not change the perceived hue, though, since white is a neutral colour having no hue itself. If lines of constant hue converge on a chromatic colour, however, we encounter a problem: Also in this case, the colours along any line of constant hue can can be regarded as a mixture of two colours. In this case, though, both of them have a hue. If lines of constant hue converge on yellow, for instance, all lines of constant hue should contain increasing amounts of yellowness, as illustrated in Colour Plate X on page 142. By way of example, a line of constant, say, red hue should contain increasing amounts of redness, which obviously means that hue changes, and changing hue along a line of constant hue is of course a contradiction in terms. Nevertheless our subjects, who were instructed to make constant hue judgements, actually produced settings suggesting that lines of constant hue converge on the chromaticity of the coloured surround. How can this enigmatic pattern of results be accounted for?

As already mentioned above, any line in chromaticity space converging on the chromaticity of the surround consist of two colour components, one of them corresponding to the colour of the surround, for instance red, and the other to the hue of the chromaticity at the other end of the line in chromaticity space. In Colour Plate X on page 142, the relative amounts of these two hue components, which we refer to as the 'surround component' and the 'contrast component', respectively, are schematically represented by the thickness of the two wedges; We refer to the colour component which corresponds to the colour of the surround as the 'surround component' and the remaining colour component as the 'contrast component'. Now, if the subjects somehow disregard the surround component and base their judgements on the contrast component only, then the result that lines of constant hue converge on the surround colour would be understandable; The hue of the contrast component may remain constant although the total colour impression changes toward the colour of the surround. This would also explain the results obtained with the enclosure criterion; any circle around the surround chromaticity would contain a full range of hues in terms of the contrast component. But why should the subjects disregard the surround component? A plausible hypothesis is that the local proximal stimulus is decomposed into two simultaneous perceived colour components, like in the case of perceptual transparency (Metelli, 1970), colour scission (Anderson, 1997) or laminar segmentation (Mausfeld, 1998): The visual system attributes the surround component to a transparent overlay covering both the surround and the target patch, and therefore only the remaining contrast component is taken into account when the hue of the target patch is judged.

Interestingly, the above explanation in terms of colour scission suggests that true achromatic settings actually should be impossible to make, because two chromatic colour components should be perceived simultaneously. Colour Plate X illustrates the problems which subjects can be expected to experience when searching for the achromatic point: At the chromaticity of the surround, the target appears in the same colour as the surround, say yellow. If the chromaticity of the target is changed in the direction of a colour complementary to the surround colour, in this case blue, one would classically expect to find the achromatic point. However, our findings suggest that as soon as one moves in this direction, the contrast component is already blue. Making it less blue is impossible without going back to the surround chromaticity. Getting rid of a noticeable yellow surround component is also impossible without having a strong bluish contrast component. In other words, the subjects are unable to cancel the yellowish surround component and the bluish contrast component against each other in order to obtain an achromatic colour impression and must revert to simply trying to balance them such that the total colour impression is close, but not identical, to an achromatic colour impression.

The scission interpretation above is not only supported by our data but also by the phenomenological observation that, especially at low chromatic contrast, the contrast colour component of the target patch seems to be covered by a "veil" that has the same colour as the surround. A good impression of what this looks like can be gained in Colour Plate XI. In the top left panel, the 'spokes' of different colours from the hue circle are all shaded towards yellow at the centre. That is, the colours along the spokes correspond to lines in chromaticity space converging on the yellow colour at the centre. When the same set of spokes is embedded in a yellow surround (top right panel), they suddenly appear to be of roughly constant hue. Furthermore, one has the impression of a yellow transparent 'haze' or 'veil' covering the spokes, particularly towards the centre of the display where the contrast between the colours of the spokes and the surround is low. The lower part of the figure shows the analogous case where a red surround is used instead of a yellow one.

The stimulus display in this demonstration is more complex than the simple centre-surround stimuli used in our experiments, so one might argue that this complexity gives rise to more complex perceptual interpretations. Obviously, this cannot be ruled out *a priori*, but the correspondence between what one can observe in this demonstration and the hypotheses we have ventured in order to account for our findings with the simple centre-surround is rather striking. Thus, the demonstration actually may be said to illustrate the most important aspects of our findings at a single glance. Some further interesting aspects of simultaneous contrast are also suggested by this demonstration. It appears that the impression of transparency and the strength of the simultaneous contrast effect are correlated. Towards the ‘hub’ of the spokes, where contrast between the spokes and the surround is low, the most pronounced impressions of transparency are evoked. Also, the strength of the simultaneous contrast effect, i.e. the perceived difference between the physically identical sets of spokes embedded in the grey and the yellow surround, respectively, is most impressive at the centre. The perceived colours of the portions of the spokes which are farther away from the centre appear almost equal in the two surrounds, suggesting that simultaneous contrast is comparatively weak or even absent at high contrast, where there is no impression of transparency.

Impossibility of asymmetric colour matching This demonstration can also be used to illustrate a further implication of our findings, namely that true asymmetric colour matches should sometimes be impossible to make. In the left and middle panels of Colour Plate XII, two identical sets of spokes are embedded in a yellow and a red surround, respectively. In this case asymmetric matching is tantamount to changing the colours of the spokes in one of the surrounds, so that the two sets of spokes appear identical although they are embedded in different surrounds. In the right panel, the colours of the spokes embedded in the yellow surround have been changed such that they appear very similar to the spokes embedded in the red surround.[2] The most salient aspect of the remaining differences in the perceived colours of the two sets of spokes is that the one embedded in the red surround seems to be covered by a reddish haze, whereas the one embedded in the yellow surround appears to be covered by yellowish haze. Since the colour of the haze is determined by the colour of the surround, it cannot be made equal for the two sets of spokes as long as the two surrounds are different, and therefore true asymmetric colour matches are not feasible. This is at odds with all the classical models of simultaneous contrast discussed in Chapter 3, and the matter will be subjected to further experimental investigation in Chapter 5.

Dimensionality of perceived colour The classical models of simultaneous contrast tacitly assume that perceived colour can be adequately represented by a point in a three-dimensional space (W. S. Stiles, 1961). According to the von Kries model, for instance, the three-dimensional vector which results from multiplying the cone excitation vector of the target with a surround-dependent diagonal matrix should yield an index of the perceived colour of the target patch. This means that perceived colour of the target can be represented by an expression of the form $\mathbf{c}(T,S) = f_S(\mathbf{e}_T)$, where $\mathbf{e}_T$ is the cone excitation vector of the target T and f_S an invertible mapping which maps triplets onto triplets and depends on the surround S. In the case of the von Kries model, f_S is the transform given by the diagonal matrix of surround-dependent von Kries coefficients. This statement can be made for the two-process model and the contrast-coding model as well, except that f_S also involves a translatory component in these cases, hence all these models assume that the perceived colour of the target can be represented by a three-dimensional vector. This assumption is essentially equivalent to the statement that asymmetric colour matches should always be possible, apart from in some well-defined cases: A target T_1 embedded in a surround S_1 should appear equal to a second target T_2 embedded in another surround S_2 whenever

$$f_{S_1}(\mathbf{e}_{T_1}) = f_{S_2}(\mathbf{e}_{T_2}). \tag{4.1}$$

Since the mappings are invertible, we may solve for $\mathbf{e}_{T_1}$, obtaining $\mathbf{e}_{T_1} = f_{S_1}^{-1}(f_{S_2}(\mathbf{e}_{T_2}))$. Accordingly, the subject may establish a match between T_1 and T_2 just by setting the cone excitation values of the former to this value. Asymmetric matches will therefore always be possible, except

---

[2] This was achieved by shading the same set of endpoint colours towards yellow instead of red.

for the cases where the vector $f_{S_1}^{-1}(f_{S_2}(\mathbf{e}_{T_2}))$ is located outside of the colour cone and therefore does not represent a possible combination of cone excitation values.

In these cases, one would expect the best possible setting to be a cone excitation vector located at the border of the colour cone: Since $f_{S_1}^{-1}(f_{S_2}(\mathbf{e}_{T_2}))$ is outside the colour cone, a point at its border will obviously be closer to it than any other point in the cone. The problems with establishing asymmetric matches which can be expected to occur based on the present findings are however clearly not of this kind. Our findings suggest that a target patch which appears approximately grey when presented in a red surround cannot truly be matched by any target presented in a green surround. The most similar alternative, though, will not be located on the border of the colour cone, but instead well within it. Accordingly, the asymmetric matching problems implied by our findings suggest that the assumption that perceived colour can be represented in a three-dimensional space is untenable. As a general hypothesis, this has been proposed by a number of investigators (Katz, 1911; Gelb, 1929; Evans, 1964, 1974; Kanizsa, 1966; Heggelund, 1974, 1992; Mausfeld, 1998) previously. Based on commonly accepted empirical assumptions, Niederée (1998) has provided formal proof that any index of perceived colour which is a continous function of the target and surround cone-excitation values must be at least four-dimensional. The present findings accord well with this conclusion.

Apparent surface colour vs. unasserted colour It has long been known that the results of asymmetric colour matching experiments may depend on how the observers are instructed (Henneman, 1935). If the observer is asked to make the two target patches appear to ‘emit the same amount of light’, different results are generally obtained than when he is instructed to make the two targets appear as though they were ‘cut of the same paper’ (Arend & Reeves, 1986; Arend & Goldstein, 1987; Arend & Spehar, 1993a, 1993b). This finding suggests that the observer can distinguish between two different perceptual aspects of the target stimulus. In the case of achromatic stimuli, which were most often used in this line of research, these two ‘aspects’ are commonly referred to as brightness and lightness, respectively, whereby ‘brightness’ is defined as the perceived intensity of the light emitted by the stimulus and ‘lightness’ is understood to be its perceived reflectance. In the more general case of coloured stimuli, the corresponding terms ‘unasserted colour’ and ’apparent surface colour’ are used to make the analogous distinction (Arend, 1994).

In the present experiments, the subjects were simply instructed to make the standard patch appear as achromatic as possible or make pairs of target patches to appear as similar in hue as possible; no reference to distinction between unasserted colour or apparent surface colour was made. Since most of the subjects were experienced psychophysical observers familiar with this distinction, some of them asked which criterion to apply. Those subjects were told to base their settings on the ‘total colour impression’, i.e. on the ‘unasserted colour‘. The differences between the directly determined achromatic point and the convergence point for lines of constant hue found in our experiments can therefore not be due to different instructions. Some informal observations may be of interest in this connection. In my own experience as an observer in these experiments, I found the task of finding the achromatic point rather difficult, and was often uncertain whether I had really been able to find the best setting. The task of making two targets appear equal in hue, on the other hand, felt very natural and there was never any doubt as to what was the best setting. The difference in the ease and ‘naturalness’ of the two tasks is reminiscent of the subjective uncertainties and unease experienced when making asymmetric matches using side-by-side displays which are contrasted by the subjective ease experienced when haploscopically superimposed displays (HSD) are used instead. Whittle (1994b, p.41) noted that the HSD “makes setting a brightness match between patches on different backgrounds of different luminance subjectively easy. The measurements cry out to be made: It is almost like doing visual photometry. In side by side displays the task is much more difficult, and subjects often feel there is no satisfactory match". It seems, therefore, that the hue matching tasks in our experiments share a number of interesting characteristics with asymmetric colour matching using the HSD:

1. The matches are subjectively easy to perform and one is never in doubt as to what is the best setting.

2. The results obtained are consistent with contrast coding models.

3. The two targets which are to be matched (for colour in the HSD and just for hue in our case) are perceived against a common surround.

It should be kept in mind, though, that although there is no doubt as to what is the best hue match, the patches nevertheless look different. As already noted, the patches appear to be covered by a transparent haze of the same colour as the surround, and this impression was most pronounced for the patch which has the least contrast relative to the surround. If the surround is red, this patch therefore appears to be covered by a more pronounced red veil than the other one. It appears, though, that this difference cannot be reduced whatever setting one chooses, and this seems to be the reason why there is no doubt as to what is the best setting: One simply has no choice but to disregard this unchangeable difference and concentrate on making the remaining parts of the colour impressions (the 'contrast component') as similar as possible. In the case of the direct achromatic settings, however, one can reduce the salience of the reddish haze by choosing a setting farther away from the the chromaticity of the surround. This tends to make the contrast component appear more greenish, though, and one is uncertain what mixture of a reddish haze and a greenish patch would be the best compromise. One way of thinking about these phenomena which I find helpful is the idea that there actually is no such thing as an 'unasserted colour' impression (in the case of our stimuli, at least): The stimulus evokes a perceptual decomposition of the proximal colour signal which cannot by any act of will be reduced to a simple unsegmented colour impression. It is a widespread notion that by assuming a particual perceptual attitude, commonly referred to as the 'proximal mode' or 'painter's mode', it is possible to overcome or actively disregard the transforming influence of higher perceptual mechanisms and see something like an 'preperceptual' local colour sensation which essentially reflects the properties of the local retinal stimulation, or proximal stimulus. When trying to make an achromatic setting in our experiments by assuming such a mode, i.e. by judging 'unasserted colour', though, one soon finds oneself reverting to very 'cognitive' criteria like "Could this balance of reddishness and greenness pass for a good grey?" or "If I disregard the reddishness and the greennness, is this setting the one which has the most salient grey content?" Thus, although one tries to adopt a proximal mode, what one acutally ends up doing does not feel very 'preperceptual' and it actually feels less natural than what one does in the hue matching task, although the latter involves 'discounting' a simultaneously perceived transparent layer.[3] These informal observations are also supported by our experimental data. It is evident from the data plots in Figure 4.5 on page 75 that the settings for the achromatic point differ substantially across the individual observers, whereas the settings for the convergence point for lines of constant hue are much more similar for the different subjects.

Stimulus complexity, perceptual scene analysis and colour scission In his treatment of simultaneous colour contrast, Helmholtz (1911), noted that two complementary colours can be perceived simultaneously on the same point of the visual field:

---

[3] The problematic issue of whether there exists such a thing as 'proximal mode' or 'unasserted colour' and what function it may serve in the perceptual process is discussed more thoroughly and from a slightly different angle by Niederée (1998) and Mausfeld (2003a, 2003b).

> "Thus one believes to see two colours at this location, namely the green, that one attributes to the pane of glass, and the roseate, that one attributes to the paper lying behind it, and both combined do indeed produce the true colour at this location, namely white." (Helmholtz, 1911, p. 242).[4]

Hering (1887a, 1887b, 1887c, 1888), however, vehemently contested this and similar observations made by Helmholtz:

> "As far as the assumed scission of white into two complementary colours is concerned, I shall return to this later. At present, I merely wish to note that this assumption directly contradicts what one actually perceives in this experiment. At the location of the patch one sees but *one* colour, namely roseate, not green and roseate simultaneously, one in front of the other; no-one notices any trace of green here." (Hering, 1887b, p. 13).[5]

In the early days of colour science, this matter was the issue of considerable debate. A major reason why it was considered important is probably that the putative existence of dual colour impressions of complementary were considered to be at odds with a basic tenet of Hering's (1920) opponent colours theory, namely that two opponent colours such as red and green cannot coexist, but instead neutralise each other, yielding an achromatic equilibrium colour (cf. Section 2.3). To Helmholtz, on the other hand, who argued for a 'psychological' explanation of simultaneous contrast in terms of 'unconscious inference', the coexistence of opponent colour impressions was not only unproblematic, but also an argument in favour of his theory: According to Helmholtz (1911), the visual system somehow estimates the colour of the illuminant based on rules learned on the basis of previous perceptual experience, and takes this estimate into account when inferring the true colour of an object. From this perspective, the existence of dual colour impressions can be taken to mean that not only the inferred true colour of an object, but also the perceptual estimate of the illuminant colour is available to conscious experience. Although this is not a necessary prerequisite for the correctness of Helmholtz' theory, since, in principle, a perceptual estimate of the illuminant may play a functional role without being directly available to conscious experience, it clearly makes the theory more intuitively plausible: That the perceptual estimate of the illuminant may actually be perceived changes its status from a merely theoretical explanatory device to a directly observable fact. As such, it becomes in itself a target of scientific explanation. To Helmholtz, both the phenomenon that the colour of the central patch is changed and the phenomenon that one has an impression of a coloured illumination or transparent veil through which the colour of the central patch is perceived are two empirical facts for which he provides a common parsimonious explanation. To Hering, on the other hand, the putative perceived illumination colour is a non-fact, only the change in the colour of the patch is real and in need of explanation. Clearly this suited him well, not only because there is then no threat to his opponent colours theory, but also because any impression of the illuminant cannot easily be accounted for by his explanation of simultaneous contrast in terms of lateral interactions between neural elements on the retina.

From the above it is clear that the factual correctness of Helmholtz' observations was of central theoretical importance, and accordingly much empirical work was directed towards confirming or

---

[4] My translation of the original German: "Man glaubt also an dieser Stelle gleichzeitig zwei Farben zu sehen, nämlich das Grün, welches man der Glasplatte zuschreibt, und das Rosenrot, welches man dem dahinter liegenden Papier zuschreibt, und beide zusammen geben in der Tat die wahre Farbe dieser Stelle, nämlich Weiß."

[5] My translation of the original German: "Was die hier angenommene Spaltung des Weiss in zwei complementäre Farben betrifft, so werde ich anderwärts darauf zurückkommen. Hier will ich nur bemerken, dass diese Annahme in directem Widerspruche mit dem steht, was man bei diesem Versuch wirklich wahrnimmt. An der Stelle des Schnitzels sieht man eben nur *eine* Farbe, nämlich Rosenroth, nicht aber Grün und Rosenroth zugleich, das eine über dem andern; von Grün bemerkt hier Niemand eine Spur."

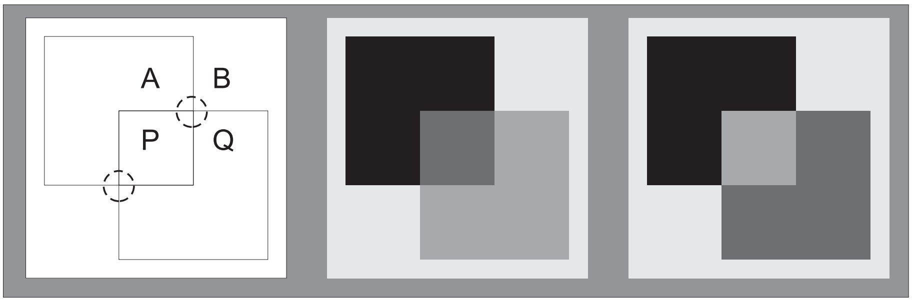

Figure 4.14: If a transparent medium corresponding to the regions P and Q covers a bipartite background, a stimulus with four regions is generated. Here, the kind of four-region stimulus traditionally employed in investigations of perceptual transparency is shown. In the middle panel, one has the impression of a transparent gray layer corresponding to the regions P and Q, covering a bipartite background A and B. In the right-hand panel, the luminances of the regions P and Q have been exchanged. The stimulus does not appear transparent anymore. Thus, the perception of transparency depends on the relation between the four colours (here luminances) of A, B, P and Q. Information about these relations is available at so-called X-junctions in the image, indicated by the dotted circles in the left-hand panel.

disconfirming it; the issue of whether two opponent colours can be perceived at the same location of the visual field became an important field of inquiry (e.g. Schumann, 1921). Historically, this line of research can be considered to be the precursor of the more modern research on perceptual transparency, which is generally based on Metelli's (1970) episcotister model (Da Pos, 1989). The central insight of this model is that a stimulus which contains only four differently coloured uniform surfaces may contain enough information for perceptual detection and reconstruction of a transparent layer. Such a four-region stimulus would result if a bipartite background is covered by a uniform transparent medium (see Figure 4.14). At the locations of the visual field where the border of the transparent medium crosses the border between the two parts of the background, so-called X-junctions are generated, and the relations between the colour signals from the four regions meeting at this junction is generally considered to be a cue for transparency (Adelson & Anadan, 1990). A theoretically important implication of this model is that four-region stimuli may be considered the simplest possible kind of stimulus providing enough cues for a perceptual parsing into background and transparent layer. Though the model may also be applied to stimuli which objectively only contain three different regions, this corresponds to cases in which figural cues lead to the perception of four different regions (Anderson, 1997; Ekroll & Faul, 2002). In these cases, so-called T-junctions, or 'implicit X-junctions' (Watanabe & Cavanagh, 1993), are thought to constitute the cues for transparency. In the case of even simpler stimuli, though, applying the Metelli model is impossible. The simple centre-surround stimuli used in the study of simultaneous contrast is a case in point: They contain neither X- nor T-junctions, hence Metelli's model instils the expectation that perceptual transparency should not be evoked. Thus, the Metelli model suggests that the visual system should treat simple centre-surround stimuli in a manner which differs fundamentally from the way it treats slightly more complex stimuli. This theoretical dichotomy is intuitively rather plausible: Stimuli containing X- or T-junctions, may evoke rather compelling impressions of transparency. One may even produce stimuli in which one has a rather distinct impression of two opponent colours coexisting on the same location of the visual field, for instance because a red background is perceived through a green transparent layer (see Colour Plate IX on page 141). The phenomena occuring in these complex stimuli can clearly not be accounted for by simple Hering-type explanations in terms of lateral inhibition and adaptation, but

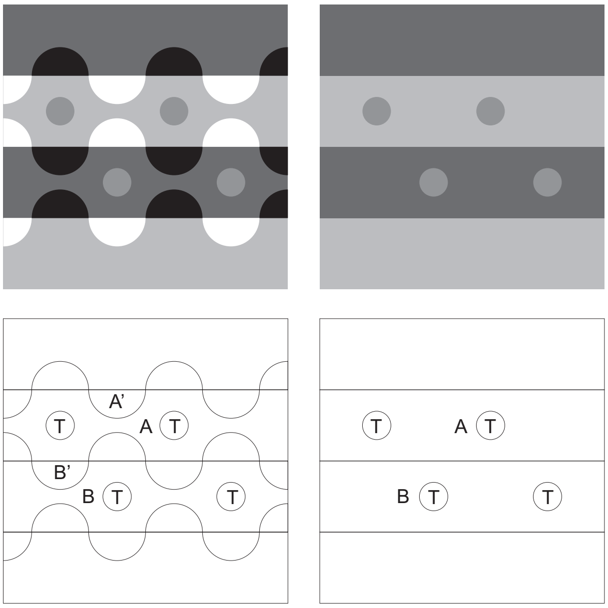

Figure 4.15: Left: Adelson’s (2000) Snake Illusion. The targets T are physically identical, yet the upper ones appear much darker than the lower ones. This effect cannot be accounted for by lateral inhibition, since the targets and the immediately surrounding regions A and B are exactly identical in the right-hand panel, where the effect is substantially reduced (actually, the only difference between the left and right displays is the luminances of the regions A’ and B’). A Helmholtzian explanation appears reasonable, though, since the upper and lower targets appear to be located in regions of brighter and darker illumination, respectively. Alternatively, one may say that the lower targets appear to be covered by a darkening transparent layer.

lend themselves naturally to an explanation in the spirit of Helmholtz since one has a distinct impression of something like a veiling illumination or a transparent layer. In simple centre-surround stimuli, on the other hand, such impressions are far less obvious, and it appears plausible that the simultaneous contrast effects occuring in these stimuli are better accounted for by lateral inhibition and adaptation.

In recent years, several theoretical developments and experimental findings suggest, in accordance with the above dichotomy, that many phenomena of brightness and colour perception occuring in stimuli containing X- or T-junctions can be accounted for by assuming that the visual system corrects for or takes into account the inferred presence of a transparent medium (Anderson, 1997; Adelson, 1993, 2000). Adelson’s (2000) well-known *Snake Illusion*, which is shown in Figure 4.15, illustrates this point nicely. The upper and lower pairs of target disks T are physically identical, yet the upper ones appear very much darker than the lower ones. This effect cannot be accounted for by a simple mechanism which takes only the luminance of the immediate surround into account: In the right panel, the same stimulus display is shown, except that the semi-disks A' and B' have been removed. Though the immediate surrounds A and B of the target disks T are physically identical in the left and the right panel, the perceived difference in brightness of the upper and lower targets is much larger in the left display. Adelson (2000) explained the enhanced effect in the left display in terms of an Helmholtzian perceptual mechanism. In the left stimulus display, one has a rather vivid impression of a shadow or a darkening transparent layer covering the regions B and B' as well as the lower pair of target disks. One may also say that the regions A and A' and the upper pair of target disks appear to be more intensely illuminated. If the visual system discounts these differences in perceptually estimated illumination or transparency in order to estimate the reflectance of the targets, one would expect it to estimate a higher reflectance for the target which appears to be covered by a darkening transparent layer, hence the lower targets should appear to be lighter. In the right panel, corresponding differences in perceived illumination (or transparency) are not evident, and accordingly the effect is much smaller. There is a small remaining effect though, and it appears reasonable to assume that this effect is due to another mechanism; A simple Hering-type mechanism of lateral inhibition suggests itself. Thus, this phenomenon supports the abovementioned theoretical dichotomy in both theoretical as well as phenomenal regards: In the left display, which contains X-junctions, compelling impressions of transparency are evoked, and there is a strong lightness illusion. In the right display, on the other hand, there are neither X-junctions nor compelling impressions of transparency, and the lightness illusion is much weaker.

The present findings are very much in line with Helmholtz' (1911), Anderson's (1997) and Adelson's (1993, 2000) emphasis on the role of colour scission in lightness and colour perception. They go beyond Anderson's and Adelson's theories and findings, however, in suggesting that colour scission plays an important role also in simple centre-surround stimuli, although they contain neither X- nor T-junctions. Our findings may be said to corroborate Helmholtz' (1911) original phenomenal observations made with simple centre-surround stimuli. Unlike Helmholtz' observations on simultaneous contrast, though, which were made with experimental apparatus actually involving perceiving the stimulus through real, physically transparent media, such as in one case a half-reflecting mirror (Scina's experiment, 1847) and in another case a transparent piece of fabric (Meyer's experiment, 1855), and therefore were easy targets for Hering's (1887a, 1887b, 1887c, 1888) critisism, the present observations were made under well-controlled experimental conditions: Even in the absence of physical transparency, Helmholtz' reports of seeing two colours in the region corresponding to the target in a centre-surround stimulus can be confirmed. Our findings also suggest an answer to the question how Helmholtz could be so positive about seeing two colours simultaneously in the region of the target whereas Hering could very decidedly claim seeing but a single colour: Perceptual transparency is only evoked in centre-surround stimuli when the contrast between the target and the surround is low. Viewing the stimulus through a physically transparent medium, as Helmholtz did, reduces the contrast between target and surround. Therefore, it is likely that his impressions of transparency were not due to the physically transparent media themselves, but to the fact that this produced a stimulus with low contrast, which in view of our observations suggests itself to be a cue in the proximal stimulus evoking the perception of transparency (see also Masin & Idone, 1981; Brenner & Cornelissen, 1991; Mausfeld, 1998; Koenderink, 2003). The corroboration of Helmholtz' observations under the present more well-controlled conditions suggest that it is necessary to draw the conclusion that Hering was eager to avoid: Lateral inhibition cannot yield a complete account of simultaneous contrast even in simple centre-surround stimuli. Our findings do not preclude that it may play a role, but at least it is clear that lateral inhibition in itself cannot account for the dual colour impressions.

Other evidence suggesting that the functionally neutral point differs from the achromatic point According to opponent colours theory, the phenomenally achromatic point is considered to be identical to a functionally neutral point relative to which neural mechanisms code all other colours. Our findings, though, suggest in accordance with contrast-coding models that colour (in this case constant hue) is computed relative to the colour of the surround, whatever colour the surround might have. Interestingly, a similar dissociation of the achromatic point and a functionally relevant neutral point has been found in a study on contrast adaptation (Webster & Wilson, 2000; Webster, 2003). Here, the experimentally determined centre of contraction in contrast adaptation was not the achromatic point but the mean of a temporally modulated adapting stimulus. Although it is not obvious whether or how their findings may be related to the present observations, the results of both their and our experiments agree in suggesting that the functionally neutral point may differ from the phenomenally defined achromatic point.

An experimental finding which appears to be more directly related to the present ones was recently reported by Sakata (2004b, 2004a). His subjects matched the colour appearance of afterimages produced by small targets presented in a uniform coloured surround. A novel feature of his experiments which allowed rather precise measurements of the colour appearance of the afterimages was the manner in which the afterimages were produced: The subjects simultanously viewed a number of small targets arranged in a circle and embedded in a uniform coloured surround, much as in our experiments. Different from our experiments, though, all of the targets had the same colour. During the viewing of this stimulus, one of the targets was always blanked, whereas all of the others remained visible. Which of the targets from the circle was blanked changed in a cyclical and rapid manner. Thus, at any point of time, the subject perceived a fresh afterimage at the location where a target was blanked. The afterimage therefore may be said to continously change its spatial position, but to remain very constant in perceived colour. The colour of a further test patch embedded in the same surround could be adjusted by the subjects such that it appeared to have the same colour as the afterimage.

Since, according to opponent colours theory, afterimages should have a colour which is the complementary to the colour of the stimulus which produced it, one may expect that the neutral point should be located somewhere on the line in chromaticity space connecting the chromaticity of the target used to produce the afterimage and the chromaticity of the patch which was judged to appear equal in colour to the afterimage. If measurements are made for several different afterimage-inducing targets, the resulting set of lines can be expected to intersect at the neutral point. Interestingly, Sakata found that this point of intersection coincided with the chromaticity of the surround, even though the surround was coloured. Accordingly, the estimate of the neutral point obtained in Sakata's experiments suggests that the chromaticity of the surround is the neutral point. Obviously, this finding parallels our results obtained with the convergence and enclosure criterion settings.

Interim Conclusions Our results strongly suggest that perceptual transparency is evoked in simple centre-surround stimuli when contrast between the target and the surround is low. The colour of one of the phenomenal layers of this dual perceptual representation seems to be well described by contrast-coding models of simultaneous contrast (Whittle & Challands, 1969; Whittle, 2003; Walraven, 1976), whereas the colour of the other phenomenal layer corresponds to the colour of the surround. Due to this laminar segmentation (Mausfeld, 1998) of the proximal stimulus it appears to be impossible to find a setting for the target patch such that the total colour impression appears truly achromatic. If subjects are nevertheless asked to do so, they arrive at some compromise, and the settings obtained in this case are consistent with Shevell's (1978) dual process model. This interpretation of the present findings may be said to provide a new perspective on the long-standing Walraven-Shevell controversy: Both models provide a good description of different aspects of the phenomenon, but they must both be considered inadequate because they fail to accomodate for the dual aspect of the colours perceived. Consistent with previous claims (Evans, 1964; Niederée, 1998), our findings strongly suggest that the colours perceived in simple centre-suround stimuli cannot be adequately described by a three-dimensional colour code: More dimensions than three are needed to obtain a complete representation of the dual colour impressions. Another way to phrase this is to say that the cone excitation values of the target and those of the surround do *not* act in a compensatory manner to determine the perceived colour of the target. This not being the case, classical psychophysical methods for measuring simultaneous contrast, such as grey settings and asymmetric colour matching can be expected to yield results fraught with artifact. This matter is pursued further in the next chapter.

Chapter 5

Kinds of simultaneous contrast

The results of the experiments reported in the previous chapter suggest that the colour changes induced into a target patch when it is embedded in a uniform coloured surround cannot always be compensated by changing the cone excitation values of the target patch. This is, however, exactly what subjects are asked to do in experiments using classical methods for quantifying simultaneous colour contrast such as grey settings and asymmetric colour matching. Clearly, if subjects are asked to do something which they are, in fact, not able to do, they may resort to doing something else, and interpreting the data obtained may lead to erroneous conclusions.

According to our findings, one would expect asymmetric colour matching to be difficult because, at low contrast, one perceives a transparent layer having the colour of the surround extending ‘across’ the target patch. Accordingly, if two differently coloured surrounds are used, the colours of the transparent layers will also be different, making a true colour match impossible. By this logic, one would expect asymmetric colour matching to be difficult whenever the surrounds used are perceived as having different colours: Only then, one should perceive differently coloured transparent layers. This means that asymmetric colour matching should be difficult when side-by-side displays are used, in which the different colours of the surrounds are readily perceived, whereas it should be unproblematic when haploscopically superimposed displays (HSD) are used, since the physically different surrounds are then merged into a common percept. Indeed, as noted by Whittle (1994b, p. 48), establishing asymmetric colour matches is subjectively much easier when the HSD is used than when the surrounds are perceived side-by-side. A further, possibly related difference between these two ways of performing asymmetric colour matches is that the HSD yields a comparatively simple pattern of results (Whittle, 1994b; Chichilnisky & Wandell, 1995; Niederée & Mausfeld, 1997; Shepherd, 1997; Richter, 2002), whereas the data obtained with side-by-side displays are much more complicated (Whittle, 1994a). Smith and Pokorny (1996), for instance, obtained a pattern of results much more complicated than the almost linear data regularly reported in studies using the HSD. Based on the above deliberations, one may surmise that the more complicated data obtained with side-by-side displays may be directly related to subjective matching problems. In the present chapter, I present evidence from asymmetric colour matching experiments demonstrating that this is indeed the case.

5.1 Experiment 4

Stimuli and procedure All stimuli were presented on a CRT computer monitor, which was colourimetrically calibrated by means of a colourimeter (LMT C1210) following a standard procedure (Brainard, 1989), and controlled by a graphics card yielding a colour resolution of 8 bits per RGB channel. Viewing distance was approximately 90 cm, and the monitor, which was the

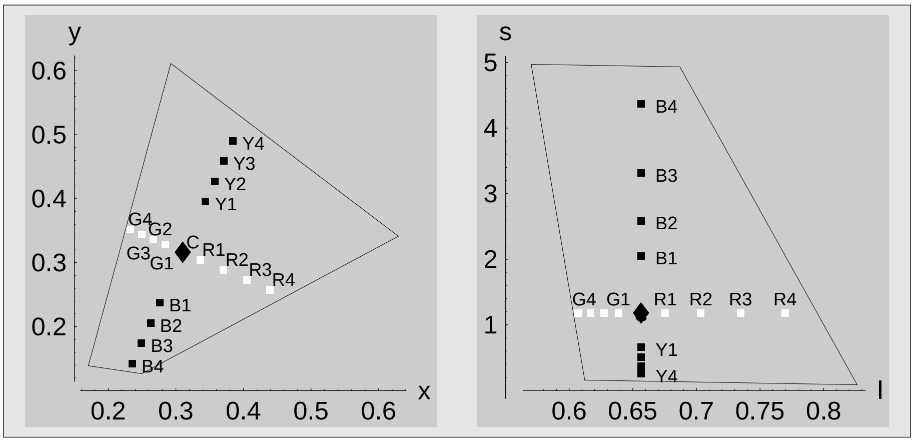

Figure 5.1: Chromaticities of the surrounds employed in experiment 4 plotted in the CIE **xy**-chromaticity space (left) and the MacLeod-Boynton (1979) chromaticity space (right). In both cases the polygon represents the gamut of the monitor at the luminance used in the experiment.

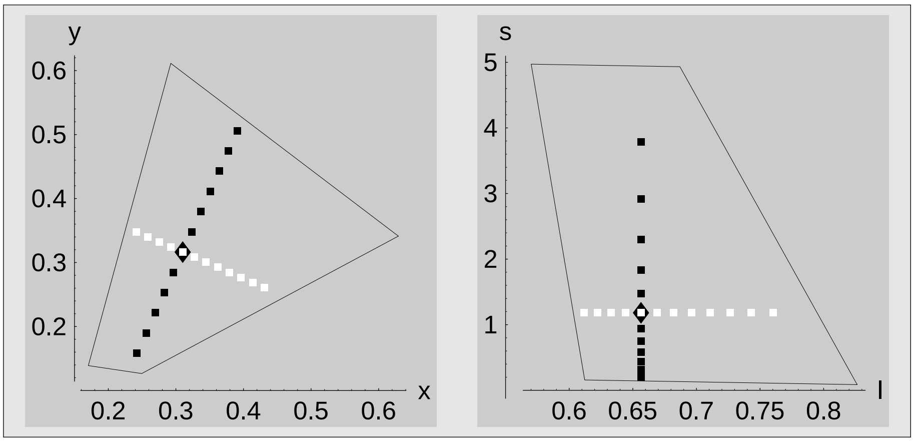

Figure 5.2: Chromaticities employed for the fixed patch in experiment 4a (white symbols) and experiment 4b (black symbols) plotted in the CIE **xy**-chromaticity diagram (left) and the MacLeod-Boynton chromaticity diagram (right).

only source of illumination in the room, was located within a viewing box covered by black velvet on the inside. The subjects viewed two square patches (subtending $1.9°$ visual angle), each centred in the middle of a square surround ($9.6°$). The centre-to-centre distance between the two centre-surround stimuli was $12.2°$. All parts of the stimuli except the dark general background were equiluminant at 10 cd/m^2. The luminance of the adjustable patch was also restricted to 10 cd/m^2, but the chromaticity could be varied freely within the gamut of the monitor (see fig. 5.1) by using the arrow keys of a keyboard. For both surrounds and the fixed patch we used chromaticities which were all located on the same cardinal axis (Krauskopf et al., 1982). The cardinal axes were defined as the lines through the (**l**,**s**) co-ordinates (0.656, 1.182) of Illuminant C (Wyszecki & Stiles, 1982) in the MacLeod-Boynton ((1979)) chromaticity diagram with constant **l** and **s** values, respectively. The computation of $\mathbf{l} := L/(L+M)$ and $\mathbf{s} := c \cdot S/(L+M)$ is based on the cone excitations L, M and S as estimated by Smith and Pokorny (Wyszecki & Stiles, 1982, p. 615). The **s** co-ordinate was scaled by a constant c such that it was 1 for equal energy white. For isolated light spots, the perceived colours corresponding to the chromaticities on the **s**-axis may be said to vary from a reddish blue to a greenish yellow, those corresponding to the chromaticities on the **l**-axis range from red to green. For brevity, and keeping in mind that these axes are not to be confused with the unique hue loci of opponent colours theory (Valberg, 2001), we shall refer to the axes as the blue-yellow and red-green cardinal axis, respectively. All the surround chromaticities used in the present experiments are shown in figure 5.1. In the left part of the figure the CIE (**x**,**y**)-chromaticities of the surrounds are plotted, and in the right part the corresponding MacLeod-Boynton (**l**,**s**)-chromaticities. In experiment 4a, we used the following pairs of chromaticities (plotted in figure 5.1) from the blue-yellow cardinal axis for the surrounds of the fixed and the adjustable patch [given in the order (fixed, adjustable)]: (C, B_i), (C, Y_i) and (B_i, Y_i), with $i = 1, 2, 3, 4$, and for the fixed patch we used the chromaticities from the same cardinal axis which are plotted as black symbols in figure 5.2. In experiment 4b we chose the pairs (C, R_i), (C, G_i) and (Ri, G_i) from the red-green axis as surround colours and for the fixed patch we used the chromaticities from the same cardinal axis plotted as white symbols in figure 5.2.

The task of the subjects was to make the perceived colours of the central patches as similar as possible by manipulating the chromaticity of the adjustable patch. In each experimental session matches were made for 12 different fixed patches and one pair of surrounds. Since each match was repeated 3 times for each subject, a single experimental session consisted of 36 single matches.

Three naïve, but experienced psychophysical observers (MK, BS and GW) participated in the experiments. All were colour-normal according to the Ishihara Tests for Colour-Blindness. Author VE also performed the experiment. His data were very similar to those of observer GW and are not reported.

Results Consistent with previous reports (Wuerger, 1996; Rinner & Gegenfurtner, 2002), the asymmetric matches[1] differed mainly in the same chromaticity co-ordinate (either **l** or **s**) as the surround pairs. Differences with respect to the other co-ordinate were negligible. In figure 5.3 every single setting made in the experiment is plotted, those for surround pairs differing in the **s** co-ordinate on the left, and those for surround pairs differing in the **l** co-ordinate on the right.

The white lines in figure 5.3 represent the constant **l** co-ordinate of the fixed patch in experiment 4a (left panel) and the constant **s** co-ordinate of the fixed patches in experiment 4b (right panel). Apart from the small deviations from these lines, plots showing only one of the chromaticity co-ordinates yield a complete representation of the data. In figure 5.5 the results for the surround pairs differing in the **s** co-ordinate are shown. The format of these plots is explained in

---

[1] In the following we use the term asymmetric match for a best match in the sense of the above instruction. Whenever perceptual identity is implied we shall speak of a *true asymmetric match.*

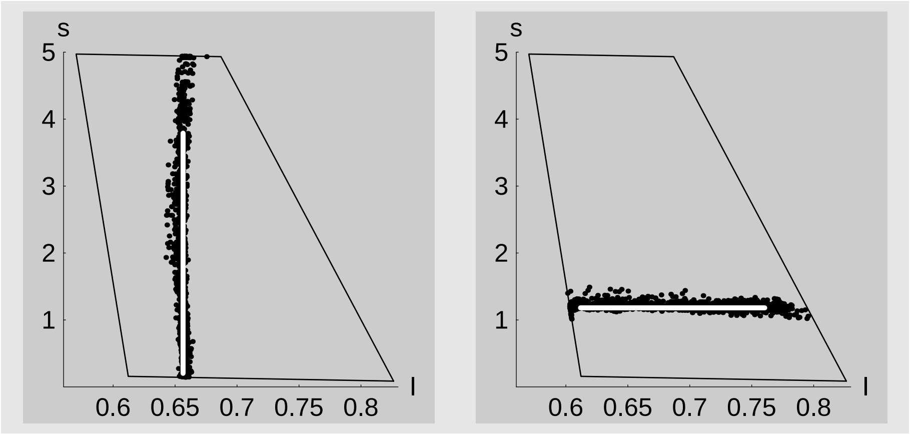

Figure 5.3: Complete plot of all the settings chosen by the subjects in experiment 4a (left) and 4b (right) plotted in the MacLeod-Boynton chromaticity diagram.

figure 5.4. The mean **s** co-ordinate of the subjects' settings (vertical axis) is plotted against the **s** co-ordinate of the fixed patch (horizontal axis). Each data point is based on nine individual settings, three from each of the three subjects, and the error bars represent one standard deviation in each direction. The top four panels represent the data for the surround pairs (C, B_i), $i = 1, 2, 3, 4$, the middle four panels for the surround pairs (C, Y_i) and the bottom four panels for the surround pairs (B_i, Y_i). In each panel, the point where the two white rectangles meet represents the **s** co-ordinates of the surround pair. The projection of this point on the horizontal axis gives the co-ordinate of the surround in which the fixed patch was presented, the projection on the vertical axis gives the co-ordinate of the surround of the adjustable patch. The two white rectangles also show the regions in which data points representing an increment-decrement match with respect to S-cone excitation would fall[2]. A center-surround stimulus is said to be incremental (decremental) with respect to dimension X whenever its center has a larger (smaller) x-value than its surround. Based on contrast-coding models, it is to be expected that increments can never be matched to decrements, neither with respect to luminance nor with respect to cone excitations (Whittle, 2003; Kingdom, 2003).

The induction effects, defined as the difference in the co-ordinates of the fixed and the adjustable patch, are apparent in the plots as the vertical distance between a given plot point and the diagonal line, which is where plot points would fall in the absence of any induction effect. Consistent with previous observations (Jameson & Huvich, 1972; Walraven, 1976; Shevell, 1978), the data cannot be accounted for by simple von Kries adaptation (Kries, 1905). If this were the case, the data in each plot should fall on a single line through the origin, which is clearly not the case. A very conspicious regularity of the data is a step-shaped bump, henceforth referred to as the 'step', where the plot points seem to graze the border of one of the white rectangles. This 'step' gets larger with increasing difference between the **s** co-ordinates of the surround pair, whereas the induction effects outside this region, represented by the plot points which do not graze the white rectangle, are but moderate in comparison. Thus, on a purely descriptive level, it appears that a major part of the induction effects is accounted for simply by stating that the plot points graze the white rectangle. This means that subjects avoid making S-cone increment-decrement matches, consistent with the predictions of contrast-coding (Whittle, 2003). However, it also re-

---

[2] Due to the equiluminance of the stimuli, an S cone increment-decrement match occurs whenever an **s** co-ordinate increment-decrement match occurs.

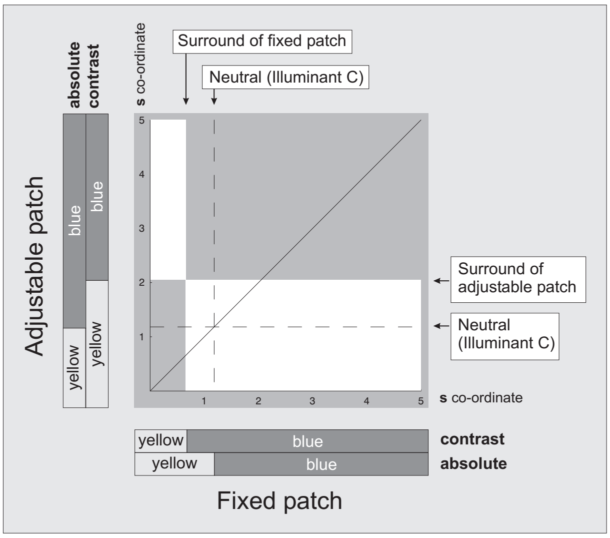

Figure 5.4: Explanation of the general format of the data plots. The horizontal and vertical axes represent the **s** co-ordinate of the fixed and the adjustable patch, respectively. The diagonal line shows where the data points would fall in the absence of any induction effect. The dashed lines show the location of a nominally neutral stimulus (Illuminant C) on each axis. In each plot the horizontal co-ordinate of the point where the two white rectangles meet represents the **s** co-ordinate of the surround of the fixed patch, the vertical co-ordinate that of the surround of the adjustable patch. The white regions show where data points would fall in case of a match between an **s** co-ordinate increment and an **s** co-ordinate decrement. The bars along the axes illustrate where colour impressions should change from blue to yellow assuming either “absolute” coding (relative to Illuminant C) or contrast coding (relative to the surround).

veals more. Referring, for example, to the top panels in figure 5.5, the horizontal portion of the data curves means that several different **s** co-ordinate increments presented in the neutral surround are matched to the same patch chromaticity in the blue surround, namely to the chromaticity of the blue surround itself.

Casual observations suggest that the reason for this curious pattern of results is that the range of **s** co-ordinate increments presented in the neutral surround, which appear less bluish than the blue surround, cannot be reproduced by any possible setting for the patch presented in the blue surround. It is, in other words, not possible to realise blue colour impressions in the blue surround which are less blue than the surround itself: As soon as one has a perceptible **s** co-ordinate decrement in the blue surround, the colour impression of the patch splits into two components having complementary hues: a yellowish contrast component, and a background component of the same colour as the blue surround (see previous chapter). Since a true match between the ‘pure’ blue colour impressions in the neutral surround and the ‘yellowish’ blue colour impressions in the blue surround is impossible, the subjects can at best try to minimise the differences in the perceived color of the central patches. Depending on which aspect of the colour impression is regarded most important, one can distinguish different strategies.

One strategy (‘A’) would be to avoid that the two central patches contain complementary colour components. This can be achieved by avoiding **s** co-ordinate decrements in the blue surround. However, this strategy has the unfortunate consequence that the patch in the blue surround now appears too blue, a discrepancy which is the more evident and disturbing the more saturated the blue surround is. An alternative strategy (‘B’) would be to accept complementary colour impressions in the two central patches, in order to avoid the differences in blueness which would result

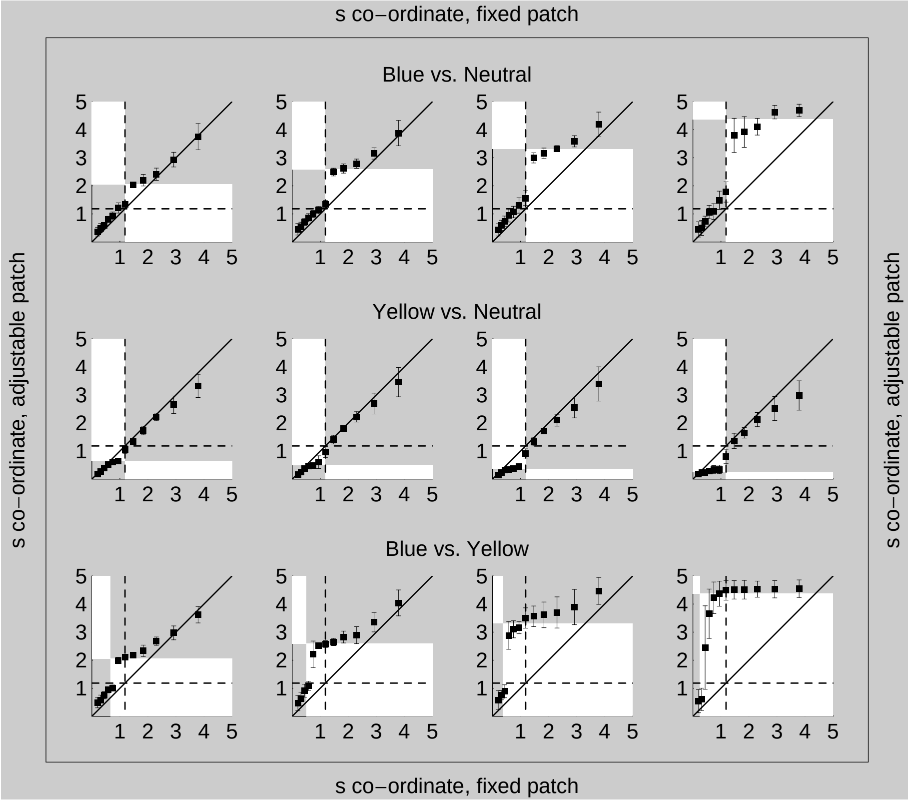

Figure 5.5: Results from experiment 4a. Each data point represents the mean of 9 settings (3 repetitions for each of the 3 subjects). Error bars represent one standard deviation in each direction. See fig 5.4 for further explanations.

under strategy A.

Consistent with this reasoning, there were marked inter-observer differences evident in the data, which may be attributed to different strategies. In figure 5.6 individual plots for one of the surround pairs differing in the **s** co-ordinate are shown. Observers GW and BS almost always avoided increment-decrement matches, whereas observer MK sometimes did not. These data suggest that when confronted with the dilemma outlined above, our subjects made their settings according to different strategies: Observers GW and BS seem to prefer strategy A, whereas observer MK seems to switch from strategy A to strategy B when the perceived differences in blueness of the central patches would get very large if one were to stick to strategy A. All of the abovementioned patterns in the data for surrounds differing in the **s** co-ordinate are also evident in the data for the surrounds differing in the **l** co-ordinate. Figure 5.7 shows the pooled data for these surround pairs, and figure 5.8 shows individual plots for one of the surround pairs demonstrating analogous inter-observer differences.

The data of experiment 4 show a rather complex pattern of results for each surround pair. However, the plots for all of the surround pairs share a characteristic ‘step’ which is associated with subjective matching problems and a general trend towards avoiding increment-decrement matches. Visual inspection of the plots in figures 5.5 and 5.7 suggests that the subset of data points which do not graze the border of the white ‘increment-decrement’ rectangles represent smaller induction effects. Furthermore, this subset of data points can in many cases be fairly well described by a line through the origin, which would be consistent with a simple von Kries mechanism. This idea is illustrated by the exemplary plot in figure 5.9. The open symbols represent data points grazing the

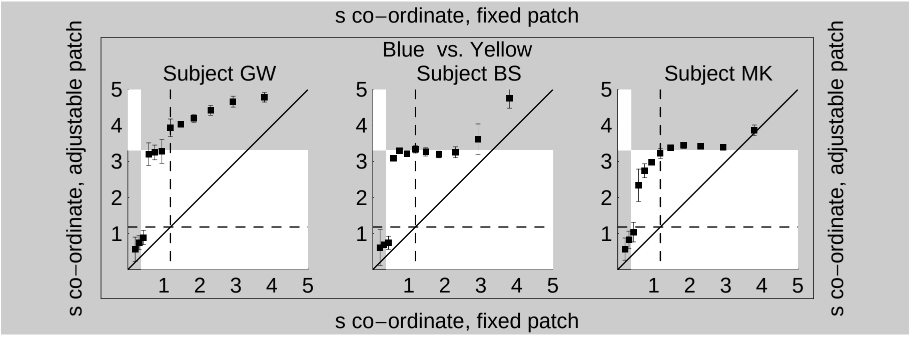

Figure 5.6: Individual plots for the surround pair $(B3, Y3)$ for the observers GW (left), BS (middle) and MK (right). Each data point is based on 3 individual settings. See caption of figure 5.5 for further explanations.

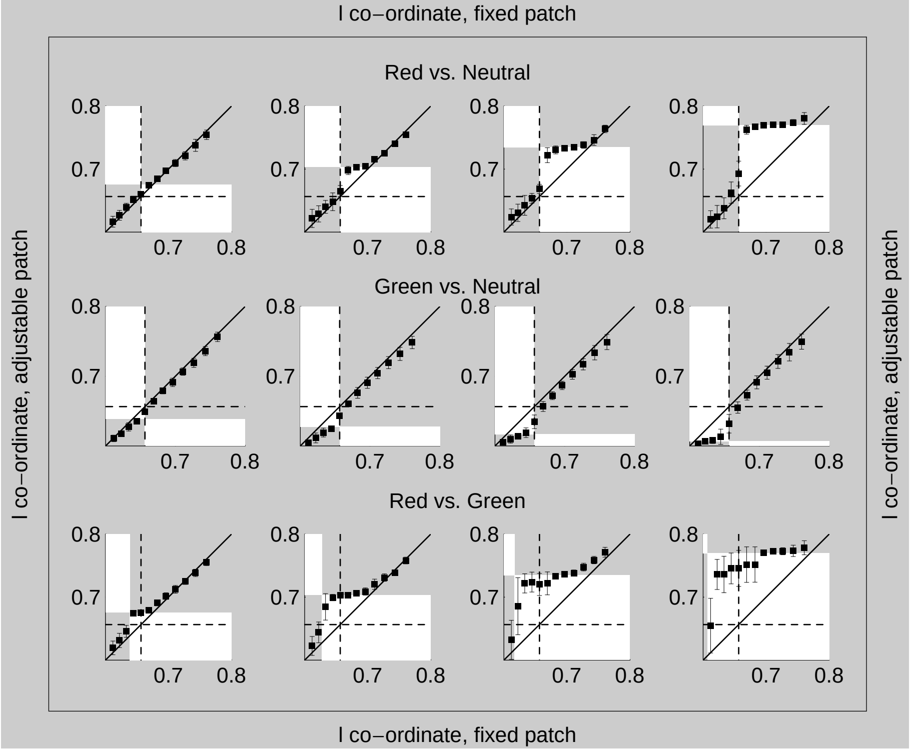

Figure 5.7: Results from experiment 4b. The plots show the results for the **l** co-ordinate. The caption of fig. 5.5 applies analogously by substituting ‘**s** co-ordinate’ by ‘**l** co-ordinate’.

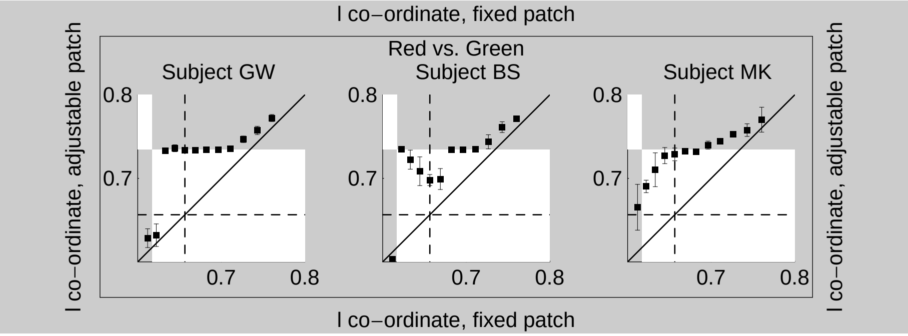

Figure 5.8: Individual plots for the surround pair $(R3, G3)$ for the observers GW (left), BS (middle) and MK (right). Each data point is based on 3 individual settings. See caption of figure 5.7 for further explanations.

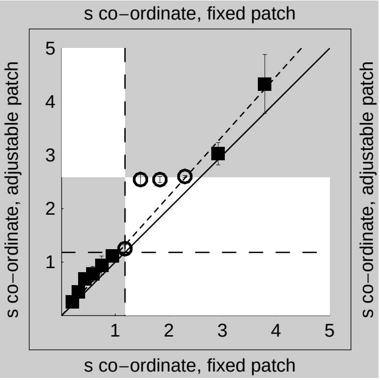

Figure 5.9: If those matches which graze the border of the white rectangle representing increment-decrement matches (open symbols) are disregarded, the remaining data points (filled symbols) can be well accounted for by simple von Kries scaling (dashed line). Data for surround pair (C, B_2), observer BS.

white increment-decrement border rectangle. The rest of the data points (filled symbols) are well fitted by a straight (dashed) line through the origin, which is consistent with von Kries adaptation. Based on the phenomenological and quantitative differences between these two subsets of the data we surmised that different perceptual mechanisms might be responsible for different subsets of the matching data.

5.2 Experiment 5

As reported in the previous chapter, simple centre-surround stimuli tend to evoke the perception of perceptual transparency. In the introduction to the present chapter, we ventured the hypothesis that the resulting dual colour impressions should make it difficult to establish true asymmetric colour matches. The results of experiment 4 indicate that this is indeed the case: The characteristic ‘step’ in the data curves appear to be related to subjective matching problems, which in turn seem to be related to the perception of transparency. These observations suggest that stimuli with uniform surrounds, which are indeed very simple in terms of a *physical description*, are far from simple in terms of the *internal semantics* of the visual system (Mausfeld, 2003a). Considering that—as noted by (Evans, 1974, p. 210)—spatially uniform surfaces come “close to being contrary to the

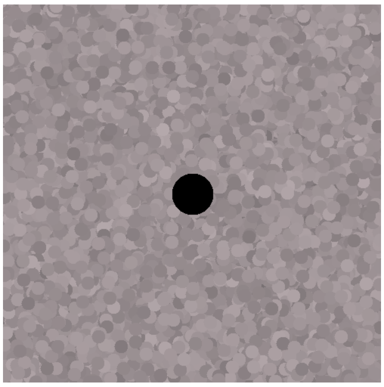

Figure 5.10: An achromatic Seurat surround (Andres, 1997; Mausfeld, 1998; Mausfeld & Andres, 2002). The figure represents the geometry of the actual surrounds used in our experiments correctly, but the Seurats actually employed in our experiments were isoluminant with purely chromatic variance.

laws of nature", this is not entirely unexpected. Indeed, natural scenes seldom give rise to uniform retinal stimulation. One rare exception, however, is the case of dense fog. Hence one may surmise that an extended uniform stimulus may serve as a fairly reliable cue for fog, which is a translucent medium. Indeed, a low variance of colour codes in the stimulus has been proposed by Brown and MacLeod (1997) as a cue for the presence of haze or fog. Since a uniform stimulus may be regarded as the limiting case of low variance, one could surmise that it also gives rise to this interpretation. If it is really the low variance that triggers an interpretation in terms of fog or haze, then, conversely, a high variance surround should make such an interpretation improbable, and any influence that may result from such an interpretation would be missing. Hence, by comparing the induction effects appearing in uniform surrounds with those appearing in variegated surrounds of high variance it should be possible to isolate the 'transparency effect', provided that both types of stimuli are comparable with respect to other possible mechanisms. Seurat-type stimuli (see Figure 5.10), previously investigated in a number of studies (Andres, 1997; Golz & MacLeod, 2002; Mausfeld & Andres, 2002; Webster, Malkoc, Bilson, & Webster, 2002), appear to be well-suited as such a 'base-line stimulus': For any uniform surround, a high-variance Seurat surround can be found which is functionally equivalent with respect to visual mechanisms which adapt to the spatial mean of the distribution of cone excitations. In accordance with these deliberations we decided to compare the induction effects appearing in uniform surrounds with those appearing in corresponding high-variance Seurat surrounds having the same spatial average of cone-excitations.

Stimuli and procedure Display apparatus and calibration technique were the same as in experiment 4. The CIE 1931 2° XYZ measurements made with a colourimeter were however converted to the presumably more realistic cone excitation values based on the Stockman-MacLeod-Johnson (1993) 2° fundamentals according to a procedure recently proposed by Golz and MacLeod (2003). Specifically, our XYZ measurements were multiplied with the matrix

$$M := \begin{bmatrix} 0.18772 & 0.60445 & -0.02517 \\ -0.14014 & 0.43056 & 0.03773 \\ 0.02017 & -0.04189 & 1.08472 \end{bmatrix}$$

in order to obtain LMS cone excitation values. The MacLeod-Boynton (1979) chromaticity co-ordinates given in this section are based on these values, i.e. $\mathbf{l} := L/(L{+}M)$ and $\mathbf{s} := S/(L{+}M)$. Luminance was defined as $L{+}M$. Again, the **s** co-ordinate is scaled to be 1 for equal energy white.

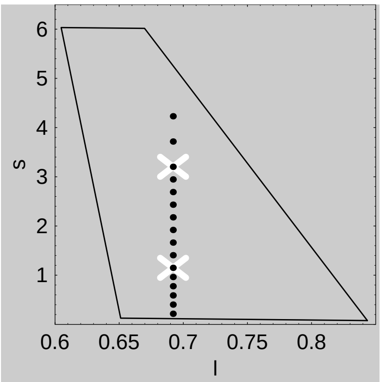

Figure 5.11: Chromaticities used in experiment 5 for the two uniform surrounds (white X's) and the fixed patch (black dots), plotted in the MacLeod-Boynton chromaticity diagram. The Seurat surrounds had the same *mean* chromaticities as the corresponding uniform surrounds.

The subjects viewed two circular patches (subtending 1.4° visual angle), each centred in the middle of a square surround (9.8°). The centre-to-centre distance between the two centre-surround stimuli was 11.7°. The surrounds were either uniform or variegated (Seurats); only surrounds of the same type were used together in a given stimulus presentation. All surfaces except the dark general background were equiluminant at $L + M = 10.97$.[3] The luminance of the test patch was also restricted to $L + M = 10.97$, but the chromaticity could be varied freely within the gamut of the monitor (see fig. 5.11) by using the arrow keys of a keyboard. For the two surrounds and the fixed patch we used combinations of chromaticities which were all located on the same cardinal axis (Krauskopf et al., 1982). Since our main objective in this experiment was to compare the effects of uniform and variegated surrounds, we only used chromaticities from the 'blue-yellow' cardinal axis (with a constant **l** co-ordinate at 0.692) for two surrounds and the fixed patch, thus keeping the experimental effort within reasonable bounds.

The MacLeod-Boynton chromaticities of the two uniform surrounds used were $(0.692, 3.203)$ and $(0.692, 1.149)$, which appeared approximately violet and white, respectively. In order to avoid cumbersome language, violet and its complementary, a yellowish-green, will henceforth be referred to as 'blue' and 'yellow', respectively. To each of the uniform surrounds, a corresponding variegated surround was constructed which had the same spatial average of MacLeod-Boynton co-ordinates. For both variegated surrounds, the variance of the **l**, **s** and luminance co-ordinates were 0.000484, 0.256036 and 0 respectively (with zero covariance between **l** and **s**). Since parts of the surround having a common contour with the central patch may potentially be more important in determining its perceived colour than more remote parts of the surround, we ensured that the spatial mean and variance of chromaticity co-ordinates was equal for several narrow ring-shaped regions around the central patch.

The chromaticities of the 16 fixed patches are plotted together with the chromaticities of the surrounds in figure 5.11. All had the same **l** co-ordinate as the surrounds, and the **s** co-ordinates ranged from 0.216 to 4.230.

On a standard account it should not matter which member of a given pair of surrounds is chosen for presenting the fixed patch and which of them is used as a surround for the adjustable patch. However, the characteristic 'step' evident in the data from experiment 4 suggests that presenting the fixed patch in only one of the surrounds yields an incomplete picture. Hence, for the pair of uniform surrounds employed, both surrounds were used for presenting the fixed patch, yielding a

---

[3] For CIE-illuminant C this value of $L + M$ corresponds to a luminance of 10 cd/m^2.

double, complementary data set. Preliminary experiments showed that when a pair of variegated surrounds were used, the data curves were much less complex, so that the use of only one of the surrounds for presenting the fixed patch yielded a sufficiently clear picture.

The 5 subjects were colour normal as assessed by the Ishihara Tests for Colour-Blindness and naïve as to the purposes of the experiments. All of the subjects were instructed to set the chromaticity of the adjustable patch such that it appeared as similar in colour to the fixed patch as possible, and in order to avoid possible misunderstandings, they were also told that if achieving the best possible match should require setting the chromaticity of the adjustable patch equal to that of its surround, they should not hesitate to do so merely because the central patch would then become indistinguishable from the surround.

As mentioned in the discussion of the results of experiment 4, true colour matches are sometimes impossible to make, and in these cases subjects seemingly revert to different strategies (A or B). In order to elucidate this point, two of the subjects (MH and GH) were told that if they experienced problems in establishing a perfect match and were uncertain how to proceed, they should at least satisfy the criterion that the two central patches should not contain traces of complementary colours, i.e. if one patch contained a shade of blue, the other should not contain a shade of yellow, and vice versa. Under the hypothesis that perfect colour matches can be established, this auxiliary instruction should of course not have any effect on the behaviour of the subjects. Based on the results of experiment 4, however, we expected that subjective problems should occur and that when confronted with them, this instruction would make the subjects pursue strategy A.

Results The results for those 3 subjects (AD, OS and SZ) who *did not* receive the auxiliary instruction are plotted in figure 5.12, those of the 2 subjects (MH and TG) who did, in figure 5.13. Each row in the figures represents the data of one subject. In the leftmost panels, the matching data for the pair of variegated surrounds are shown. For each subject, these data were fitted by a line through the origin (dotted line). These fits describe the data rather well, with values of R^2 between 0.94 and 0.97 for the 5 subjects. Furthermore, the individual slopes are rather similar, ranging from 1.21 to 1.26.

The simplicity of these data curves are contrasted by the more complicated data curves obtained with the corresponding pair of uniform surrounds (middle and right hand panels). The filled symbols in the middle panels represent the data obtained when the fixed patch was presented in the ’blue’ surround, and those in the right-hand panels represent the data obtained when it was presented in the ’neutral’ surround. To ease comparison with the data from the variegated surrounds, the dotted regression lines of the left-hand panels are redrawn in these panels. Data points which fall within the white regions in the plots represent a match between an **s** co-ordinate (i.e, S-cone) increment and an **s** co-ordinate decrement. The data for the uniform surrounds (middle and right-hand plots) are in general rather similar to the data for the variegated surrounds (left-hand plot) except that the white increment-decrement regions appear to be avoided, resulting in a ‘step’ in the data curve: the data curves appear to ‘get stuck’ at the left hand (middle panels) and upper (right-hand panels) borders of the (lower) white rectangle. This feature of the data is most pronounced in the data of the subjects who received the auxiliary instruction (figure 5.13). In the right hand panels of figure 5.13, none of the data points are within the increment-decrement region, and in the middle panels only one data point for each observer is clearly within, namely the data point immediately below the horizonal border between the two white rectangles. A possible reason for this exception is that the contrast between the fixed patch and its surround was so low that the observers occasionally did not recognize the central patch at all, and instead perceived a uniform field of the same colour as the surround. In these cases, one would expect a setting closer to the diagonal in the plots. The high variances of the data points in question support this interpretation.

The ‘step’, i.e. the vertical and horizontal portion of the data curves obtained with uniform

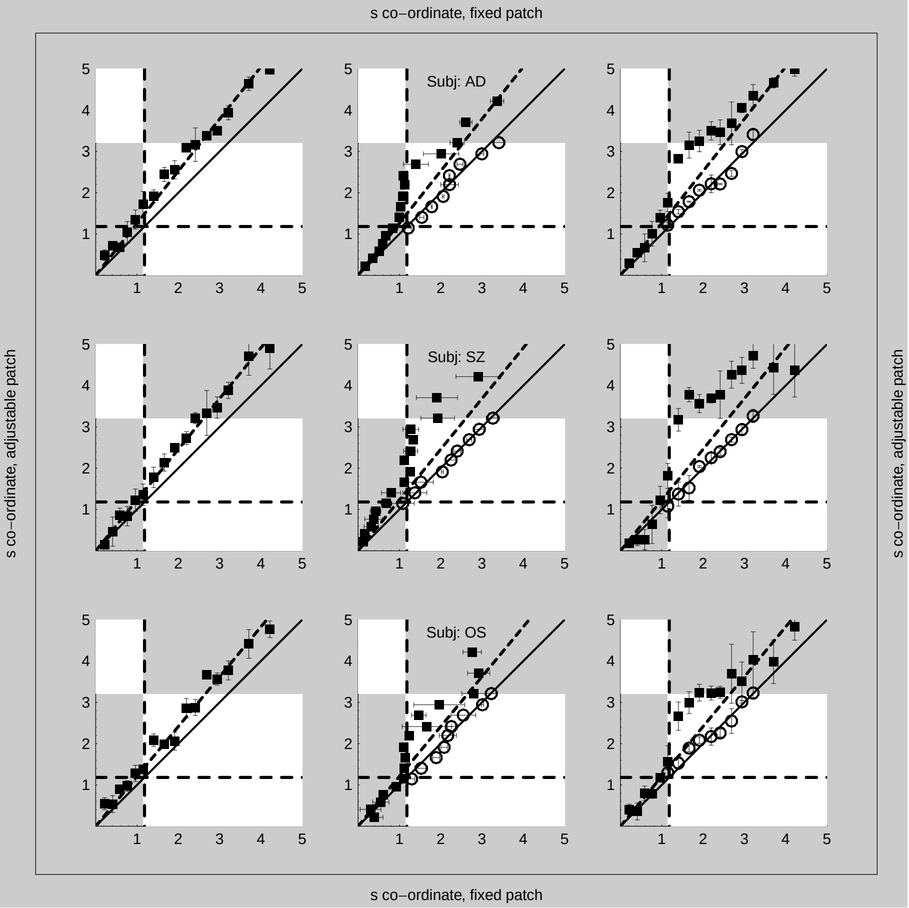

Figure 5.12: Results from experiment 5 for the subjects AD (top), SZ (middle) and OS (bottom). In the left and right panel the horizontal axis represents the **s** co-ordinate of the fixed patch and the vertical axis the **s** co-ordinate of the settings for the adjustable patch. In the middle panel the axes are interchanged. Each plot point represents the mean based on three individual settings, error bars represent one standard deviation in each direction. As in figure 5.5, the point where the white rectangles meet represents the **s** co-ordinates of the surrounds and the white rectangles the region where increment-decrement matches would be located. Left panels: Results obtained with *variegated* surrounds are well fitted by von Kries scaling (dashed line). Middle and right-hand panels: Results obtained with corresponding *uniform* surrounds show a characteristic 'step' which is associated with the avoidance of the white region which represents increment-decrement matches. Outside of this region, the matches fall close to the dotted line representing the von Kries scaling obtained with the variegated surrounds. The middle panels represent the data obtained when the fixed patch was presented in the neutral surround, the right-hand panel represents the data obtained when it was presented in the blue surround. Open symbols represent *symmetric* matches with a pair of identical neutral surrounds (middle panels) or with a pair of identical blue surrounds (right-hand panels).

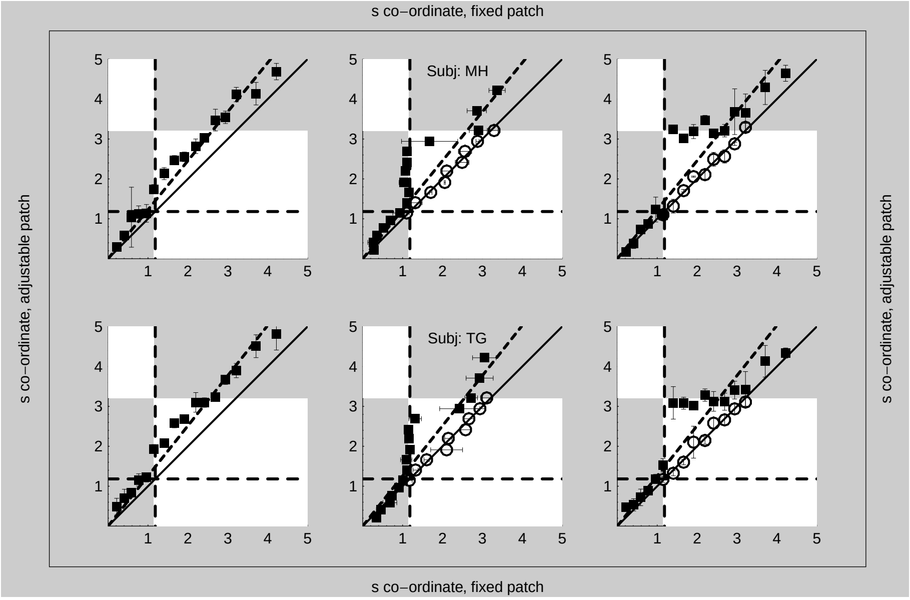

Figure 5.13: Results from experiment 5 for the subjects MH (top) and TG (bottom). See caption of figure 5.12 for further explanations.

surrounds shown in figures 5.12 and 5.13, means that several different chromaticities of the fixed patch were matched by the same chromaticity of the adjustable patch. This finding is *a priori* open to two interpretations. Either the asymmetric matches made by the subjects do not fulfill the requirements of an equivalence relation, which means that they cannot be taken to reflect perceptual identity, i.e. they cannot be true matches, or, alternatively, the physically different standard chromaticities which were matched to the same test chromaticity were perceptually indistinguishable.[4]

The open symbols in the middle and right hand plots of figure 5.12 and 5.13 represent *symmetric* matches with either the blue uniform surround used on both sides (middle panels) or the neutral uniform surround used on both sides (right panel). Since a pair of identical surrounds is used in both cases these matches should be colourimetric and hence fall on the diagonal of the plots, which is indeed the case. For our purposes, the interesting feature of these data is however not their veridicality, but their precision. The fact that these data points are well-ordered and monotonic indicates that the different fixed patches are readily discriminable. Since the set of chromaticities used for the fixed patch in these symmetric matches include those that were matched to the same chromaticity of the adjustable patch when the other uniform surround was used, we can conclude that the original asymmetric matches do not fulfill the requirements of an equivalence relation, and thus cannot, in general, be considered true colour matches.

5.3 Discussion

The results of experiment 5 suggest that there are fundamental differences between the simultaneous contrast effects evoked by variegated surrounds and those evoked by uniform surrounds.

---

[4] Let $(a, A) \equiv (b, B)$ denote the empirical statement that the central patch a in the surround A is matched by the central patch b in the surround B. Formally, our empirical finding means that there are two physically different central patches a and a' such that $(a, A) \equiv (b, B)$ and $(a', A) \equiv (b, B)$. If $\equiv$ is an equivalence relation, it follows that $(a, A) \equiv (a', A)$. Obviously, if this is not the case, $\equiv$ is not an equivalence relation and hence cannot be taken to represent perceptual identity.

In the case of variegated surrounds, a rather simple picture emerges: Firstly, the data curves are simple and consistent with von Kries scaling, and secondly, neither impressions of transparency nor subjective matching problems were reported. In the case of uniform surrounds, however, matters seem to be much more complicated, and we shall devote some space here to discuss how the phenomena occuring in uniform surrounds can be described.

5.3.1 How uniform surrounds are special

As reported in the previous chapter, equiluminant centre-surround stimuli with low chromatic contrast between centre and surround evoke impressions of transparency in which two simultaneously perceived colour components can be distinguished, namely a contrast component and a background component. The background component always appears in the same colour as the surround, whereas the perceived colour of the contrast component depends on the contrast between centre and surround. If the surround is achromatic, then the background component is also achromatic. This is illustrated in the left panel of Colour Plate XIII, where the background component is visible as a kind of achromatic haze, through which the 'spokes' are perceived. The colours of the spokes themselves is what we refer to as the contrast component. Note that the hazy background component is only perceived in the region close to the centre of the spokes, where contrast with the achromatic surround is low. At larger contrast, no perceptual transparency is evident; The regions of the spokes further away from the centre do not appear to be covered by any transparent medium. In the right-hand panel of the Colour Plate, the same set of spokes is embedded in a variegated surround. In this case, one has no impression of transparent haze. Comparing the colour appearance of the two sets of spokes in the left and right panels, respectively, one can make a very noteworthy observation. When the spokes are embedded in the variegated surround, the colours of the spokes grow increasingly similar to gray towards the centre. In the uniform surround, however, the same increase in gray content is attributed to the hazy gray background component only. The spokes themselves, as perceived 'through' the grayish haze, do not appear to vary in grayishness: They appear rather, to be devoid of any grayishness altogether. Thus, although the spokes may be said to vary in colourfulness, they do not vary in saturation.[5] As the colours of the spokes become more similar to the surround, they become less colourful or less pronounced but *not less saturated*. The contrast component cannot be made arbitrarily similar to white, hence the range of contrast component colour impressions that can be evoked in a uniform surround does not include colour impressions of low saturation. Compared to the range of colour impressions which can be evoked in the variegated surround, we may say that colour impressions of low saturation are missing for all hues.

Now, if a uniform surround is coloured instead of achromatic, the situation is slightly different. Also in this case certain colour impressions can be said to be missing, but they are not exactly the same as the ones that are missing in the case of an achromatic surround. Compared to the set of contrast component colour impressions that can be realised in a grey uniform surround, the ones that can be realised in a blue uniform surround may be said to, on the one hand, include *more desaturated yellows*, but, on the other hand, even *less desaturated blues*. Analogous statements can be made with respect to any arbitrarily coloured uniform surround. In a green surround, for instance, *a greater set of low-saturation reddish colour impressions* and *a smaller set of low-saturation greenish ones* will be available than in a grey surround. These observations are

---

[5] Wyszecki and Stiles (1982, p. 487) define *saturation* as "the attribute of a visual sensation which permits a judgement to be made of the degreee to which a chromatic stimulus differs from an achromatic stimulus regardless of their brightness" and *colourfulness* (= chromaticness) as "the attribute of a visual sensation according to which the (perceived) colour of an area appears to be more of less chromatic." Note that the former definition makes reference to an achromatic sensation, i.e. grey (or white) content, whereas the latter does not.

demonstrated in Colour Plates XIV and XV on pages 145 and 146 . In short, there is always a set of colour impressions that can be evoked in a neutral uniform surround but not in a coloured uniform surround, and vice versa. Clearly, this statement can be generalised for any pair of surrounds of different hue. It is thus clear that true asymmetric colour matches should sometimes be impossible to make. An important point which needs to be made here, is that there seems to be not only one reason why asymmetric colour matching should be difficult when uniform surrounds are used. Based on the analysis of the experiments in the previous chapter, we assumed that matching problems should occur because differently coloured surrounds lead to the perception of differently coloured background components perceived as transparent haze covering the targets. Such problems were indeed experienced in the present experiments, but the above observations demonstrate that even if one pays attention only to the contrast component, asymmetric colour matches are difficult because the set of contrast components which can be perceived depends on the colour of the surround. In the following, we shall consider a descriptive model of the latter phenomenon which may account for the curious ‘step’ in the data curves of our experiments. Afterwards, in section 5.4, we shall see that this model makes rather good predictions for other experimental data as well.

5.3.2 Saturation scale extension and truncation

The descriptive model to be developed in the following is based on the above observation that the set of colour impressions that can be evoked in a uniform surround depends on the surround colour. The model refers not as much to what colours are actually perceived as to what aspects of the total colour impressions subjects seem to rely on when making their settings. In cases where no transparency is perceived, this obviously requires no further specification. In cases where transparency is perceived, though, both a background component and a contrast component is perceived. Nevertheless, as already noted in the previous chapter, the subjects seem to base their settings predominantly on the contrast component. Accordingly, in cases involving dual colour impressions, the descriptive model refers only to the contrast component and neglects the background component.

Now, consider—in the equiluminance plane—a line of chromaticities through equal energy white, ranging from, say, blue to yellow. We then have a scale of colour impressions which is divided into a saturation scale of yellows and a saturation scale of blues at the chromaticity of the neutral point. In a graphic representation of the colour impressions on this line, we may think of each of these two half-axes as a wedge, whereby the thickness of this wedge represents the saturation of the colour impression represented. Thus we obtain a blue wedge and a yellow wedge which abut at the neutral point (see Colour Plate XVI on page 147, left middle panel). It is obvious from this representation that the line of chromaticities in question allows for colour impressions of arbitrarily low saturation for both hues, in accordance with classical assumptions. In our descriptive model, however, we depart from this assumption, and posit that it is not possible to produce colour impressions of arbitrarily low saturation when the patch is embedded in an equiluminant uniform neutral surround. This feature of our descriptive model, which we shall refer to as ‘pretruncation’, is accommodated in the graphic representation by using wedges which have their tips cut off (see Colour Plate XVI, right middle panel). We further posit that when a coloured surround is used, the wedges representing the colour impressions for the central patches remain essentially *unchanged*. The graphic representation is merely modified in the following way: the wedge which represents colours of the same hue as the surround is truncated at the location of the surround colour, and the wedge representing colours of the complementary colour impression is correspondingly extended (see Colour Plate XVI, right panel). This feature of the descriptive model will be referred to as saturation scale extension and truncation.

Features of the data captured by the notions of extension and truncation From our descriptive model we may indeed derive expected data curves for asymmetric matches that closely resemble the actual data. This is illustrated for asymmetric matches between a neutral and a blue surround in the upper left panel of Colour Plate XVII on page 148. In this case it is assumed that the fixed patches are presented in the neutral surround. As in our data plots, the axes in these figures represent the chromaticity of the fixed and the adjustable patch, respectively. Along the horizontal axis a schematic representation of expected perceived colour for these chromaticities is drawn, given that the central patch is presented in a neutral surround. Along the vertical axis a corresponding representation for expected perceived colour for central patches presented in a blue surround is shown. Any point in this 'plot' which projects to locations on the two axes representing the same colour impression, as indicated by the schematic saturation scales drawn along them, represents an expected match. The tips of the *solid* vertical arrows represent such expected matches. It will be noted, however, that the scale drawn at the horizontal axis contains blues of low saturation which have no *identical* counterpart in the scale drawn at the vertical axis. For these colours a true match should not be possible. The tips of the *dotted* vertical arrows drawn in the figure represent the best possible settings that can be expected under the assumption that any two colour impressions which have the same hue are more similar than two complementary colour impressions.[6] If the same pair of surrounds is used, but the fixed patch is presented in the blue surround (Colour Plate XVII, lower left panel), an analogous construction can be made. In this case we observe that the scale drawn at the vertical axis contains yellows of low saturation which have no identical counterpart in the scale drawn at the horizontal axis. The tips of the horizontal arrows represent the expected matches, whereby those arrows which are drawn as dotted lines again represent imperfect, but best possible matches. In the right-hand panels of Colour Plate XVII analogous expected matching curves have been derived for a blue and a yellow surround, exhibiting the same general features.

It should now be clear that the somewhat unorthodox notions of saturation scale truncation and extension, which are based on the observation that some low-saturation colour impressions do not appear in certain surrounds, allow for an elegant and parsimonious description of a very prominent feature of our results, namely the 'step' in the data curves which arises because several different chromaticities of the fixed patch are matched to the same chromaticity of the adjustable patch. An obvious but unimportant and easily remediable shortcoming of this description is that it predicts colourimetric (i.e. 'veridical') matches outside the region of the 'step': A closer correspondence to this subset of the actual data can be achieved by including a simple von Kries scaling operation prior to truncation and extension. In the above we refrained from doing so merely for reasons of expository simplicity. Furthermore, the amount of von Kries scaling necessary to account for the actual data points in our plots would be rather small.

5.4 Experiment 6

Though the descriptive 'saturation scale truncation and extension' model developed in the previous section provides a simple and principled account of both phenomenal observations and the quantitative pattern of results, it is, admittedly, both provisional and *ad hoc*. On the other hand, however, it is also sufficiently specific to derive expectations for further experiments. In this section, we consider a phenomenon originally studied by Brown and MacLeod (1997): A set of differently coloured targets appear more saturated when they are embedded in a uniform grey surround than

---

[6] As discussed in the results section of experiment 4, subjects may pursue two different strategies (A or B) when confronted with this lack of a true match: either ensure that hue is correct, or that saturation differences are minimised. The derived expectations are based on the assumption that strategy A is pursued.

when they are embedded in a variegated surround with the same average colour. This effect, for which Brown and MacLeod coined the term 'gamut expansion' is demonstrated in Colour Plate XVIII on page 148. In order to measure the effect quantitatively, Brown and MacLeod (1997) performed an asymmetric colour matching experiment: Target patches were presented in the uniform surround, and the observers adjusted the purity of corresponding patches presented in the variegated surround in order to make them appear equal to the targets. For this experimental situation, we can make the predictions illustrated in the left panel of Colour Plate XIX on page 149. For concreteness, the predictions are derived for target patches along the red-green axis in chromaticity space. The horizontal and vertical axes represent the same scale of chromaticities along the red-green axis for targets embedded in the uniform and the variegated surround, respectively. Alongside the axes, wedges representing the colour impressions we expect to occur are drawn: Only in the case of the uniform surround, pretruncation is expected to occur, whereas in the variegated surround, colour impressions of arbitrarily low saturations should be realisable. A perceptual match should plot at the points where the horizontal and vertical wedges have the same thickness (representing saturation), and these points are given by the heavy black lines. The dashed diagonal line represents the identity $y = x$, hence the vertical distance between the dashed and the heavy lines can be interpreted as a measure of the expected induction effect. Consequently, the illustration shows that the expected induction effects are largest for target patches with low purity, and, as the target purity increases, approach zero. Brown and MacLeod (1997) used target patches of low purity, and did indeed find rather large induction effects. Since they did not investigate target patches of higher purities, however, the data they reported do not allow further evaluation of the above predictions. In this section, I describe an experiment yielding data in good accordance with the above expectation. This experiment was conducted by Gunnar Wendt in collaboration with Franz Faul and myself (Wendt, 2003; Wendt, Faul, & Ekroll, submitted).

Stimuli and procedure The stimuli were displayed on a CRT computer monitor. The subjects viewed two centre-surround configurations with a horizontal centre-to-centre distance of $9.49°$. The circular target patches had a radius of $0.516°$ and the square surrounds had a width of $7.64°$. One of the surrounds was uniform with the chromaticity of CIE Illuminant C (MacLeod-Boynton coordinates $l = 0.692$, $s = 1.149$ based on the Stockman et al., 1993, 2°-cone fundamentals) and a luminance $L + M = 10$ (corresponding to a value of Y = 9.12 cd/m^2 in the CIE 1931 system). The other, variegated surround always had the same spatial mean as the uniform surround and the fixed covariance matrix

$$Cov = \begin{pmatrix} 4.4283 & 1.9361 & 0 \\ 1.9361 & 0.9401 & 0 \\ 0 & 0 & 10.24 \end{pmatrix}$$

of LMS-coordinates. This covariance matrix describes the largest ellipsoidal distribution in colour space realisable within the monitor gamut for the given mean. In each trial, the colour of the central patch in the uniform surround was fixed and the task of the subject was to adjust the luminance and purity of the central patch in the variegated surround such that both patches appeared as similar as possible. In order to ease the task of the subjects the settings were restricted to the line in chromaticity space through Illuminant C and the chromaticity of the fixed patch. This restriction, which was also made in Brown and MacLeod's study, was found to be unproblematic in a pilot experiment (Wendt, 2003) in which unrestricted settings showed no significant deviation from this "constant hue" line. The MacLeod-Boynton chromaticity coordinates of the fixed patches used are shown in Fig. 5.14. Their luminance was always 9.12 cd/m^2, i.e. they were equiluminant to the surround. Each of the 24 different test colours was presented 8 times, resulting in 192 trials which were presented in random order. The assignment of the uniform and variegated stimulus to the left or right hand side of the monitor was balanced in order to avoid any bias due to slight

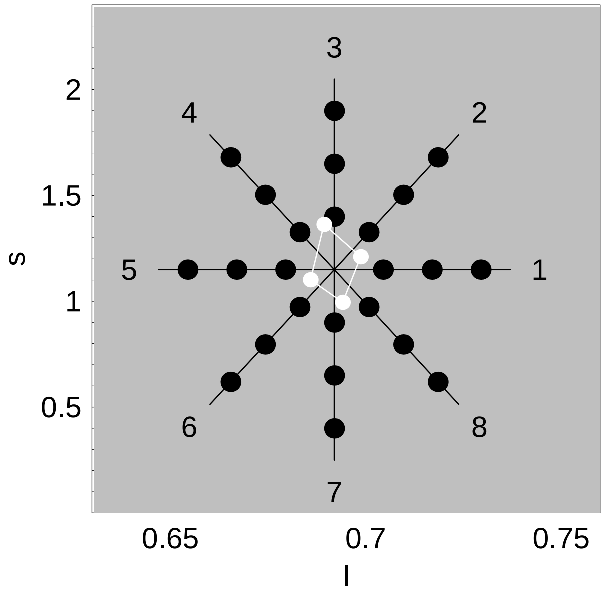

Figure 5.14: MacLeod-Boynton chromaticities used for the fixed patch in experiment 6 (black dots). On each hue direction, numbered from 1 through 8, 3 different purity levels were investigated. For comparison, the chromaticities used by Brown and MacLeod (1997) are shown as white dots.

monitor inhomogeneities. The computer monitor was colourimetrically calibrated by means of a colourimeter (LMT C1210) following a standard procedure (Brainard, 1989). The monitor was controlled by a graphics card with a colour resolution of 8 bits per channel. The viewing distance was 100 cm. The monitor was the only source of illumination in an otherwise dark room. Four naïve subjects and one of the investigators (GW) performed the experiment. All had normal colour vision according to the Ishihara Tests for Colour-Blindness (Ishihara, 1967).

Results The mean results, pooled across all subjects, are shown in Figure 5.15. In each panel, the data for one pair of complementary hue directions are shown, whereby the chromaticity of the neutral surround is chosen as origin of the coordinate system, i.e. $\Delta s := s - s_c$, $\Delta l := l - l_c$, where s_c, l_c are the chromaticity coordinates of the surround (Illuminant C). In these plots the effect corresponds to the difference between the data points and the diagonal line as indicated by the grey region. With increasing purity (corresponding to the absolute values of Δs or Δl), the data points approach the diagonal. It appears natural to regard the diagonal as an asymptote since a reversal of the effect with increasing purity seems highly improbable. The asymptotic behaviour of the data points are well described by a function $y = a\exp(-xb) + x$, where x and y stand for the purity of the fixed and adjustable patch, respectively. The grey regions are the area between the diagonal and the curve fitted to the data points. The general shape of the data curves is rather similar to the expectations based by the notion of saturation scale truncation (see Colour Plate XIX on page 149, left): The induction effect is largest for targets of low purity, and approaches zero for larger purities. This is the opposite of what one would expect based on a simple model of contrast gain control (Colour Plate XIX, right). According to this model, the size of the induction effect should increase, not decrease, with purity. In Figure 5.15, the best fit to the contrast gain control model is indicated by the dotted lines (unequal gain factors for increments and decrements have been permitted, hence the changes in slope at the origin). The predictions of the truncation model are clearly in much better accordance with the data.

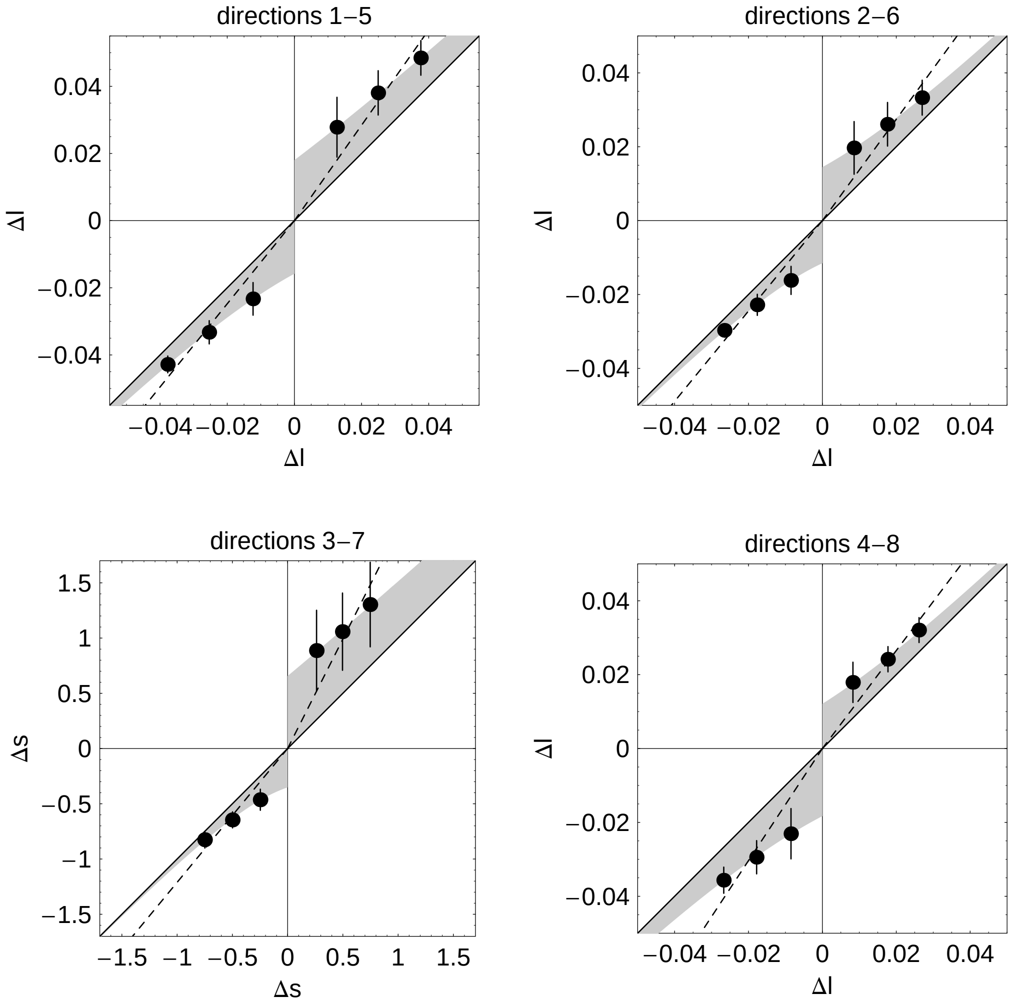

Figure 5.15: Dependence of the gamut expansion effect on the purity level. Each panel shows the mean settings of all subjects for two complementary hue directions (see Fig. 5.14) in MacLeod-Boynton coordinates, relative to the coordinates of Illuminant C. Hence, $|\Delta l|$ and $|\Delta s|$ are proportional to purity. For directions 2-6, and 4-8, for which Δs and Δl covary, only Δl is given. The error bars correspond to $\pm$ 1 SD. On the abscissa the relative coordinates of the fixed patch are shown, on the ordinate those of the subjects' settings. The grey regions show the gamut expansion effect (according to the fitting procedure described in the text). The dotted lines show the best-fitting predictions of the gain control model, which clearly fail to account for the data.

As shown in Colour Plate XIX, the truncation model predicts a linear decrease in the size of the effect with purity. This is merely to be understood as a provisional simplifying assumption, so it is not very surprising that the data suggest a slight nonlinearity. The essential point of the truncation model is that colour of low saturation cannot be evoked in the uniform surround, and that one therefore would expect a discontinuity, or step, in the data curve at the chromaticity of the surround.

As in the previous experiments, impressions of transparency were experienced when the contrast between target and surround was low. This observation was also made in the original experiments of Brown and MacLeod (1997, p. 848), who noted that “the appearance of low contrast displays often had the phenomenological quality of rich colors seen through a thick fog”. The reader may get an impression of this in Colour Plate XIII on page 144. This demonstration may be said to informally demonstrate the main findings of this experiment: At low contrast between target and surround, a) the induction effect is most pronounced, and b) impressions of transparent haze are evident. Also, when the spokes are embedded in the uniform surround they do not demonstrate the ‘natural’ impressions of desaturation evident when they are embedded in the variegated surround.

Chapter 6

General Discussion

As discussed in chapter 3, quantitative models of simultaneous contrast classically assume that the effect of the surround on the perceived colour of the target can be compensated by changing the cone excitation vector of the target. Consequently, the influence of the surround is usually modelled as a surround-dependent transformation of target cone excitation vectors. The experiments described in chapter 4, however, strongly suggested that this assumption is unwarranted. The basis for this conclusion was the seemingly paradoxical finding that experimentally determined lines of constant hue converged on the chromaticity of the coloured surround instead of on the chromaticity of a target which appears achromatic. Although this finding appears paradoxical when viewed from a classical perspective, it was shown that it can be rationally accounted for under the assumption that simple centre-surround stimuli evoke dual colour impressions reminiscent of perceptual transparency. A basic ingredient in this explanation was the notion that, at low contrast between target and surround, the observer perceives a transparent 'veil' or 'haze' covering the target patch, and that this transparent veil is perceived to have basically the same colour as the surround. Based on the principles of this explanation, which was also supported by direct phenomenological observations, the display shown in Colour Plate XI on page 143 was devised, which can be said to provide an informal demonstration of the basic observations.

A straighforward implication of the interpretation offered was that it should sometimes be impossible to make two targets embedded in differently coloured surrounds appear exactly equal in colour, which would mean that the classical psychophysical method of asymmetric colour matching should be problematic. This matter was investigated further in the experiments of chapter 5. We expected subjective matching problems to occur because targets embedded in differently coloured surrounds should appear to be covered by differently coloured transparent media. This was indeed noted by the subjects, but these experiments also revealed a second, more disturbing, source of subjective discontent with the matches. Even disregarding the differently coloured transparent 'veils', it was sometimes impossible to match the perceived colours of the remaining colour impression (the 'contrast component'). Based on the observations made in the experiments, it was concluded that this was because in a given uniform surround, a certain set colour impressions cannot be evoked by any target chromaticity whatsoever, and that this set of 'missing colours' is different for differently coloured surrounds. To describe these phenomena of 'missing colours', the notions of 'saturation scale extension' and 'truncation' were introduced (section 5.3.2). A simple descriptive model based on these notions was then shown to account naturally for the distinct 'step' in the data curves (section 5.3.2). The results of experiment 5 (section 5.2) suggest that both the occurence of impressions of transparency as well as the phenomenon of saturation scale extension and truncation (reflected in the distinct 'step' in the data curves) are particular to uniform surrounds. When comparable variegated surrounds are used instead, none of the above was observed, and the data were found to be compatible with simple von Kries scaling. It was noted that the more complex data obtained with uniform surrounds could be understood as a combination of this von Kries scaling and saturation scale extension and truncation, whereby the latter overrides the former.

Based on the observation that saturation scale extension and truncation only occured in the uniform surrounds, predictions were derived for the ‘gamut expansion’ effect originally studied by Brown and MacLeod (1997). The dependence of the gamut expansion effect on the purity of the target was found to be in good qualitative agreement with these expectations, which further demonstrates the usefulness of the notions of saturation scale extension and truncation.

The most important conclusion which can be drawn based on the present findings is that uniform surrounds evoke induction effects of a very peculiar nature, which are, contrary to widespread and seemingly innocuous assumptions, *not representative* of colour induction effects at large. This observation is of central importance, given the widespread use of uniform surrounds in studies of colour vision. Various aspects of the findings also have implication for a number of other current research issues. Before this is discussed in more detail, though, we shall briefly consider how the notions of saturation scale extension and truncation can account for an aspect of simultaneous colour contrast known as Meyer’s effect.

6.1 Accounting for Meyer’s effect

An effect related to simultaneous contrast which has received little attention in modern colour science, but was often discussed in the earlier literature is Meyer’s effect, which is also known as *tissue contrast* (German: ‘Florkontrast’, Meyer, 1855; Hering, 1887b; Helmholtz, 1911; Perls, 1932; Walls, 1960; Brown, 2003; Mausfeld, 2003a). The basic observation is that the inducing effect a coloured uniform surround has on a central patch which appears grey when viewed in isolation may appear equally or even more impressive when the centre-surround stimulus is viewed through a transparent tissue although this generally reduces the purity of the surround in terms of the proximal stimulus. In order to observe this effect it is not necessary to use a real transparent tissue; The mere simulation of a transparent medium will also work, as demonstrated in Colour Plate XX on page 150. The two centre-surround stimuli on the left constitute the typical demonstration of simultaneous colour contrast where two physically identical ‘grey’ targets are embedded in differently coloured surrounds. In the middle, this demonstration is partially covered by a simulated grey transparent layer[1]. The colour of the transparent layer was chosen to be identical to the colour of the targets; hence the transparent layer leaves the cone excitation values of the target patches unchanged, whereas it reduces the purity (saturation) of the coloured surrounds. In spite of this, the two targets do not appear less different than in the original demonstration on the left. The demonstration of the right-hand side shows that simply reducing the purity of the surrounds by the same amount as the transparent layer would do has basically the same effect.

The interesting feature of this demonstration is that it is at odds with the conventional wisdom that more saturated surrounds have a stronger inducing effect. Instead, when the purity of the surround is reduced, the colour induced into a nominally gray target seems to remain essentially the same. This is documented in a study by Kinney (1962). Her results are replotted in Figure 6.1. It can be seen that the strength of the induction effect, which is present already with surrounds of very low purity, remains remarkably stable in spite of large increases of surround purity. In fact, Kinney reported that for some of her subjects the induction effect even *decreased* with increasing surround purity.

Considering that experiments on simultaneous contrast, including the present ones, suggest, in accordance with contrast-coding models (Whittle, 2003), that changing the colour of the surround

---

[1] The simulation is based on Metelli’s (1970) model of additive transparency.

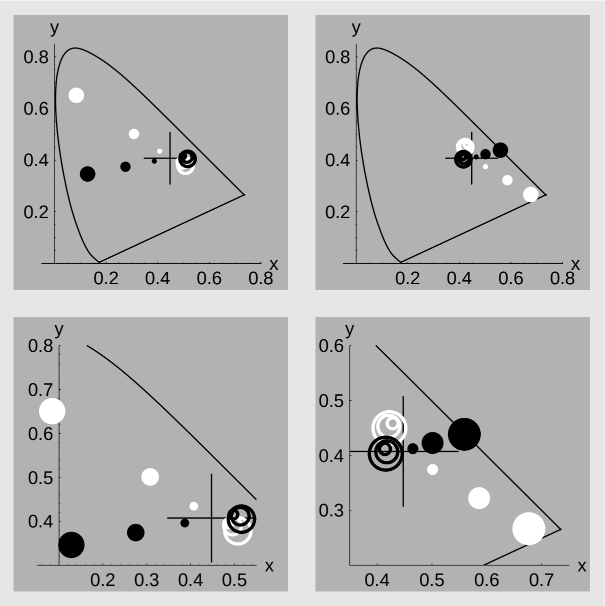

Figure 6.1: Matching data from Kinney (1962), replotted in the CIE 1931 **xy**-chromaticity diagram. A patch with the chromaticity of illuminant A (large cross) was presented in surrounds with different purity (filled circles). Subjects adjusted the colour of an isolated patch in order to establish a match with the embedded patch. The mean results from four observers are shown as open circles. Correspondence between surround condition and subjects' settings is denoted by size and colour of the disks and circles, respectively. The lower graphs are blown-up versions of the upper ones. Obviously, the subjects' settings change but marginally with large changes in the purity of the surround.

can have quite dramatic influences on the perceived colour of the target, this observation may appear rather surprising. The saturation scale truncation and extension model introduced in section 5.3.2, however, accounts naturally for all of these observations. As shown in Figure 6.2 saturation scale truncation and extension predicts that the perceived colour of the central, nominally grey patch should appear the same yellow when presented in a blue surround, however saturated the surround is, and the same blue when presented in a yellow surround, whatever its saturation. In cases where the target is not nominally neutral, though, changing the purity of the surround may well have the dramatic effect of turning the original colour appearance of the target into its complementary (e.g. from yellow to blue).

6.2 The role of contrast-coding in colour vision

Traditionally, one has sought to understand colour vision primarily in terms of the absolute cone excitations produced by the local stimulus. Accordingly, context effects have been assumed to be due to changes in the sensitivity of the receptors (or postreceptoral mechanisms), brought about by changes in the stimulation of neighbouring regions of the retina. Viewed from this perspective, the absolute, local, colour signal plays a more fundamental role than the colour signals from surrounding regions, since the role of the latter is merely to modify the former (see chapter 3). In more modern treatments of colour vision, though, the colour signals from the surround are often assumed to play an equally fundamental role as those evoked by the target itself: Colour is coded in terms of the difference (or, contrast) between adjacent regions, and since the cone excitation vectors at both sides of an edge contribute equally to that difference, neither can be said to be

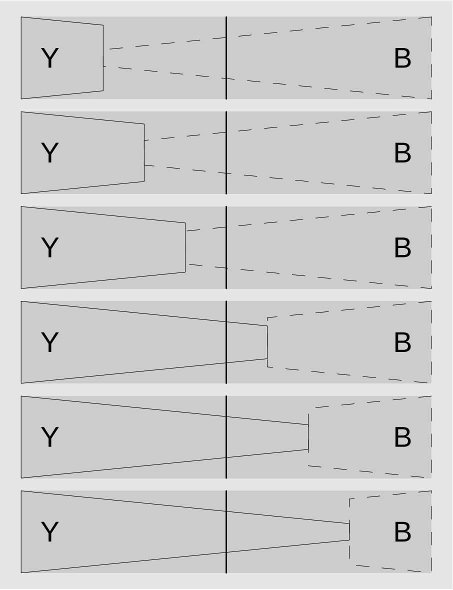

Figure 6.2: Prediction of Meyer's effect based on saturation scale extension and truncation. The three upper scales represent the expected colour impressions for central patches from the blue-yellow axis in chromaticity space when viewed in three yellow surrounds of different saturation. The three lower scales show the same for three blue surrounds. The vertical lines denote the chromaticity of a nominally grey central patch. As evident from the illustration, such a neutral patch should appear the same blue (yellow) when presented in any of the yellow (blue) surrounds.

more fundamental than the other.

The two abovementioned general hypotheses are often treated as though they were mutually exclusive possibilities, and depending upon which of the coding principles – absolute or contrast-based – is regarded as basic, very different theoretical significance is attributed to the relevant empirical observations. Some experiments and informal observations indicate strong effects of context on colour appearance, whereas in other cases perceived colour is found to be almost independent of context. If absolute coding is assumed to be basic, the former observations are in need of an explanation, whereas the latter is exactly what one would expect. If, on the other hand, contrast coding is regarded as basic, strong simultaneous contrast effects are in no need of further explanation. Also, the well-documented and biologically important phenomenon of 'colour constancy' (Gelb, 1929; Brainard, 2003), i.e. the approximate independence of perceived surface colour in spite of illuminant changes, is to be expected, since the contrast between neighbouring surfaces is largely invariant across illuminant changes. However, the cases in which colour perception is found to be largely independent of context poses a challenge to the contrast-coding perspective. A case in point is the observation that when an object moves such that the background against which it is seen, and consequently also the contrast at its border changes, its perceived colour does not appear to change appreciably (Whittle & Challands, 1969; Gilchrist et al., 1999; Whittle, 2003; Brown, 2003). Here we see that if the contrast-coding perspective is adopted, it becomes necessary to provide an explanation for observations which correspond exactly to what every layperson would naïvely expect. To the layperson the constant colour appearance of an object moving against a background of differently coloured objects is completely unspectacular, but for the scientist viewing it from the contrast-coding perspective it becomes a phenomenon interesting enough to reserve a special term for it, namely 'background-independent'-, or 'type 2'-

constancy (Gilchrist et al., 1999; Whittle, 2003). The qualifiers ‘background-independent’ or ‘type 2’ are used to avoid confusion with the classical phenomenon of ’illumination-independent’-, or ‘type 1’-colour constancy. Background-independent constancy is typically accounted for in terms of mecanisms of ‘integration’ which counteract the effects of contrast-coding, according to the idea that the “retina differentiates; the brain can integrate” (Whittle & Challands, 1969, p. 1109, see also Arend, 1973). On this view, understanding the putative mechanisms of integration becomes a primary aim of colour vision theory: Why does the brain sometimes integrate, and sometimes not? As discussed in section 3.3.6, it is simple to explain why integration should break down under the special experimental conditions of haploscopically superimposed displays (Whittle & Challands, 1969), but explaining why it sometimes also breaks down under ordinary, everyday viewing conditions is more tricky. A prominent theoretical notion is that the visual system may classify edges in the proximal stimulus as being due either to a change in reflectance, or to a change in illumination, and disregard the latter in the integration process (Gilchrist, Delman, & Jacobsen, 1983): This could be an ’intelligent’ strategy used by the visual system in order to achieve colour constancy.

In order to explain strong simultaneous contrast effects in terms of a ‘breakdown of integration’, it seems necessary to assume that the ‘illumination edge’ which is neglected in the integration process corresponds to another edge in the proximal stimulus than that between the target and its immediate surrounding. For, if the visual system were to neglect the edge between the target and its surround, it would neglect the very contrast signal which is assumed to be responsible for the strong simultaneous contrast effect. Hence, in order to account for the strong simultaneous contrast effect observed in a simple centre-surround stimulus presented in complete darkness, one is left with the option of assuming that the outer edge of the surround is neglected. Yet this does not seem to be the case, because then a surround of high purity should appear similar to the surrounding context of darkness: It generally does not[2], yet strong simultaneous contrast effects are observed, as demonstrated by the results of the experiments reported in chapter 4. Furthermore, even if one is nevertheless willing to uphold the assumption that a failure of integration at the outer border of the surround is responsible for the strong simultaneous contrast effect, one would encounter severe difficulties when trying to explain why the strength of the simultaneous contrast effect varies dramatically even when only the colour of the target is varied. This phenomenon is a clear feature of data of the asymmetric colour matching experiments with uniform surrounds reported in chapter 5, and a central property of the saturations scale extension and trunctation model used to described them. It is also demonstrated in Colour Plate XXI. Different degrees of integration cannot explain these dramatic variations in the strength of the simultaneous contrast effect because everything that can reasonbly be assumed to influence it is kept constant.

From the above, it is clear that it is difficult to account for our findings in terms of the ‘contrast-coding-plus-integration’ scheme. How, then, can we account for the fact that our experiments revealed both strong simultaneous contrast effects indicative of contrast coding, as well as very weak ones indicative of background-independent constancy? It would appear natural to give up the assumption that colour perception in general is *either* fundamentally contrast coded *or* fundamentally absolutely coded, and instead assume that the visual system uses both coding schemes from the very start. If we assume that the visual system also has an absolutely coded colour signal at its disposal, then there would be no need for a process of integration, and we could dispose of the ‘there-and-back again’ explanation of background-independent constancy: Instead of assuming that background-independent constancy is due to integration counteracting the putatively more basic contrast coding, we could simply assume that it is due to absolute colour coding.

A revealing aspect of the present findings is that in the cases where *uniform* surrounds were

---

[2] In the case of prolonged viewing, a strongly coloured surround may appear less saturated, but this is not necessary for strong simultaneous contrast effects to occur; the major part of the simultaneous contrast effects appearing in uniform surrounds occurs practically *instantaneously* (Rinner & Gegenfurtner, 2000).

used, very strong evidence for contrast-coding was obtained, whereas the results obtained with variegated surrounds are in no way indicative of contrast coding. This apparent dichotomy between uniform and variegated surrounds is particularly interesting considering that previous psychophysical evidence for contrast coding stems, to the best of my knowledge, *exclusively* from studies using uniform surrounds (Whittle & Challands, 1969; Walraven, 1976; Werner & Walraven, 1982; Whittle, 1992; Mausfeld & Niederée, 1993; Nerger et al., 1993; Chichilnisky & Wandell, 1995, 1996; Niederée & Mausfeld, 1997; Shepherd, 1997, 1999; Richter, 2002). Assuming, as the present findings suggest, that contrast-coding is a special feature of uniform surrounds throws new light on the observation that the perceived colour of an object does not change appreciably when it moves through space although this changes the background against which it is seen: In nature, objects are but seldom located in front of a perfectly uniform background[3], and if they move against a variegated background, the present findings suggest that only rather slight simultaneous contrast effects should occur. In fact, even in the seldom cases where the background consists of two or more really uniform surfaces, the present findings suggest that compelling effects should only be observed in a few cases. We may say, therefore, that none of the present findings are at odds with everyday experience.

Taken together, the present findings suggest that the colour induction effects appearing in uniform surrounds are a combination of two distinct effects that may be understood as resulting from "absolute" and "contrast" coding, respectively, namely simple von Kries scaling and saturation scale truncation and extension. The critical observation is that the von Kries effect is common to both uniform and variegated surrounds whereas the presumably contrast-based truncation and extension seems to be specific to low contrast stimuli with uniform surrounds. This suggest that not only contrast coding can be observed in a "pure form" (e.g. Whittle & Challands, 1969; Wuerger, 1996; Shepherd, 1997; Beer & MacLeod, 2001) but also absolute, von Kries coding. It may therefore be premature and potentially misleading to attribute a more fundamental role to contrast coding than to absolute coding. It appears more reasonable to assume that both codings schemes play specific functional roles in visual processing that may be triggered by different stimulus conditions. Such a "dual coding" could for instance be implemented as two parallel channels of colour information (Arend, 1973, p. 391).

6.3 On the functional equivalence of surrounds

It is often assumed that the colour induction effects appearing in uniform and variegated surrounds are basically of the same nature. A prominent theoretical concept reflecting this general notion is that of *functionally equivalent surrounds*.

Functional equivalence can be defined both on a theoretical level and on a performance level. From a theoretical point of view, the concept of functionally equivalent stimuli is defined with respect to a *particular mechanism*: Two stimuli are functional equivalent whenever the output of this particular mechanism is the same for both stimuli. On the level of psychophysical performance, functional equivalence is understood with respect to a *behavioral criterion*: Two stimuli are regarded as functional equivalent in this sense, if they lead to indistinguishable experimental results. Failures of functional equivalence on the behavioral level can only be taken as an unequivocal indicator of functional equivalence with respect to a specific mechanism if the behavior of the subjects depends exclusively on this mechanism. If, on the other hand, a further mechanism also

---

[3] Based on my own experience, it seems that in the seldom cases in which objects are perceived against a uniform background, 'background-independent constancy' does indeed break down. I recall, for instance, perceiving birds flying against a background of clouds. As they crossed the border between a rather dark and a much lighter cloud, their perceived brightness changed rather impressively.

plays a role, failure of functional equivalence on the behavioral level could result even if functional equivalence with respect of one of the mechanisms holds.

In the colour vision literature, functional equivalence is often discussed with reference to mechanisms of colour constancy. If, for instance, the colour of the illuminant is estimated based on the average colour, all stimuli—variegated or uniform—should be functionally equivalent whenever they have the same mean colour code. In order to test the viability of this theoretical concept, it has been translated into the empirical question whether it is possible to find, for any arbitrary surround A, a corresponding uniform surround A' such that A and A' induce the "same colour shifts" (Valberg & Lange-Malecki, 1990), in the sense that *any* conceivable central patch should look the same whether it is embedded in A or A' (Andres, 1997; Brill, 2000). A number of previous studies (e.g. Schirillo & Shevell, 1996; Brown & MacLeod, 1997; Shevell & Wei, 1998; Barnes et al., 1999), suggest that this is not possible. The present findings also speak strongly against the possibility of finding a uniform surround which is functionally equivalent to a given variegated surround in this behavioral sense: The data curves obtained in our asymmetric matching experiment with pairs of variegated surrounds are approximately linear and consistent with simple von Kries scaling. If there exists a pair of corresponding functionally equivalent uniform surrounds, it should be possible to obtain the same simple data curve with this pair of uniform surrounds. This does not appear feasible, however, since all of our data curves obtained with pairs of uniform surrounds demonstrate strong non-linearities.

However, the present findings questions the assumption that colour induction effects can be ascribed to a single unique mechanism. Instead, the data curves obtained with uniform surrounds seem to be due to two distinct mechanisms, namely one yielding a simple and weak von Kries scaling and another leading to saturation scale extension and truncation. It seems that the former mechanism operates in the same simple manner both in the case of uniform and variegated surrounds, whereas the latter only plays a role in uniform surrounds. These findings suggest that even seemingly simple stimuli may trigger more than one colour coding mechanism. Accordingly, failures of functional equivalence found in previous as well as the present study should be interpreted cautiously; they do not preclude the possibility of functional equivalence with respect to specific contributing mechanisms.

6.4 The validity of grey settings and matching data

The methods of asymmetric colour matching and grey settings are the most common experimental techniques for measuring simultaneous contrast and context effects in colour vision at large. They are also widely used in studies of colour constancy (e.g. Brainard, Kraft, & Longère, 2003; Brainard, 1998; Golz & MacLeod, 2002). A central implicit assumption providing the rationale for these methods is the notion that the effect of the surround on the colour appearance of the target can be compensated for by adjusting the cone excitation values of the target. The results of the present experiments strongly suggest that when uniform surrounds are used, this assumption is unwarranted. Consequently, in this case, the interpretation of grey settings and matching data is fraught with problems. As the results of experiments reported in chapter 4 demonstrate, a grey setting made by a subject can not uncritically be assumed to imply that the target actually appeared achromatic. Instead, it must be interpreted as the setting at which the subject was able to achieve the possible approximation to a truly achromatic colour appearance of the target. The same holds, *mutatis mudandis*, for asymmetric colour matches: The fact that a subject has made a certain match does not imply that he was acutally able to make the two targets appear equal, only that they appear as similar as possible. Consider, for instance, the findings of Smith and Pokorny (1996; see also Miyahara et al., 2001). In asymmetric colour matching experiments rather similar to the present ones, they found that targets with chromaticities approximately in the range between the chromaticity of equal energy white and the chromaticity of the surround in which they were presented were all matched to approximately the target chromaticity in the second surround. If these matches are interpreted as representing perceptual indentity, the interpretation offered by Smith and Pokorny (1996), namely that there is a 'hiatus' in chromaticity space in which colour appearance does not vary greatly with the chromaticity of the target, seems inevitable. In light of the present findings, though, it appears that colour appearance does indeed vary appreciable within the 'hiatus', but that it varies along a range of low saturation levels, none of which can be produced in the other surround due to saturation scale extension and truncation, hence they are all matched to the best available alternative, which is the same for all.

6.5 The representativity of grey settings

Beyond questioning the validity of grey settings and asymmetric matches resulting from experiments in which uniform surrounds are used, the present results also have implications for an issue recently discussed by Speigle and Brainard (1999), namely whether grey settings yield an estimate of the size of simultaneous contrast or colour constancy effects, which is representative also for targets which do not appear grey. Speigle and Brainard (1999) compared asymmetric colour matches with grey settings obtained under identical experimental conditions and found this to be the case: Their asymmetric matching data could be well described by a simple von Kries transform, and the corresponding grey settings could be used to predict the parameters of that transform accurately (up to a common scale factor, which by necessity must remain indeterminate). Hence, their finding suggests that both patches which appear achromatic as well as patches of all other colours are subject to the same general principles of colour induction, namely von Kries scaling. The present findings suggest a different picture, though. The notion of saturation scale truncation and extension, which captures our findings with uniform surrounds very well, implies that measuring the transition point between complementary hues should, as a result, yield the point corresponding to the corner of the 'step' in the data curve. It is clear that the transition point will never actually be chosen as a grey setting, since it corresponds to the chromaticity of the surround. Nevertheless, the results of the experiments reported in chapter 4, as well as further available data suggest that actual grey settings fall rather close to the chromaticity of the surround (Werner & Walraven, 1982; Chichilnisky & Wandell, 1999), which means that they are located well within the region of the 'step'. As can be seen in the schematic representation in Colour Plate XVII as well as in the data plots of Experiments 4 and 5, interpreting any point within the 'step' as a representative measurement, would lead to an overestimation of the overall induction effect. This overestimation is at its largest and most misleading for points close to the transition point.

Thus, based on the present findings we may conclude that, when uniform equiluminant surrounds are used, grey settings can be expected to overestimate the general induction effect. Speigle and Brainard (1999), however, used more naturalistic stimuli, and it is reasonble to assume that they obtained a different result because their stimuli did not evoke saturation scale extension and truncation.

6.6 The relation to colour constancy

According to a widespread idea which dates back at least to Helmholtz (1911), the phenomenon of simultaneous contrast can be understood as "a misdirected attempt to obtain colour constancy" (Walraven, Benzschawel, & Rogowitz, 1987, p. 269). While the general idea does not seem unreasonable, the present findings suggest that one should be cautious about interpreting the induction effects obtained in uniform surrounds in terms of colour constancy: Saturation scale extension and truncation, which seems to be responsible for the strong effects obtained with equiluminant uniform surrounds, appears to be a feature which is particular to this type of stimuli. Therefore the results obtained with uniform surrounds may not be representative of induction effects occurring in stimuli that resemble natural situations more closely, such as those used by Speigle and Brainard (1999). Since it is difficult to conceive of a sensible functional role in mechanisms correcting for illuminations changes for the distinguishing feature of saturation scale truncation and extension, namely the ‘missing colours’ phenomenon, it appears more natural to regard it as an additional effect, that is not directly related to a basic colour constancy mechanism. In the light of our findings, as well as those of Speigle and Brainard (1999), a more plausible candidate for such a mechanism is simple von Kries scaling, which is consistent with all the present data sets obtained with variegated surrounds, and also seems to play a basic, yet not exlusive, role in uniform surrounds.

If correct, this interpretation implies that uniform surrounds should be avoided in studies of colour constancy. If the immediate surround of the target patches is uniform, as is common in many investigations of colour constancy (e.g. Walraven et al., 1987; Bäuml, 1995, 2001; Lucassen & Walraven, 1993, 1996), it can be expected that saturation scale extension and truncation is evoked, and to the extent that this effect is unrelated to colour constancy, the results obtained may be misleading.

6.7 The role of perceptual transparency

The present findings provide, in keeping with several previous observations (Masin & Idone, 1981; Brenner & Cornelissen, 1991; Brown & MacLeod, 1997; Mausfeld, 1998; Niederée, 1998), strong evidence that impressions reminscent of perceptual transparency are evoked by the simple centre-surround stimuli typically employed in studies of simultaneous colour contrast. There are several indications that the strong colour inductions effects which are here described in terms of saturation scale extension and truncation, and appear to be particular to uniform surrounds, are intimately related to transparency-like colour scission.

A question of central theoretical importance is whether the correlation between the strength of the colour induction effect and the phenomenal impression of transparency observed in the present experiments indicates a causal link between the perception of transparency and colour induction. This question is not easy to address experimentally, since the perception of transparency as a putatively causal variable can only be manipulated indirectly by changing properties of the stimulus, which theoretically may influence the colour induction effect directly. In the present experiments, the strength of the colour induction effects, as well as the impression of transparency were found to covary with the uniformity of the surround, so both the uniformity of the surround in itself as well as the perception of transparency can be regarded as candidate causes of the strong induction effect. Taking other evidence into consideration, though, it appears that the impression of transparency constitutes a more relevant variable than the uniformity of the surround. In the present experiments it was found that eliminating the impressions of transparency by using variegated surrounds instead of uniform ones eliminated the strongest colour induction effects. If the perception of transparency is the relevant variable, however, it should also be possible to obtain strong colour induction effects in variegated surrounds, provided that the impression of transparency is supported. In the case of variegated surrounds, impressions of transparency may be reinstated by using a variegated target instead of a uniform one, as evidenced by the Chubb illusion (Chubb, Sperling, & Solomon, 1989; Lotto & Purves, 2001). Hurlbert and Wolf (2004) investigated this case, and did indeed find strong induction effects, which may be taken to suggest that the perception of transparency is more relevant for strength of the colour induction effect than the uniformity of the surround as such.

6.8 Relation to simultaneous brightness contrast

Although simultaneous brightness contrast and chromatic simultaneous colour contrast are clearly intimately related phenomena, they are often investigated separately, in different research traditions. One reason for this is presumably that a ’divide and conquer strategy’ is pursued: Try to understand simultaneous brightness contrast, and chromatic simultaneous contrast, each in isolation, and then turn to the more general case, which can be expected to pose a greater challenge: As is well known, the chromatic and achromatic aspects of simultaneous contrast interact in rather complex manners (Evans, 1964, 1974; Valberg, 1975; Lotto & Purves, 1999). Studying the achromatic and chromatic aspects of simultaneous contrast simplifies matters considerably. A fact which is well-known to the extent that it is often regarded as trivial (which it certainly is not), is that when an achromatic patch is embedded in another achromatic surround, it does not cease to appear achromatic: The only aspect of its colour appearance that is affected is its brightness (or lightness). Abusing mathematical terminology, one may say that the set of achromatic colour impressions is ‘closed’ under ‘the operations of simultaneous contrast’. This actually quite remarkable fact makes it possible to study simultaneous brightness contrast in isolation, without ever having to consider the chromatic aspects of simultaneous contrast, as has been done so extensively that it makes sense to speak of a separate research tradition (Gilchrist et al., 1999). In much the same way, any set of isoluminant colours appears to be closed under the operation of simultaneous contrast. Accordingly, when only isoluminant stimuli are employed, it is possible to study the purely chromatic aspects of simultaneous contrast without having to take simultaneous brightness contrast into consideration. This strategy was pursued in the present investigations. It should be kept in mind, though, that the reason for this was primarily the wish to keep things simple. Once one also allows for luminance differences between the target and the surrounds, I have little doubt that matters will become vastly less clear-cut and more complicated than the present results suggest. In particular, simultaneous contrast effects tend to become less dramatic when luminance contrast and chromatic contrast are mixed, than when they are each studied in isolation (Kirschmann, 1891; Olkkonen, Saarela, Peromaa, & Laurinen, 2002).

Although simultaneous brightness contrast was not investigated here, there is reason to believe that there is something to be learned about it based on the present findings, simply because purely achromatic simultaneous contrast and purely chromatic simultaneous contrast seem to be analogous in many respects. In both cases, there is clear evidence for contrast coding (Whittle & Challands, 1969; Whittle, 1994b; Arend, Buehler, & Lockhead, 1971; Beer & MacLeod, 2001; Kingdom, 2003). Furthermore, inspection of the data curves from some of the asymmetric brightness matching experiments reported in the literature (Takasaki, 1966; Schirillo, 1999) suggests the existence of a charateristic ‘step’ analogous to the ones revealed in the present experiments. This step could perhaps be accounted for in terms of blackness (or whiteness) scale truncation and extension, in analogy to the notions of saturation scale extension and truncation. As demonstrated in Figure 6.3, which is constructed in a manner analogous to Colour Plates XIV and XV, increments invariably appear whitish in some way, whereas decrements appear blackish. Yet the range of whitish colour impressions which can be produced in a bright surround is truncated relative to the range of whitish colour impressions which can be produced in a dark surround. There seems to be, in essence, no target which appears slighly less whitish than the surround, although the surround is very whitish, since any decrement looks blackish. The same holds, *mutatis mutandis*, with respect to blackish colour impressions.

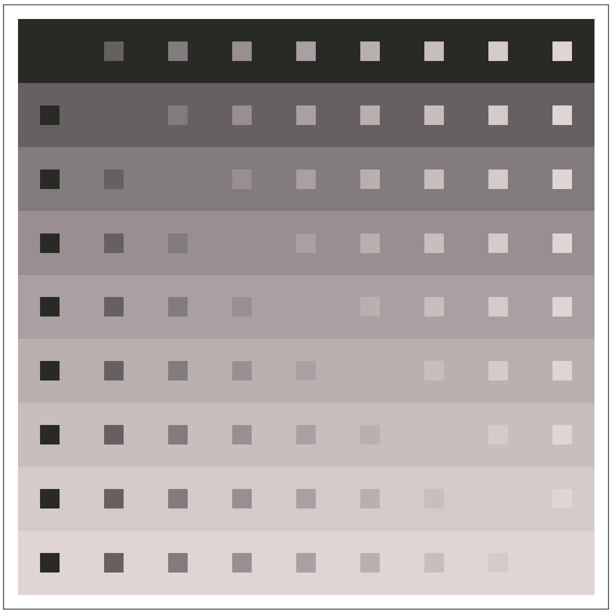

Figure 6.3: Within each coloum, the small squares are physically identical, that is, in each row actually the same scale of colours from dark to bright is reproduced. Note that all of the targets above the diagonal, which are luminance decrements appear whitish, while those below the diagonal, which are luminance increments, appear blackish. Note also that the upper scales, which are embedded in the dark surrounds, contain very dark whitish targets, particularly in the darkest surround. The scales that are embedded in the brighter surround, on the other hand, contain no dark-appearing whitish targets. In exchange, though, they contain more bright-appearing blackish targets than the scales embedded in the dark surrounds.

6.9 Conclusions

The results of the present experiments show that simultaneous contrast is a much more complex and multifaceted phenomenon than generally appreciated. While the simultaneous contrast effects evoked by complex, variegated surrounds appear to be of a relatively simple nature, uniform surrounds seem to trigger more complex perceptual phenomena involving transparency-like dual colour impressions.

The perhaps most important insight that can be gained from the present investigations is that classical models of simultaneous contrast, which are based upon the assumption that the colour appearance changes induced into a target by a uniform coloured surround can be compensated by adjusting the colour-coordinates of the target, can at best yield an incomplete description of the relevant phenomena. In the case of uniform surrounds, this ‘compensation assumption’ is clearly unwarranted, and correspondingly, psychophysical methods which are based upon it, such as asymmetric colour matching and grey settings, yield results flawed by artifact. Experimental data obtained with these methods must therefore be interpreted with care.

Though the present findings demonstrate that classical models of simultaneous contrast are based on an erroneous assumption, they also provide strong evidence for colour-coding mechanisms responding to the difference between the cone excitation vectors of the target and the surround, as envisioned in the works of Whittle (1994b, 2003) and Walraven (1976). Contrary to the original assumptions, though, the contrast-coding scheme does not describe the total colour impression, but only a phenomenally distinct part of it. The total colour impression consists of two phenomenally distinct and simultaneously perceived layers, one of which is perceived ‘through’ the other. One of these colour components behaves according to the principles of contrast coding, whereas the other, which has the phenomenal quality of a transparent haze covering the target, appears in the colour of the surround. Once it is recognised that contrast coding determines but a part of the total colour impression, the apparently contradictory experimental findings of many studies, in particular those reported in connection with the Walraven-Shevell controversy (Walraven, 1976, 1979; Shevell, 1978, 1980; Drum, 1981; Adelson, 1981; Davies et al., 1983) become intelligible.

Though there is, in the above sense, clear evidence for contrast coding in the present data, the total pattern of findings suggests that contrast coding does not play the fundamental, all-encompassing role in colour vision often attributed to it (Whittle, 2003). Rather, it seems that absolute colour coding plays an equally important role. Clear evidence for contrast coding is only obtained when uniform surrounds are used, while the induction effects obtained with comparable variegated surrounds are considerably smaller than one would expect based on contrast coding. It seems therefore, that contrast-coded colour signals are exploited by the visual system only under special conditions. One may speculate that conditions supporting the perception of transparency may trigger the utilisation of the contrast signal.

The induction effects obtained with uniform surrounds appear to a be a combination of two effects, one of which is identical to the relatively small ‘von Kries-type’ perturbation of the target’s absolute colour signal observed with comparable variegated surrounds. The other effect, which is particular to uniform surrounds, appears to be of a rather peculiar nature. Basically, any uniform surround seems to prohibit the perception of a certain set of colour impressions. The set of such ’missing colours’ depends on the colour of the surround in characteristic ways captured by a descriptive model introduced here as ‘saturation scale extension and truncation’. This model provides a rather good description of many curious aspects of simultaneous contrast in uniform surrounds, such as the vast variation in the size of the effect for different target colours, Meyer’s effect (Meyer, 1855; Helmholtz, 1911; Brown, 2003), and the gamut expansion effect of Brown and MacLeod (1997).

6.10 Outlook

The primary focus of the present study was on the kind of simultaneous colour contrast that occurs in uniform surrounds. While this is presumably one of the most well-known demonstations of the context-dependence of colour perception, it is by no means the only one. Several different stimulus displays demonstrating context effects have been reported and investigated in the literature, such as the Mach bands (Mach, 1866; Fiorentini, 1972), the Craik-O’Brien illusion (Wachtler & Wehrhahn, 1997), the Hermann Grid (Hermann, 1870; Spillmann, 1994), Wertheimer-Benary cross (Benary, 1924), the Munker-White effect (Munker, 1970; White, 1979), neon colour spreading (Munker, 1970; van Tuijl, 1975), Adelson’ snake and tile illusions (Adelson, 1993, 2000), the watercolour illusion (Pinna, Brelstaff, & Spillmann, 2001), and various forms of colour and brightness assimilation (Bezold, 1874; de Weert & van Kruysbergen, 1997; Shapley & Reid, 1985; Monnier & Shevell, 2003).

The ultimate of goal vision science is to understand the lawful workings of the visual mechanisms which manifest themselves both in everyday colour perception as well as in these demonstrations. Towards this aim it would be of great value to know which phenomena share a common mechanism and which do not, i.e. to develop a natural taxonomy consisting of a limited set of visual mechanisms as opposed to a meaningless list of phenomena. Yet the task of categorising the phenomena in a way that carves nature at its joints is by no means a trivial one, given that phenomena which appear dissimilar at a superficial level may nevertheless be due to the same visual mechanisms, and vice versa. It is therefore not surprising that rather different proposals have been made concerning which phenomena share a common mechanism and which do not (e.g. Anderson, 1997; Gilchrist et al., 1999; Gilchrist & Economou, 2003; Kingdom, 2003; Blakeslee & McCourt, 2003).

Among the phenomena of colour perception which have most often been linked to various other phenomena, simultaneous contrast definitely tops the list. It is often linked to Mach bands, the Craik-O'Brien illusion, and the Hermann Grid because they all involve the perceptual enhancement of perceived contrast. It has also been proposed that the mechanisms behind classical simultaneous contrast are at work also in the Munker-White illusion (Gilchrist et al., 1999) and neon colour spreading (da Pos & Bressan, 2003). The insights gained in the present investigations may be used to evaluate the plausibility of such proposed links, and qualify their meaning. Since the kind of simultaneous contrast occuring in uniform surrounds appear to be due to two distinct mechanisms, any proposal of a link between simultaneous contrast and any other phenomenon should not be made without stating which of the two 'simultaneous contrast mechanisms' one refers to. Once such a commitment has been made, one has a better basis for evaluating the plausibility of the proposal in psychophysical experiments. I am particularly optimistic that the regularities described in this thesis as saturation scale extension and trunctation may serve as a psychophysical signature of a special contrast-coding mechanism, and that checking whether these regularities are present in other visual illusions may help answer the question whether they share a mechanism with classical simultaneous contrast. I am tempted to venture the hypothesis that saturation scale extensions and truncation (and/or their analogues in the brightness domain) can be documented in the Munker-White illusion, as well as in Adelson's snake illusion. This remains a question for future research, though.

In conclusion, I would like to raise one final issue which I think may be worth considering for future research on colour vision: Although one is not primarily (or even at all) interested in simultaneous contrast in many lines of psychophysical research on colour vision, it is a trivial fact that any target which is to be judged, scaled or matched must, by necessity, be presented against a background of some sort, even if it is a background of complete darkness. More often than not, a uniform background is used, because of a wish for simplicity and well-specified viewing conditions. In view of the present findings, which suggest that quite special and rather complex phenomena occur in the presence of uniform surrounds, this might not always be a very good idea, and for many purposes it may turn out to be more advantageous to use variegated, unstructured surrounds. In applications where subjects are to match the colour of a target by adjusting the colour of an adjustable patch, the latter should not be presented in a uniform surround, since this may preclude the possibility of realising some colour impressions.

6.11 Acknowledgements

I am indebted to Johannes Andres, Jürgen Golz, Donald I. A. MacLeod, Rainer Mausfeld and Daniel Wollschläger for stimulating discussions and helpful suggestions. Reinhard Niederée, Eike Richter and Gunnar Wendt freely contributed insights from their own work on related issues. This work would not have been possible without the extensive and all-encompassing advice and mentorship of Franz Faul. It would be futile to express my gratitude in words. This work was supported by the grants FA-425/1-1 and 1-2 from the Deutsche Forschungsgemeinschaft to Franz Faul.

References

Abney, W., & Festing, E. R. (1886). Colour photometry. *Philosophical Transaction of the Royal Society of London*, *177*, 423-456.

Adelson, E. H. (1981). Looking at the world through a rose-colored Ganzfeld. *Vision Research*, *21*(5), 749-750.

Adelson, E. H. (1993). Perceptual organization and the judgment of brightness. *Science*, *262*(5142), 2042-2044.

Adelson, E. H. (2000). Lightness perception and lightness illusions. In M. Gazzaniga (Ed.), *The new cognitive neurosciences, 2 edition* (p. 339-351). Cambrigde, MA: MIT Press.

Adelson, E. H., & Anadan, P. (1990). Ordinal characteristics of transparency. In *AAAI Workshop on Quantitative Vision* (p. 77-81). Boston, MA.

Anderson, B. L. (1997). A theory of illusory lightness and transparency in monocular and binocular images: the role of contour junctions. *Perception*, *26*, 419-452.

Andres, J. (1997). *Formale Modelle der Farbkonstanz und ihre Untersuchung durch die Methode der stetigen Szenenvariation*. Post-doctoral thesis, Christian-Albrechts-Universität zu Kiel.

Arend, L. E. (1973). Spatial differential and integral operations in human vision: Implications of stabilized retinal image fading. *Psychological Review*, *80*(5), 374-395.

Arend, L. E. (1994). Surface colors, illumination, and surface geometry: Intrinsic-image models of human color perception. In A. L. Gilchrist (Ed.), *Lightness, brightness and transparency* (p. 159-213). Lawrence Erlbaum Associates.

Arend, L. E., Buehler, J. N., & Lockhead, G. R. (1971). Difference information in brightness perception. *Perception & Psychophysics*, *9*, 367-370.

Arend, L. E., & Goldstein, R. (1987). Simultanous constancy, lightness and brightness. *Journal of the Optical Society of America A*, *4*, 2281-2285.

Arend, L. E., & Reeves, A. (1986). Simultaneous color constancy. *Journal of the Optical Society of America A*, *3*(10), 1743-1751.

Arend, L. E., & Spehar, B. (1993a). Lightness, brightness, and brightness contrast: 1. Illuminance variation. *Perception & Psychophysics*, *54*(4), 446-456.

Arend, L. E., & Spehar, B. (1993b). Lightness, brightness, and brightness contrast: 2. Reflectance variation. *Perception & Psychophysics*, *54*(4), 457-468.

Barnes, C. S., Wei, J., & Shevell, S. K. (1999). Chromatic induction with remote chromatic contrast varied in magnitude, spatial frequency, and chromaticity. *Vision Research*, *39*, 3561-3574.

Bäuml, K. H. (1995). Illuminant changes under different surface collections: examining some principles of color appearance. *Journal of the Optical Society of America A*, *12*(2), 261-71.

Bäuml, K. H. (2001). Increments and decrements in color constancy. *Journal of the Optical Society of America A*, *18*(10), 2419-29.

Beer, R., & MacLeod, D. (2001). Afterimage appearance and adaptation [abstract]. *Journal of Vision*, *1*(3), 458a, http://journalofvision.org/1/3/458, DOI: 10.1167/1.3.458.

Benary, W. (1924). Beobachtungen zu einem Experiment über Helligkeitskontrast. *Psychologische Forschung*, *5*, 131-142.

Bezold, W. von. (1874). *Die Farbenlehre*. Braunschweig: Westermann.

Billock, V. A., Gleason, G. A., & Tsou, B. H. (2001). Perception of forbidden colors in retinally stabilized equiluminant images: an indication of softwired cortical color opponency? *Journal of the Optical Society of America A*, *18*(10), 2398-2403.

Blakeslee, B., & McCourt, M. E. (2003). A multiscale spatial filtering account of brightness phenomena. In L. Harris & M. Jenkin (Eds.), *Levels of perception* (p. 47-72). Springer-Verlag.

Brainard, D. H. (1989). Calibration of a computer controlled color monitor. *Color Research and Application*, *14*, 23-34.

Brainard, D. H. (1998). Color constancy in the nearly natural image. 2. Achromatic loci. *Journal of the Optical Society of America A*, *15*(2), 307-325.

Brainard, D. H. (2003). Color constancy. In L. Chalupe & J. Werner (Eds.), *The visual neurosciences* (p. 948-961). MIT Press.

Brainard, D. H., Kraft, J. M., & Longère, P. (2003). Color constancy: developing empirical tests of computational models. In R. Mausfeld & D. Heyer (Eds.), *Colour perception: Mind and the physical world* (p. 307-328). Oxford University Press.

Brenner, E., & Cornelissen, F. W. (1991). Spatial interactions in color vision depend on distances between boundaries. *Naturwissenschaften*, *78*, 70-73.

Brill, M. H. (2000). Understanding color matches: What are we taking for granted. In S. Davis (Ed.), *Color perception: Philosophical, psychological, artistic, and computational perspectives* (Vol. 9, p. 141-151). Oxford University Press.

Brown, R. O. (2003). Backgrounds and illuminants: the yin and yang of color constancy. In R. Mausfeld & D. Heyer (Eds.), *Colour: Mind and the physical world* (p. 247-271). Oxford University Press.

Brown, R. O., & MacLeod, D. I. A. (1997). Color appearance depends on the variance of surround colors. *Current Biology*, *7*(11), 844-849.

Burns, S. A., Elsner, A. E., Pokorny, J., & Smith, V. C. (1984). The Abney effect: chromaticity coordinates of unique and other constant hues. *Vision Research*, *24*(5), 479-489.

Chichilnisky, E. J., & Wandell, B. A. (1995). Photoreceptor sensitivity changes explain color appearance shifts induced by large uniform backgrounds in dichoptic matching. *Vision Research*, *35*(2), 239-254.

Chichilnisky, E. J., & Wandell, B. A. (1996). Seeing gray through the ON and OFF pathways. *Visual Neuroscience*, *13*(3), 591-596.

Chichilnisky, E. J., & Wandell, B. A. (1999). Trichromatic opponent color classification. *Vision Research*, *39*(20), 3444-3458.

Chubb, C., Sperling, G., & Solomon, J. A. (1989). Texture interactions determine perceived contrast. *Proceedings of the National Academy of Sciences U S A*, *86*(23), 9631-9635.

Crane, H., & Piantanida, T. P. (1983). On seeing reddish green and yellowish blue. *Science*, *221*, 1078-1080.

Da Pos, O. (1989). *Trasparenze*. Icone s.r.l.

da Pos, O., & Bressan, P. (2003). Chromatic induction in neon colour spreading. *Vision Research*, *43*(6), 697-706.

Davies, S. E., Faivre, I. A., & Werner, J. S. (1983). Transient processing in chromatic induction. *Vision Research*, *23*(7), 707-712.

de Weert, C. M., & van Kruysbergen, N. A. (1997). Assimilation: central and peripheral effects. *Perception*, *26*(10), 1217-24.

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

De Valois, R. L., Abramov, I., & Jacobs, G. H. (1966). Analysis of response patterns of LGN cells. *Journal of the Optical Society of America*, *56*, 966-977.

Drum, B. (1981). Additive effect of backgrounds in chromatic induction [Letter to the editor]. *Vision Research*, *21*, 959-961.

D'Zmura, M., Colantoni, P., Knoblauch, K., & Laget, B. (1997). Color transparency. *Perception*, *26*, 471-492.

Eagleman, D. M. (2001). Visual illusions and neurobiology. *Nature Reviews Neuroscience*, *2*(12), 920-926.

Eisner, A., & MacLeod, D. I. A. (1980). Blue-sensitive cones do not contribute to luminance. *Journal of the Optical Society of America*, *70*(1), 121-123.

Ekroll, V., & Faul, F. (2002). Perceptual transparency in neon color spreading displays. *Perception & Psychophysics*, *64*(6), 945-955.

Evans, R. M. (1964). Variables of perceived color. *Journal of the Optical Society of America*, *54*(12), 1467-1474.

Evans, R. M. (1974). *The perception of color*. New York: Wiley.

Faul, F. (1997). *Theoretische und experimentelle Untersuchung chromatischer Determinanten perzeptueller Transparenz*. Dissertation, Christian-Albrechts-Universität, Kiel.

Faul, F., & Ekroll, V. (2002). Psychophysical model of chromatic perceptual transparency based on subtractive color mixture. *Journal of the Optical Society of America A*, *19*(6), 1084-1095.

Finlayson, G. D., Drew, M. S., & Funt, B. (1994). Color constancy: Generalized diagonal transforms suffice. *Journal of the Optical Society of America A*, *11*, 3011-3020.

Fiorentini, A. (1972). Mach band phenomena. In D. Jameson & L. M. Hurvich (Eds.), *Handbook of sensory physiology* (Vol. VII/4 Visual Psychophysics, p. 188-201). Berlin: Springer-Verlag.

Fuchs, W. (1923). Experimentelle Untersuchungen über das simultane Hintereinandersehen auf derselben Sehrichtung. *Zeitschrift für Psychologie*, *91*, 195-235.

Gelb, A. (1929). Die „Farbenkonstanz" der Sehdinge. In A. Bethe, G. von Bergman, G. Embden, & A. Ellinger (Eds.), *Handbuch der der normalen und pathologischen Physiologie* (p. 594-687). Berlin: Springer. (Bd. 12, 1. Hälfte. Rezeptionsorgane II)

Gilchrist, A., Delman, S., & Jacobsen, A. (1983). The classification and integration of edges as critical to the perception of reflectance and illumination. *Perception & Psychophysics*, *33*(5), 425-436.

Gilchrist, A., & Economou, E. (2003). Dualistic versus monistic accounts of lightness perception. In L. Harris & M. Jenkin (Eds.), *Levels of perception* (p. 11-22). Springer-Verlag.

Gilchrist, A., Kossyfidis, C., Bonato, F., Agostini, T., Cataliotti, J., Li, X., et al. (1999). An anchoring theory of lightness perception. *Psychological Review*, *106*(4), 795-834.

Golz, J., & MacLeod, D. I. A. (2002). Influence of scene statistics on colour constancy. *Nature*, *415*(6872), 637-640.

Golz, J., & MacLeod, D. I. A. (2003). Colorimetry for CRT displays. *Journal of the Optical Society of America A*, *20*(5), 769-781.

Hartline, H. K., Wagner, H. G., & Ratliff, F. (1956). Inhibition in the eye of Limulus. *Journal of General Physiology*, *39*, 651-673.

Heggelund, P. (1974). Achromatic color vision. II. Measurement of simultaneous contrast within a bidimensional system. *Vision Research*, *14*, 1081-1088.

Heggelund, P. (1992). A bidimensional theory of achromatic color vision. *Vision Research*, *32*(11), 2107-2119.

Helmholtz, H. von (Ed.). (1911). *Handbuch der Physiologischen Optik*. Hamburg und Leipzig: Verlag von Leopold Vos.

Helson, H., & Michels, W. C. (1948). The effect of chromatic adaptation on achromaticity. *Journal of the Optical Society of America*, *38*(12), 1025-1032.

Henneman, R. H. (1935). A photometric study of the perception of object color. *Archives of Psychology*, *XXVII*, 1-89.

Hering, E. (1887a). Ueber die Theorie des simultanen Contrastes von Helmholtz: I. Der Versuch mit den farbigen Schatten. *Pflügers Archiv für Physiologie*, *40*, 172-191.

Hering, E. (1887b). Ueber die Theorie des simultanen Contrastes von Helmholtz: II. Der Contrastversuch von H. Meyer und die Versuche am Farbenkreisel. *Pflügers Archiv für Physiologie*, *41*, 1-29.

Hering, E. (1887c). Ueber die Theorie des simultanen Contrastes von Helmholtz: III. Der Spiegelcontrastversuch. *Pflügers Archiv für Physiologie*, *41*, 358-367.

Hering, E. (1888). Ueber die Theorie des simultanen Contrastes von Helmholtz: IV. Die subjective „Trennung des Lichtes in zwei complementäre Portionen”. *Pflügers Archiv für Physiologie*, *43*, 1-21.

Hering, E. (1890). Beitrag zur Lehre vom Simultankontrast. *Zeitschrift für Psychologie und Physiologie der Sinnesorgane*, *1*, 18-28.

Hering, E. (1920). *Grundzüge der Lehre vom Lichtsinn*. Berlin: Verlag von Julius Springer.

Hermann, L. (1870). Eine Erscheinung simultanen Contrastes. *Pflügers Archiv für die gesamte Physiologie*, *3*, 13-15.

Hoffman, D. D. (1998). *Visual intelligence. How we create what we see*. New York - London: W. W. Norton & Company.

Hurlbert, A., & Wolf, K. (2004). Color contrast: a contributory mechanism to color constancy. *Progress in Brain Research*, *144*, 147-160.

Hurvich, L. M., & Jameson, D. (1955). Some quantitative aspects of an opponent-colors theory. II. Brightness, saturation, and hue in normal an dichromatic vision. *Journal of the Optical Society of America*, *45*(8), 602-616.

Ishihara, S. (1967). *Tests for Colour-Blindness*. Tokio: Kanehara Shuppan Co., Ltd.

Ives, H. E. (1912). The relation between the color of the illuminant and the color of the illuminated object. *Transactions of the Illuminating Engineering Society*, *7*, 62-72.

Jameson, D., & Hurvich, L. M. (1955). Some quantitative aspects of an opponent-colors theory. I. Chromatic responses and spectral saturation. *Journal of the Optical Society of America*, *45*(7), 546-552.

Jameson, D., & Hurvich, L. M. (1959). Perceived color and its dependence on focal, surrounding, and preceding stimulus variables. *Journal of the Optical Society of America*, *49*, 890-898.

Jameson, D., & Hurvich, L. M. (1961). Opponent chromatic induction: Experimental evaluation and theoretical account. *Journal of the Optical Society of America*, *51*(1), 46-53.

Jameson, D., & Huvich, L. M. (1972). Color adaptation: Sensitivity, contrast, after-images. In D. Jameson & L. M. Hurvich (Eds.), *Handbook of sensory physiology* (Vol. VII/4 Visual Psychophysics, p. 568-581). Berlin: Springer-Verlag.

Judd, D. B. (1949). Response functions for types of vision according to the Müller theory. *U. S. Department of Commerce. National Bureau of Standards. Part of the Journal of Research of the National Bureau of Standards*, *42*(1-16).

Kaiser, P. K., & Boynton, R. M. (1996). *Human color vision*. Washington, DC: Optical Society of America.

Kanizsa, G. (1966). Die Erscheinungsweisen der Farben. In W. Metzger & H. Erke (Eds.), *Allgemeine Psychologie, I. Der Aufbau des Erkennens* (p. 161-191). Göttingen: Verlag für Psychologie, Dr. C. J. Hogrefe. (Volume 1 of Handbuch der Psychologie)

Katz, D. (1911). *Die Erscheinungsweisen der Farben und ihre Beeinflussung durch die individuelle Erfahrung*. Leipzig: Barth.

Kingdom, F. A. (1997). Simultaneous contrast: the legacies of Hering and Helmholtz. *Perception*, *26*(6), 673-677.

Kingdom, F. A. (2003). Levels of brightness perception. In L. Harris & M. Jenkin (Eds.), *Levels of perception* (p. 23-46). Springer-Verlag.

Kinney, J. A. S. (1962). Factors affecting induced color. *Vision Research*, *2*, 503-525.

Kirschmann, A. (1891). Ueber die quantitativen Verhältnisse des simultanen Helligkeits- und Farben-Contrastes. *Philosophische Studien*, *6*, 417-491.

Koenderink, J. J. (2003). Another Babel. *Perception*, *32*(4), 391-394.

Kohlrausch, A. (1935). Zur Photometrie farbiger Lichter. *Licht*, *5*, 259-279.

König, A., & Dieterici, C. (1886). Die Grundempfindungen und ihre Intensitätsvertheilung im Spectrum. *Sitzungsberichte der Akademie der Wissenschaften in Berlin*, *Juli*, 805-829.

Krantz, D. H. (1975a). Color measurement and color theory: II. Opponent color theory. *Journal of Mathematical Psychology*, *12*, 304-327.

Krantz, D. H. (1975b). Color measurement and color theory: I. Representation theorem for Grassmann structures. *Journal of Mathematical Psychology*, *12*, 283-303.

Krauskopf, J., Williams, D. R., & Heeley, D. W. (1982). Cardinal directions of color space. *Vision Research*, *22*(9), 1123-1131.

Kries, J. von. (1905). Die Gesichtsempfindungen. In W. Nagel (Ed.), *Handbuch der Physiologie des Menschen* (Vols. 3, Physiologie der Sinne, p. 109-282). Braunschweig: Friedrich Vieweg und Sohn.

Kuehni, R. G. (2004). Variability in unique hue selection: A surprising phenomenon. *Color Research and Application*, *29*(2), 158-162.

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

Larimer, J., Krantz, D. H., & Cicerone, C. M. (1975). Opponent process additivity. II. Yellow/blue equilibria and nonlinear models. *Vision Research*, *15*(6), 723-731.

Le Grand, Y. (1957). *Light, colour and vision*. London: Chapman & Hall Ltd.

Lennie, P., Pokorny, J., & Smith, V. C. (1993). Luminance. *Journal of the Optical Society of America A*, *10*(6), 1283-1293.

Lotto, R. B., & Purves, D. (1999). The effects of color on brightness. *Nature Neuroscience*, *2*(11), 1010-1014.

Lotto, R. B., & Purves, D. (2001). An empirical explanation of the Chubb illusion. *Journal of Cognitive Neuroscience*, *13*(5), 547-555.

Lucassen, M. P., & Walraven, J. (1993). Quantifying color constancy: evidence for nonlinear processing of cone-specific contrast. *Vision Research*, *33*(5-6), 739-757.

Lucassen, M. P., & Walraven, J. (1996). Color constancy under natural and artificial illumination. *Vision Research*, *36*(17), 2699-2711.

MacAdam, D. L. (1956). Chromatic adaptation. *Journal of the Optical Society of America*, *66*(7), 500-513.

MacAdam, D. L. (1961). Chromatic adaptation. *Vision Research*, *1*, 9-41.

Mach, E. (1866). Über die Wirkung der räumlichen Vertheilung des Lichtreizes auf die Netzhaut. *Sitzungsberichte der kaiserlichen Akademie der Wissenschaften, mathem.-naturwis. Klasse*, *52*, 303-322.

MacLeod, D. I. A., & Boynton, R. M. (1979). Chromaticity diagram showing cone excitation by stimuli of equal luminance. *Journal of the Optical Society of America*, *69*(8), 1183-1186.

Malkoc, G., Kay, P., & Webster, M. A. (2001). Individual differences in unique and binary hues [abstract]. *Journal of Vision*, *2*(10), 32a, http://journalofvision.org/2/10/32, DOI: 10.1167/2.10.32.

Masin, S. C., & Idone, A. M. (1981). Studio sperimentale sulla percezione della trasparenza con figura e sfondo achromatici e omogenei. *Giornale Italiano di Psicologia*, *8*, 265-277.

Mausfeld, R. (1998). Color perception: From Grassman codes to a dual code for object and illumination colors. In W. G. K. Backhaus, R. Kliegl, & J. S. Werner (Eds.), *Color vision* (p. 219-250). Berlin: De Gruyter.

Mausfeld, R. (2003a). 'Colour' as part of the format of different perceptual primitives: the dual coding of colour. In R. Mausfeld & D. Heyer (Eds.), *Colour perception: Mind and the physical world* (p. 381-429). Oxford University Press.

Mausfeld, R. (2003b). Conjoint representations and the mental capacity for multiple simultaneous perspectives. In H. Hecht, R. Schwartz, & M. Atherton (Eds.), *Looking into pictures: An interdisciplinary approach to pictorial space* (p. 17-60). MIT Press.

Mausfeld, R., & Andres, J. (2002). Second-order statistics of colour codes modulate transformations that effectuate varying degrees of scene invariance and illumination invariance. *Perception*, *31*(2), 209-224.

Mausfeld, R., & Niederée, R. (1993). An inquiry into relational concepts of colour based on incremental principles of colour coding for minimal relational stimuli. *Perception*, *22*, 427-462.

Maxwell, J. C. (1860). On the theory of compound colours, and the relations of the colours of the spectrum. *Philosophical Transactions of the Royal Society of London*, *150*, 57-84.

Metelli, F. (1970). An algebraic development of the theory of perceptual transparency. *Ergonomics*, *13*, 59-66.

Meyer, H. (1855). Über Kontrast- und Komplementärfarben. *Annalen der Physik*, *XCV*, 170-171.

Miyahara, E., Smith, V. C., & Pokorny, J. (2001). The consequences of opponent rectification: the effect of surround size and luminance on color appearance. *Vision Research*, *41*(7), 859-871.

Mollon, J. D. (2003). The origins of modern color science. In S. Shevell (Ed.), *Color science* (p. 1-39). Optical Society of America.

Mollon, J. D., & Cavonius, C. R. (1987). The chromatic antagonisms of Opponent Process Theory are not the same as those revealed in studies of detection and discrimination. In G. Verriest (Ed.), *Colour vision deficiencies VIII* (p. 473-483). The Hague: W. Junk.

Monnier, P., & Shevell, S. K. (2003). Large shifts in color appearance from patterned chromatic backgrounds. *Nature Neuroscience*, *6*(8), 801-802.

Munker, H. (1970). *Farbige Gitter, Abbildungen auf der Netzhaut und übertragungstheoretische Beschreibung der Farbwahrnehmung.* Postdoctoral thesis, Ludwig-Maximilians-Universität München.

Nerger, J. L., Piantanida, T. P., & Larimer, J. (1993). Color appearance of filled-in backgrounds affects hue cancellation, but not detection thresholds. *Vision Research*, *33*(2), 165-172.

Newton, I. (1704/1952). *Opticks*. New York: Dover Publications, Inc.

Niederée, R. (1998). *Die Erscheinungsweisen der Farben und ihre stetigen Übergangsformen: Theoretische und Experimentelle Untersuchungen zur relationalen Farbkodierung und zur Dimensionalität vollständiger perzeptueller Farbcodes.* Postdoctoral thesis, Philosophische Fakultät der Christian-Albrechts-Universität zu Kiel, Kiel.

Niederée, R., & Mausfeld, R. (1997). Increment-decrement asymmetry in dichoptic matching with haploscopically superimposed backgrounds. *Vision Research*, *37*(5), 613-616.

Olkkonen, K. M., Saarela, T. P., Peromaa, T. L., & Laurinen, P. I. (2002). Effects of chromatic contrast on brightness perception [abstract]. *Perception*, *31*(Supplement), 184d, http://perceptionweb.com/ecvp02/0519.html.

Palmer, G. (1777). *Theory of colours and vision*. London: S. Leacroft.

Perls, L. (1932). *Die Erscheinungen des simultanen Kontrastes und der Eindruck der Feldbeleuchtung*. Dissertation, Johann Wolfgang Goethe-Universität, Frankfurt am Main.

Pinna, B., Brelstaff, G., & Spillmann, L. (2001). Surface color from boundaries: a new 'watercolor' illusion. *Vision Research*, *41*(20), 2669-2676.

Richter, E. (2002). *Modelle der Kontrastcodierung*. Ph. D. thesis, Martin-Luther-Universität Halle-Wittenberg.

Rinner, O., & Gegenfurtner, K. R. (2000). Time course of chromatic adaptation for color appearance and discrimination. *Vision Research*, *40*(14), 1813-1826.

Rinner, O., & Gegenfurtner, K. R. (2002). Cone contributions to colour constancy. *Perception*, *31*(6), 733-746.

Sakata, K. (2004a). The effects of background colour on chromatic afterimage [abstract]. *Journal of Vision*, *4*(11), 80a, http://journalofvision.org/4/11/80, DOI: 10.1167/4.11.80.

Sakata, K. (2004b). The influence of chromatic adaptation on colour afterimage [abstract]. *Perception*, *33*(Supplement), 62a.

Schirillo, J. A. (1999). Surround articulation. I. Brightness judgments. *Journal of the Optical Society of America A*, *16*(4), 793-803.

Schirillo, J. A., & Shevell, S. K. (1996). Brightness contrast from inhomogeneous surrounds. *Vision Research*, *36*(12), 1783-1796.

Schrödinger, E. (1920a). Grundlinien einer Theorie der Farbenmetrik im Tagessehen. Erste und Zweite Mitteilung. *Annalen der Physik*, *63*(21), 397-456.

Schrödinger, E. (1920b). Über das Verhältnis der Vierfarben- zur Dreifarbentheorie. *Sitzungsberichte der Akademie der Wissenschaften in Wien. Mathematisch-naturwissenschaftliche Klasse, Abteilung 2a*, *134*, 471-490.

Schumann, F. (1921). Die Dimensionen des Sehraumes. *Zeitschrift für Psychologie*, *86*, 253-277.

Scina, R. (1847). Su taluni fenomeni che presentano i cristalli colorati. *Raccolta fisico-chimica italiana*, *II*, 207-210.

Shapley, R., & Reid, R. C. (1985). Contrast and assimilation in the perception of brightness. *Proceedings of the National Academy of Sciences U S A*, *82*(17), 5983-5986.

Sharpe, L. T., & Stockman, A. (1999). Cone spectral sensitivities and color matching. In K. R. Gegenfurtner (Ed.), *Color vision: From genes to perception* (p. 53-87). Cambridge: Cambridge University Press.

Shepherd, A. J. (1997). A vector model of colour contrast in a cone-excitation colour space. *Perception*, *26*(4), 455-470.

Shepherd, A. J. (1999). Remodelling colour contrast: implications for visual processing and colour representation. *Vision Research*, *39*(7), 1329-1345.

Shevell, S. K. (1978). The dual role of chromatic backgrounds in color perception. *Vision Research*, *18*, 1649-1661.

Shevell, S. K. (1980). Unambiguous evidence for the additive effect in chromatic adaption. *Vision Research*, *20*, 637-639.

Shevell, S. K., & Wei, J. (1998). Chromatic induction: border contrast or adaptation to surrounding light? *Vision Research*, *38*, 1561-1566.

Sinden, R. H. (1923). Studies based on spectral complementaries: I. subjective saturation of spectral hues. II. Individual variations of the normal color sense. *Journal of the Optical Society of America*, *7*(12), 1123-1153.

Smith, V. C., & Pokorny, J. (1996). Color contrast under controlled chromatic adaptation reveals opponent rectification. *Vision Research*, *36*(19), 3087-3105.

Speigle, J. M., & Brainard, D. H. (1999). Predicting color from gray: The relationship between achromatic adjustment and asymmetric matching. *Journal of the Optical Society of America A*, *10*, 2370-2376.

Spillmann, L. (1994). The Hermann grid illusion: a tool for studying human perspective field organization. *Perception*, *23*(6), 691-708.

Stiles, W., & Burch, J. M. (1955). Interim report to the Commission Internationale de l'Eclairage, Zurich, 1955, on the National Physical Laboratory's investigation of colour-matching. *Optical Acta*, *2*, 168.

Stiles, W. S. (1961). Adaptation, chromatic adaptation, colour transformation. *Anales de la Real Sociedad Española de Fisica y Quimica, Seria A-Fisica*, *57*, 245-271.

Stockman, A., MacLeod, D. I., & Johnson, N. E. (1993). Spectral sensitivities of the human cones. *Journal of the Optical Society of America A*, *10*(12), 2491-2521.

Svaetichin, G. (1956). Spectral response curves from single cones. *Acta Physiologica Scandinavica*, *39*(Suppl. 134), 17-46.

Takasaki, H. (1966). Lightness change of grays induced by change in reflectance of gray background. *Journal of the Optical Society of America*, *56*(4), 504-509.

Teller, D. Y. (1984). Linking propositions. *Vision Research*, *24*(10), 1233-1246.

Tschermak, A. (1903). Über Kontrast und Irradiation. *Ergebnisse der Physiologie*, *2*, 726-798.

Valberg, A. (1975). *Color induction. A study of lateral interactions in human vision*. Ph. D. thesis, Institute of Physics, University of Oslo.

Valberg, A. (2001). Unique hues: an old problem for a new generation. *Vision Research*, *41*(13), 1645-1657.

Valberg, A., & Lange-Malecki, B. (1990). ”Colour constancy” in Mondrian patterns: a partial cancellation of physical chromaticity shifts by simultaneous contrast. *Vision Research*, *30*(3), 371-380.

van Tuijl, H. F. J. M. (1975). A new visual illusion: Neonlike color spreading and complementary color induction between subjective contours. *Acta Psychologica*, *39*, 441-445.

Wachtler, T., & Wehrhahn, C. (1997). The Craik-O'Brien-Cornsweet illusion in colour: quantitative characterisation and comparison with luminance. *Perception*, *26*(11), 1423-30.

Walls, G. L. (1960). Land! Land! *Psychological Bulletin*, *57*, 29-48.

Walraven, J. (1976). Discounting the background - The missing link in the explanation of chromatic induction. *Vision Research*, *16*, 289-295.

Walraven, J. (1979). No addive effect of backgrounds. *Vision Research*, *19*, 1061-1063.

Walraven, J., Benzschawel, T. L., & Rogowitz, B. E. (1987). Color-constancy interpretation of chromatic induction. *Die Farbe*, *34*, 269-273.

Walraven, J., & Werner, J. S. (1991). The invariance of unique white; a possible implication for normalizing cone action spectra. *Vision Research*, *31*(12), 2185-2193.

Watanabe, T., & Cavanagh, P. (1993). Transparent surfaces defined by implicit X junctions. *Vision Research*, *33*(16), 2339-2346.

Webster, M. A. (2003). Light adaptation, contrast adaptation, and human color vision. In R. Mausfeld & D. Heyer (Eds.), *Colour perception: Mind and the physical world* (p. 67-110). Oxford University Press.

Webster, M. A., Malkoc, G., Bilson, A. C., & Webster, S. M. (2002). Color contrast and contextual influences on color appearance. *Journal of Vision*, *2*(6), 505-519, http://journalofvision.org/2/6/7/, DOI: 10.1167/2.6.7.

Webster, M. A., Miyahara, E., Malkoc, G., & Raker, V. E. (2000). Variations in normal color vision. II. Unique hues. *Journal of the Optical Society of America A*, *17*(9), 1545-1555.

Webster, M. A., & Wilson, J. A. (2000). Interactions between chromatic adaptation and contrast adaptation in color appearance. *Vision Research*, *40*(28), 3801-3816.

Wendt, G. (2003). *Der Einfluß der Umfeldvarianz auf die Wahrnehmung chromatischer Infelder*. Master's thesis, Institut für Psychologie, Christian-Albrechts-Universität zu Kiel.

Wendt, G., Faul, F., & Ekroll, V. (submitted). Colour appearance: The influence of uniformity and variance in the surround. *Vision Research*.

Werner, J. S., & Walraven, J. (1982). Effect of chromatic adaptation on the achromatic locus: the role of contrast, luminance and background color. *Vision Research*, *22*, 929-943.

Werner, J. S., & Wooten, B. R. (1979). Opponent chromatic mechanisms: relation to photopigments and hue naming. *Journal of the Optical Society of America*, *69*(3), 422-34.

West, G., & Brill, M. H. (1982). Necessary and sufficient conditions for von Kries chromatic adaptation to give colour constancy. *Journal of Mathematical Biology*, *15*, 249-258.

White, M. (1979). A new effect of pattern on perceived lightness. *Perception*, *8*, 413-416.

Whittle, P. (1992). Brightness, discriminability and the "crispening effect". *Vision Research*, *32*(8), 1493-507.

Whittle, P. (1994a). Contrast brightness and ordinary seeing. In A. L. Gilchrist (Ed.), *Lightness, brightness and transparency* (p. 111-158). Lawrence Erlbaum Associates.

Whittle, P. (1994b). The psychophysics of contrast brightness. In A. L. Gilchrist (Ed.), *Lightness, brightness and transparency* (p. 35-110). Lawrence Erlbaum Associates.

Whittle, P. (2003). Contrast colours. In R. Mausfeld & D. Heyer (Eds.), *Colour perception: Mind and the physical world* (p. 115-138). Oxford University Press.

Whittle, P., & Challands, P. D. C. (1969). The effect of background luminance on the brightness of flashes. *Vision Research*, *9*, 1095-1110.

Wilson, M. H., & Brocklebank, R. W. (1955). Complementary hues of after-images. *Journal of the Optical Society of America*, *45*, 293-299.

Wright, W. D. (1972). Colour mixture. In D. Jameson & L. M. Hurvich (Eds.), *Handbook of sensory physiology* (Vol. VII/4 Visual Psychophysics, p. 434-454). Berlin: Springer-Verlag.

Wuerger, S. M. (1996). Color appearance changes resulting from iso-luminant chromatic adaptation. *Vision Research*, *36*(19), 3107-3118.

Wyszecki, G., & Stiles, W. (1982). *Color science: Concepts and methods, quantitative data and formulae* (2nd ed.). John Wiley and Sons.

Young, T. (1802). The Bakerian lecture. On the theory of light and colours. *Philosophical Transactions of the Royal Society of London*, *92*, 12-48.

Appendix A

Colour Plates

Colour Plate I: Simultaneous colour contrast. The central squares are physically identical, yet they appear rather different, since they are embedded in different surrounds.

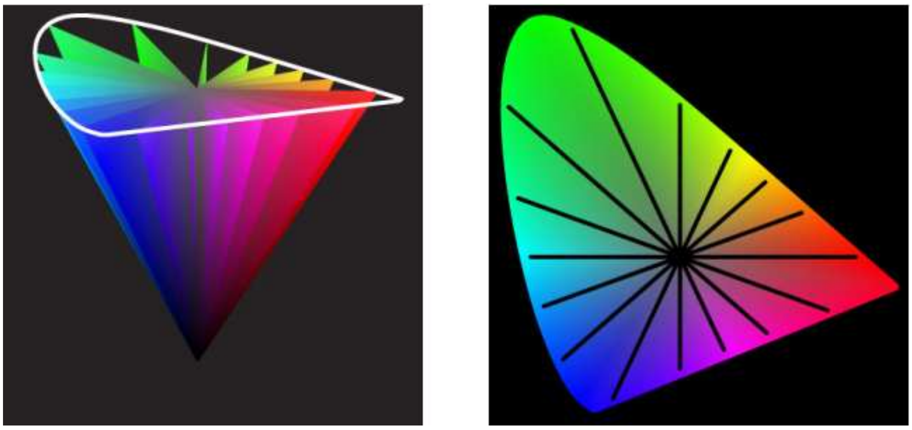

Colour Plate II: Schematic illustration of how points in colour space are thought to correspond to variables of perceived colour. Left: Planes of constant hue meet at the central ‘achromatic axis’. Brightness corresponds grossly to the ‘height’ in this three-dimensional representation. Left: In a plane of constant brightness (luminance), lines of constant hue converge on the achromatic point. Saturation it thought to correspond to the radial distance from the achromatic point.

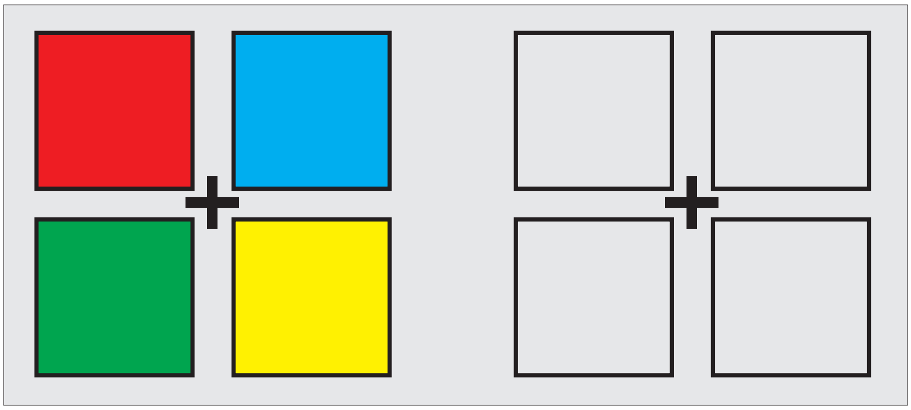

Colour Plate III: Afterimages. Stare at the central cross in the left-hand configuration for about ten seconds, and then look at the right-hand fixation cross. You should now experience fairly vivid afterimages of the coloured squares.

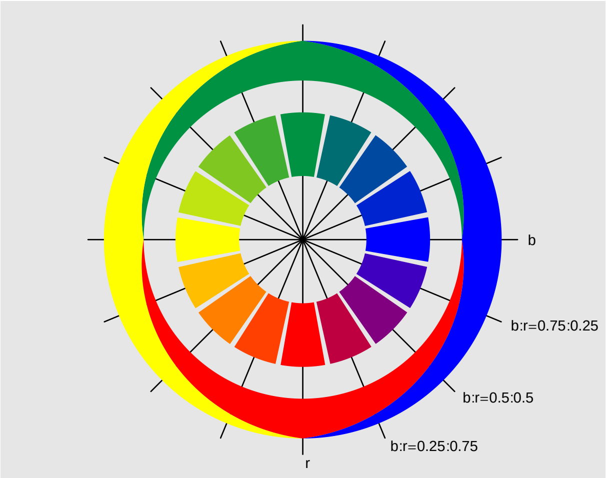

Colour Plate IV: Illustration of how all hues (inner circle) can be conceived of as perceptual mixtures of the four unique hues blue, green, yellow, and red (outer circle). A prototypical violet, for instance, contains equal amounts of blue and red. Adapted after Hering (1920).

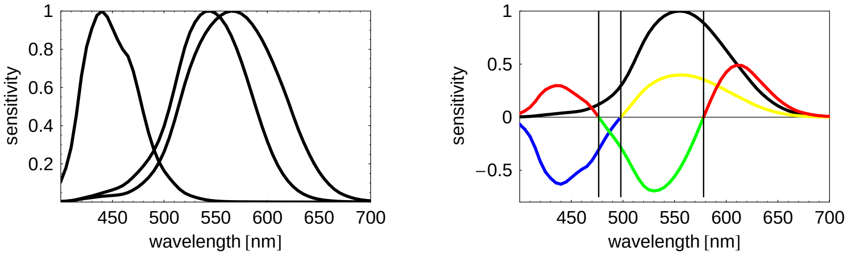

Colour Plate V: Left: The spectral sensitivity curves of the cones. Right: The spectral sensitivity curves, also called valence-curves, of the hypothetical opponent mechanisms red-green, blue-yellow and black-white. Note that whereas the cone sensitivities are all positive, each of the two colour-opponent sensitivity curves have negative and positive lobes. The opponent sensitivity function drawn in black represents the black-white and is assumed to be identical to the spectral sensitivity curve $V(\lambda)$ of the luminance mechanism. All three opponent sensitivity curves are just linear sums of cone senstivity curves on the left. To highlight the fact that they are thought to represent perceived colour in a simple way, they are drawn here in color. The vertical lines mark the zero-crossings of the colour-opponent mechanisms, i.e. the spectral loci of unique blue, unique green and unique yellow (from left to right). As is generally accepted, no monochromatic stimulus evokes the impression of unique red. Instead, unique red is evoked by a binary mixture of the monochromatic stimuli at the endpoints of the spectrum.

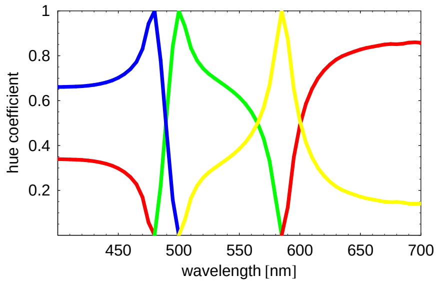

Colour Plate VI: Values of the hue coefficients of opponent colour theory for monochromatic lights across the visual spectrum. These curves are computed based on the valance curves in Color Plate V and can be read in the following way: At any wavelength, there are two hue responses, except at the loci of the unique hues, where there is only one hue response. The curves give the amount of one of the hue responses divided by the total hue response. Accordingly, the hue response is 1 at the loci of the unique hues (unique red is extraspectral), and between 0 and 1 for all other wavelengths. Note that the two curves are symmetric by definition: For instance, red/(red+blue)=1-blue/(red+blue). The points at which the two curves cross, i.e. where the hue efficients equals 0.5, correspond to to the binary hues turquoise, chartreuse and orange (from left to right). The binary hue violet is extraspectral.

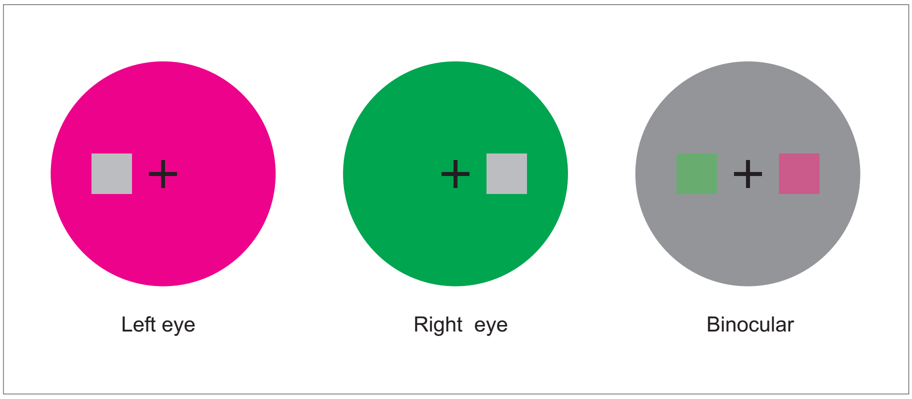

Colour Plate VII: The experimental technique of haploscopically superimposed displays (after Whittle, 1994b). The left configuration is presented to the left eye, the one in the middle to the right eye. The right panel gives an impression of what the observer then perceives. The square targets in the left and middle panels are physically identical, yet appear somewhat different due to simultaneous contrast. This effect is greatly enhanced in the binocularly fused image. The reader may try to verify this by free-fusing the left and middle panels.

Colour Plate VIII: The zero-contrast problem: The contrast-coding hypothesis predicts that the two central regions should appear equal in colour since they have the same (zero) contrast with their surrounds (The dashed contours are merely drawn to indicate where the target squares would be visible if they were differently colored than their surrounds).

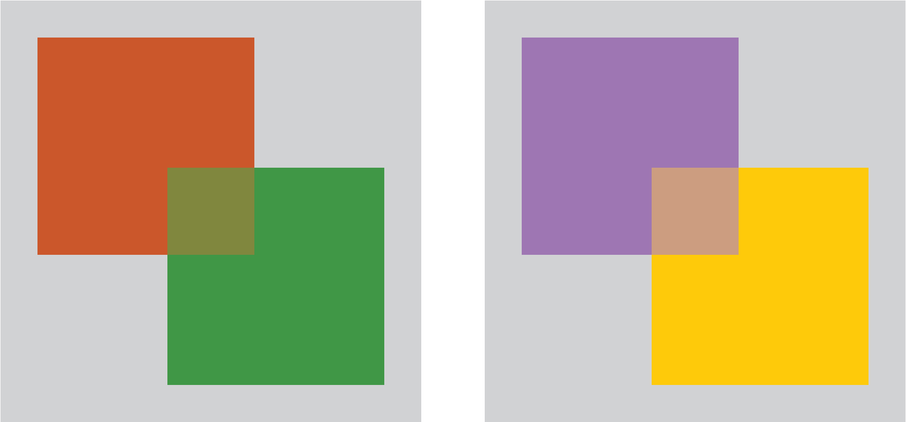

Colour Plate IX: Two four-region stimuli of the kind traditionally used in investigations of perceptual transparency. In the left configuration, one may have the impression of a greenish transparent filter in front of a red object, and in the overlapping region one therefore has the impressions of two colours coexisiting at the same location of the visual field. The left configuration illustrates the same point with other colours.

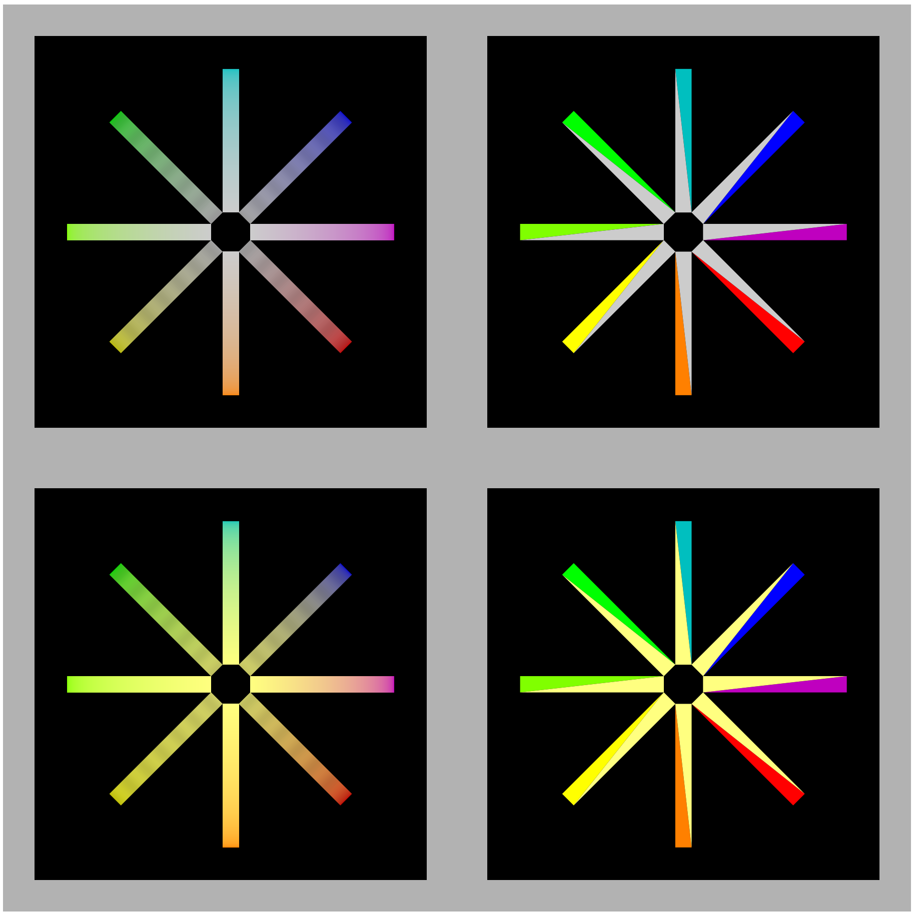

Colour Plate X: Top panels: The spokes on the left representing lines of corresponding hue can, as illustrated in the left-hand panel, each be considered to be a graded mixture of a chromatic colour and a neutral (grey) one. The thickness of the wedges represents the relative amounts of the two colours. Clearly, the admixture of grey does not change the perceive hue. Bottom panels: In the case of a graded mixture of two chromatic colours, however, the hue must of course change.

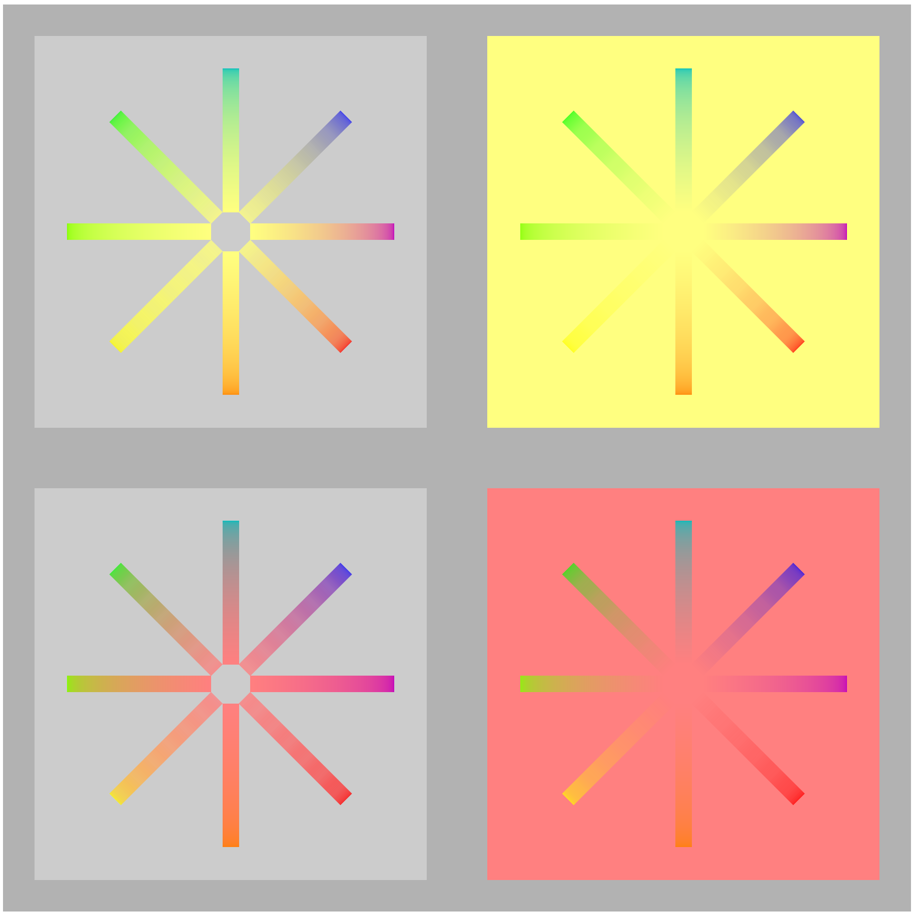

Colour Plate XI: The spokes in the upper left panel are shaded towards yellow in the middle, and can therefore not be thought of as lines of constant hue. If the same set of spokes is embedded in a yellow surround, though, they appear to be of roughly constant hue, and as the contrast between the spokes and the surround decreases towards the centre one has the impression of yellowish transparent haze covering the spokes. In the lower panels, the analogous case is shown for spokes shaded towards red and a red surround. Note that the left sets of spokes are identical to the corresponding right ones. Perceptually, though, the two sets of spokes on the right side appear much more similar, although they are just as different as the two left-hand ones. The main perceptual difference between the two is the colour of the haze appearing to cover the spokes. Note that the intended colours in this demonstration may be incorrectly reproduced in print.

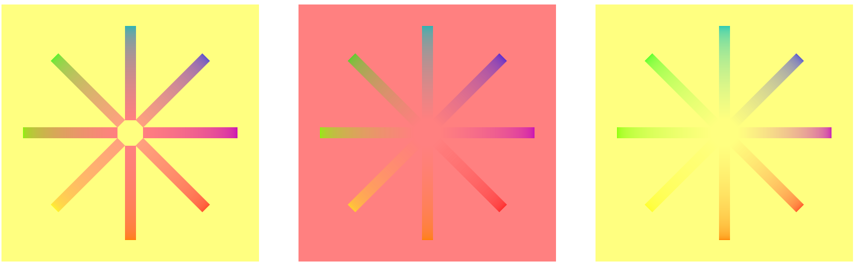

Colour Plate XII: The two left hand sets of spokes are identical (they are both shaded towards red), but look rather different since they are embedded in different surrounds. The right-hand set of spokes are shaded towards yellow instead, and look very similar to the set of spokes in the middle, which are shaded towards red. Note that the main remaining perceptual difference between the two leftmost sets of spokes is the reddish and the yellowish transparent haze. Changing the colours of the spokes further in order to make them appear exactly equal would be futile, since the colour of the haze is a property determined by the surround, and thus cannot be changed. Note that the intended colours in this demonstration may be incorrectly reproduced in print.

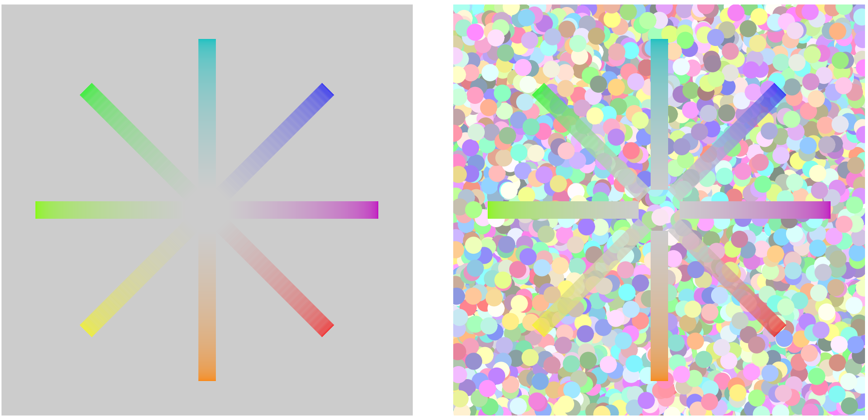

Colour Plate XIII: The two sets of spokes are identical, but appear rather different. Note that the effect is most impressive at low contrast (towards the centre), where one has impressions of transparency in the uniform surround. Note also that the spokes in the variegated surround get arbitrarily desaturated (similar to gray) towards the centre, while all of the greyness in the left-hand spokes appear to belong to the transparent haze, leaving the spokes themselves appear rather saturated even towards the centre. Note that the intended colours in this demonstration may be incorrectly reproduced in print.

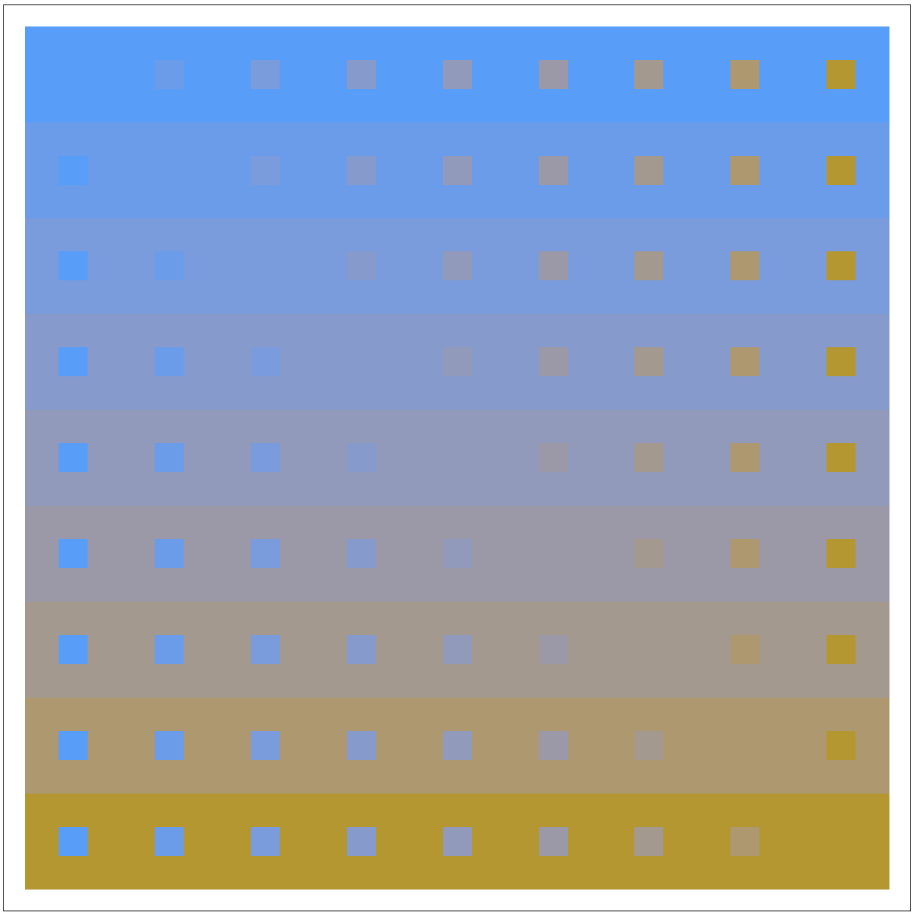

Colour Plate XIV: Within each coloum, the small squares are physically identical, that is, in each row actually the same scale of colours from blue to yellow is reproduced. Note that all of the targets above the diagonal, which are decrements with respect to the blue-yellow mechanism appear yellowish, while those below the diagonal, which are increments, appear bluish. Note also that the upper scales, which are embedded in blue surrounds, contain very desaturated yellowish targets, particularly in the most saturated blue surround. The scales that are embedded in the yellowish surrounds, on the other hand, contain no yellowish targets of very low saturation. In exchange, though, they contain more desaturated bluish targets than the scales embedded in the blue surrounds. Note that the intended colours in this demonstration may be incorrectly reproduced in print.

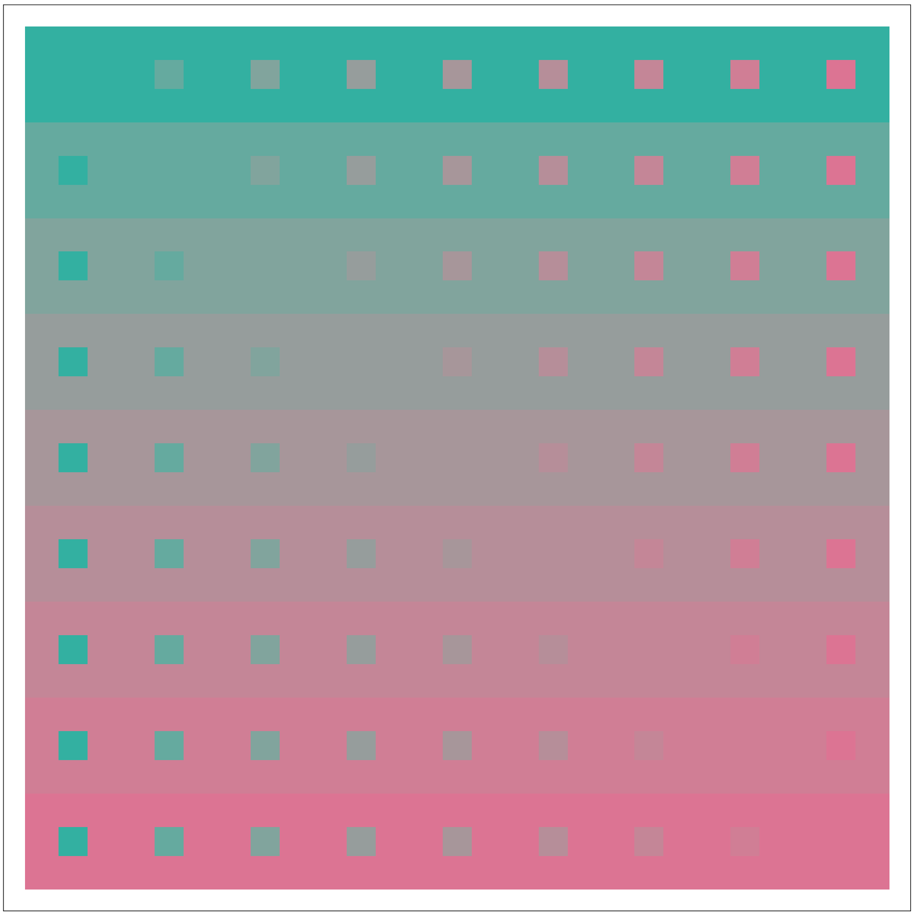

Colour Plate XV: Within each coloum, the small squares are physically identical. See caption of Colour Plate XIV for explanations, which apply *mutatis mutandis*.

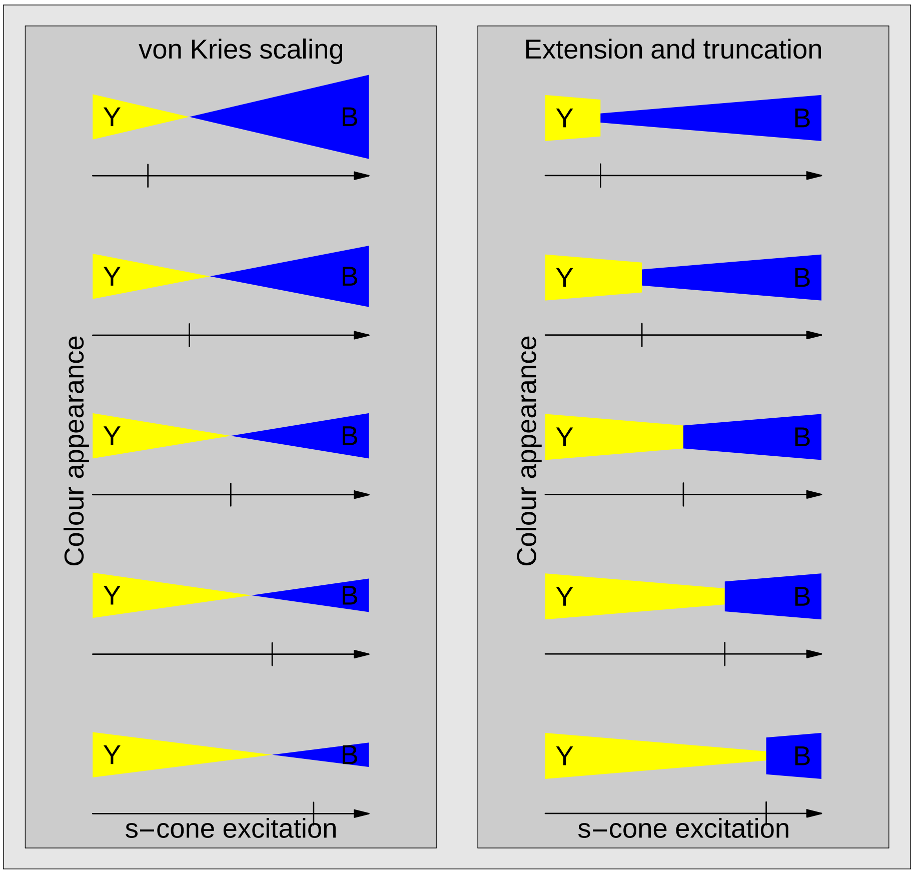

Colour Plate XVI: Classical (left) vs. truncation-extension (right) view of how colour impressions should vary along a line from yellow to blue in chromaticity space for patches embedded in surrounds of different colours. In each sub-panel a pair of wedges represents the perceived colours of the opponent hues blue and yellow. The thicker the wedge the more saturated the colour impression. Below each pair of wedges a scale representing S cone excitation is shown with a vertical tick mark at the location of the surround. From top to bottom the illustrations refer to yellow, neutral and blue surrounds. A crucial difference between the classical and the truncation-extension view is that, in the latter case, the saturation never approaches zero.

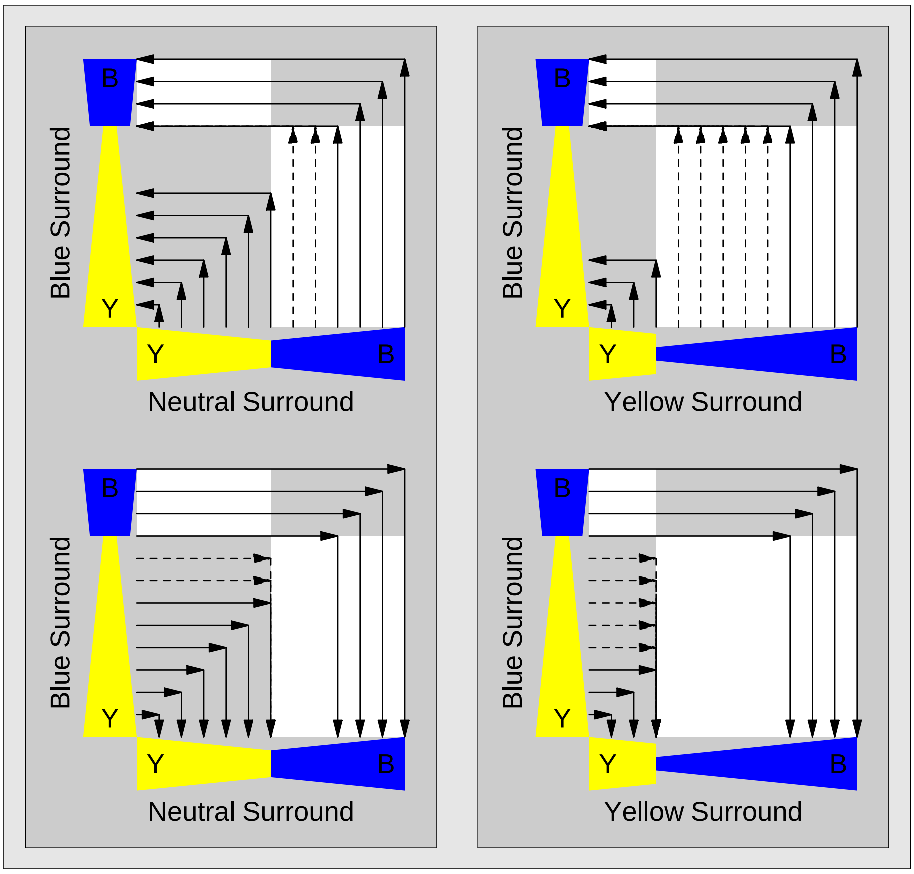

Colour Plate XVII: Top left panel: The matches which may be expected based on extension and truncation when fixed patches are presented in neutral surrounds, and the settings are made in a blue surround. Lower left panel: Expected matches when the fixed patches are presented in the blue surround. Right-hand panels: Analogous expectations derived for matches made with a yellow and a blue surround.

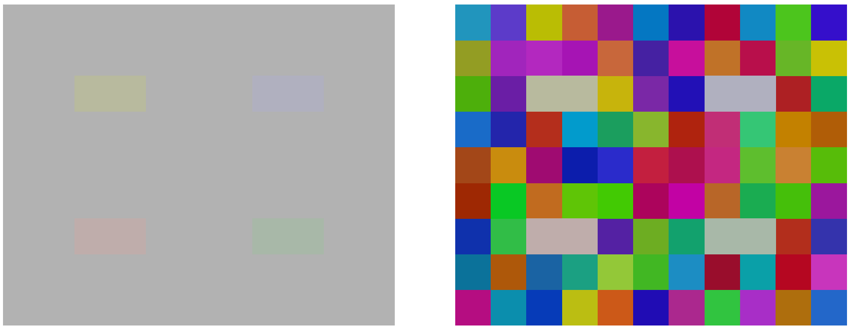

Colour Plate XVIII: The gamut expansion effect of Brown and MacLeod (1997). The four targets in the uniform surround are physically identical to the corresponding four targets in the variegated surround, but the former appear more saturated than the latter. Note that the intended colours in this demonstration may be incorrectly reproduced in print.

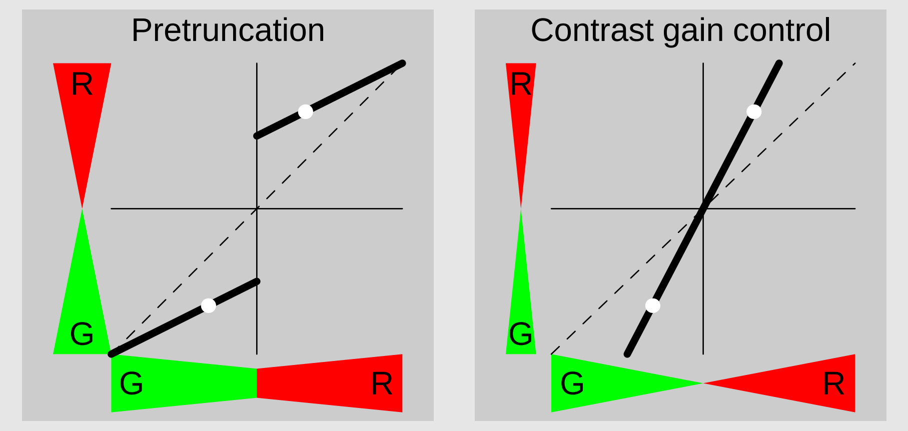

Colour Plate XIX: Predictions based on a simple gain control model (right) and ’pretruncation’ (left). The wedges on the horizontal axis represent perceived saturation for central patches with chromaticities from the red-green axis in chromaticity space when presented in a uniform neutral surround. In the case of pretruncation the wedges would have their tips cut of (left panel). The wedges on the vertical axis represent the perceived saturation for the same central patches when embedded in a variegated neutral surround. The different steepness of the horizontal and vertical wedges in the right panel represents gamut expansion. The locations of expected matches are drawn as thick lines. The white dots demonstrate that the result of measuring the effect for only one purity of the test patch, as Brown and MacLeod (1997) did, is compatible with either hypothesis.

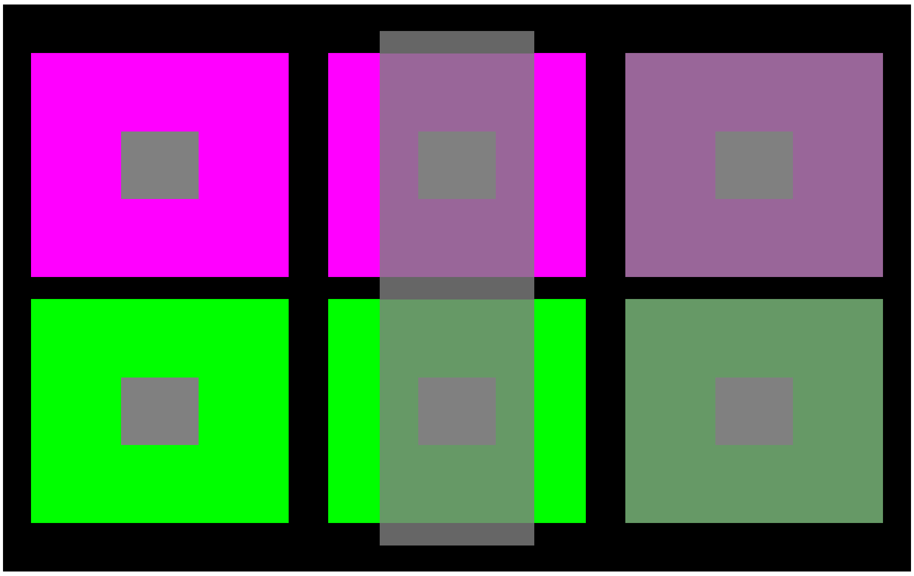

Colour Plate XX: Meyer's effect. The leftmost pair of centre-surround stimuli constitute a classical demonstration of simultaneous contrast, where the targets are printed in the same gray colour and the surrounds have saturated complementary colours. In the middle, the same stimulus pair is (partially) covered by a gray transparent layer, as simulated by using a standard transparency model. Since the colour of the transparent layer was chosen to be the same gray as the targets, the latter were left unchanged by this operation. The saturation of the immediate surrounds, however is substantially reduced. In spite of this, the perceived difference in colour between the upper and lower targets is equally convincing as in the leftmost display, perhaps even a bit more impressive. This observation constitutes Meyer's effect. In the right-most display, two centre-surround stimuli with exactly the same immeditate surrounds as in the middle panel are shown. Also in this case, the effect is equally convincing as with the more saturated surrounds.

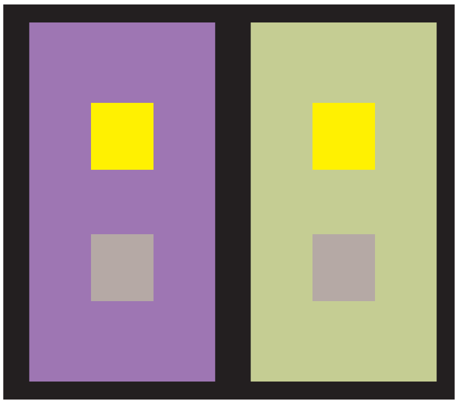

Colour Plate XXI: The two lower squares are mutually identical, as are the lower ones. The perceived difference between the latter is however much more impressive than that of the upper ones. Hence, even when the surrounds are fixed, the simultaneous contrast effect may be strong or weak, depending on the colour of the target.

Appendix B

Zusammenfassung (German abstract)

Einführung

Die zwei mittleren Quadrate in Farbtafel I (S. 137) sind physikalisch gleich. Da sie in unterschiedlich gefärbte Umfelder eingebettet sind, scheinen sie aber trotzdem unterschiedlich eingefärbt zu sein. Dieses seit der Antike bekannte Phänomen – der simultane Farbkontrast – war, insbesondere in den letzten zwei Jahrhunderten, Gegenstand einer beeindruckenden Anzahl empirischer und theoretischer Untersuchungen (Tschermak, 1903; Whittle, 2003). Die Anlässe für diese Forschungsarbeiten waren vielfältig, und reichten von rein praktischen Belangen bis hin zu philosophischen Fragen. In den wissenschaftlichen Diziplinen, die sich mit der Funktionsweise des Gesichtssinns beschäftigen, entspringt die Hauptmotivation der grundsätzlichen Vorstellung, dass die Untersuchung solcher Sinnestäuschungen einen Schlüssel zum Verständnis derjenigen Mechanismen bietet, die die mageren und mehrdeutigen Reizmuster auf der Netzhaut in nützliche und verlässliche Informationen über die Außenwelt umwandeln (Mach, 1866; Hoffman, 1998; Eagleman, 2001). Dieser Vorstellung entsprechend wurde schon früh ein enger Zusammenhang zwischen dem simultanen Farbkontrast und biologisch wichtigen Mechanismen der Farbkonstanz vermutet. Laut Herings (1920) klassischer Theorie, beispielsweise, beruhen sowohl die Farbkonstanz als auch der Simultankontrast auf Wechselwirkungen zwischen benachbarten Rezeptoren auf der Netzhaut, d.h. auf dem neuronalen Mechanismus der lateralen Hemmung (Hartline, Wagner, & Ratliff, 1956). Auch Helmholtz (1911) ging davon aus, dass der Simultankontrast und die Farbkonstanz auf gemeinsame Mechanismen zurückzuführen seien. Anders als Hering führte Helmholtz beide Phänomene auf eine unbewusste schlussfolgernde Tätigkeit des Gehirns zurück, wobei diese unbewussten Schlüsse auf der Basis früherer Erfahrung mit den Regularitäten der Umwelt gezogen würden. Der Simultankontrast ist nach Helmholtz darauf zurückzuführen, dass das Wahrnehmungssystem durch das farbige Umfeld zu einer unrealistischen Annahme über die Beleuchtungsverhältnisse verleitet wird. Da das Wahrnehmungssystem auf die Eigenfarbe von Gegenständen dadurch schließe, dass es den Einfluss der Beleuchtung herausrechne, führe diese fehlerhaft Schätzung der Beleuchtung zu einer verfälschten Schätzung der Eigenfarbe des Gegenstandes.

Die Theorien von Hering und Helmholtz, von denen die Mehrheit der heutigen Theorien und Modelle wesentliche Aspekte übernommen hat, unterscheiden sich in vielen Punkten, über die heute noch diskutiert wird (Kingdom, 1997, 2003). Hering hob die Rolle von einfachen, angeborenen physiologischen Mechanismen hervor, während Helmholtz’ Theorie die Bedeutung von intelligenten, flexiblen und lernfähigen Mechanismen unterstrich. Einig waren sich Hering und Helmholtz jedoch in einem wesentlichen Punkt, nämlich dass es sich beim Sehen „nicht um ein Schauen der Strahlungen als solcher [...] handelt, sondern um das durch diese Strahlungen vermittelte Schauen der Außendinge; das Auge hat uns nicht über die jeweilige Intensität oder Qualität des von den Außendingen kommenden Lichtes, sondern über diese Dinge selbst zu unterrichten” (Hering, 1920, S. 13)[1], und dass die Mechanismen, die uns dieses ’Schauen der Außendinge’ ermöglichen, seien es laterale Hemmung oder unbewusste Schlüsse, auch für den Simultankontrast verantwortlich seien.

Vor allem aufgrund der offensichtlichen biologischen Wichtigkeit der Farbkonstanzmechanismen hat sich die Forschung zu diesem Themengebiet seit den wegweisenden Arbeiten von Katz (1911) und Gelb (1929) zu einem breit angelegten Unterfangen entwickelt, mit zahlreichen Beiträgen aus vielen unterschiedlichen Disziplinen wie etwa Sinnesphysiologie, Wahrnehmungspsychologie und der Künstlichen Intelligenz. Dies hat auch aufgrund des angenommenen Zusammenhangs zwischen dem Simultankontrast und der Farbkonstanz, zur Folge gehabt, dass inzwischen eine beeindruckende Fülle von Forschungsergebnissen zum Simultankontrast vorliegt. In psychophysikalischen Untersuchungen geht die Parallelität zwischen Simultankontrast und Farbkonstanz sogar so weit, dass Experimente, die durchgeführt wurden um Farbkonstanz zu untersuchen, meist nicht zu unterscheiden sind von Experimenten, die der Erforschung des Simultankontrasts dienen sollten und *vice versa*: Sowohl der Versuchsaufbau, die Reize, die dargeboten werden, als auch die Aufgaben der Versuchspersonen sind meist dieselben oder zumindest sehr ähnlich.

Der schieren Menge des vorliegenden Datenmaterials zum Trotz sucht man in der Literatur vergeblich nach einem quantitativen Modell des Simultankontrastes mit zufriedenstellender Vorhersagekraft und mit Allgemeinheitsanspruch. Obwohl bedeutsame Einsichten gewonnen wurden, ist unser wissenchaftliches Verständnis und die technologische Beherrschung des Simultankontrastes heute noch bestenfalls als lückenhaft zu bezeichnen: Die fundierten und umfassenden Übersichtsarbeiten von Whittle (1994b, 1994a, 2003) vermitteln einen Eindruck davon, welchen Schwierigkeiten man sich bei dem Versuch, aus den vielen Einzelbefunden ein kohärentes und allgemein gültiges Modell aufzustellen gegenübergestellt sieht (siehe auch Mausfeld & Niederée, 1993). Die Gründe für diese Schwierigkeiten sind ohne Zweifel vielfältig. Es ist zum Beispiel seit langem bekannt, dass der Simultankontrasteffekt von einer Vielzahl von Einflussgrößen abhängt, wie etwa Beobachtungsbedingungen, Reizgrößen, perzeptuelle Organisation sowie Urteilskriterien und Instruktion der Beobachter. In der vorliegenden Dissertation werden jedoch grundlegendere, konzeptuelle Probleme, mit denen klassische Modelle des Simultankontrastes behaftet sind, in Angriff genommen.

Klassische Modelle des Simultankontrastes

Als Ausgangspunkt für die vorliegenden psychophysikalischen Untersuchungen werden in Kapitel 3 einige einflussreiche und als klassisch geltende Modelle des Simultankontrastes vorgestellt und deren Unterschiede und Gemeinsamkeiten aufgezeigt. Dabei werden das von Kriessche ‘Umstimmungsmodell der Photorezeptoren’ (Kries, 1905), das im Rahmen der opponenten Farbtheorie formulierte Zwei-Prozess-Modell von Jameson und Hurvich (1959) sowie eine Klasse von sogenannten Kontrastkodierungsmodellen diskutiert. Letztere gehen davon aus, dass Farbe nicht anhand des lokalen absoluten Farbsignals kodiert wird, sondern grundsätzlich durch die *Differenzen* von auf der Netzhaut benachbarten Farbsignalen (Whittle & Challands, 1969; Walraven, 1976).

Aufgrund einer Vielzahl negativer experimentelle Befunde gilt das von Kriessche Modell seit längerem als unzureichend für die Erklärung des Simultankontrastes, wird aber trotzdem berücksichtigt, unter anderem weil es nicht ausgeschlossen werden kann, dass ein Mechanismus dieser Art im gesamten Prozess mitwirkt: In den meisten anderen Modellen wird eine Art von Kriessche

---

[1] „Ihre wichtigste Bedeutung haben die Farben für uns, insofern sie Eigenschaften der Körper sind, und als Erkennungszeichen der Körper benutzt werden können. Wir gehen deshalb bei unseren Beobachtungen mit dem Gesichtssinne stets darauf aus, uns ein Urteil über die Körperfarben zu bilden, und dabei die Verschiedenheiten der Beleuchtung, unter der sich ein Körper uns darbietet, zu eliminieren” (Helmholtz, 1911, S. 243).

Rezeptor-Umstimmung als Teilprozess vorausgesetzt, und die Ergebnisse der vorliegenden Experimente deuten in der Tat darauf hin, dass diese Annahme nicht gänzlich unrealistisch ist. Historisch gesehen wurde das Zwei-Prozess-Modell von Jameson und Hurvich als Alternative zu dem empirisch stark geschwächten von Kries Modell vorgeschlagen. Während das von Kries-Modell ausschließlich eine multiplikative Reskalierung des retinalen Farbsignals vorsieht, geht das Zwei-Prozess-Modell davon aus, dass zusätzlich auch eine subtraktive Umkodierung stattfindet. Auf der Ebene der Photorezeptoren findet zwar eine Umstimmung im von Kriesschen Sinne statt, auf der nachgeschalteten, sogenannten opponenten Ebene der Farbverarbeitung des visuellen Systems werde jedoch zusätzlich ein Anteil des Farbsignals abgezogen, der dem Farbsignal in unmittelbar benachbarten Regionen der Netzhaut direkt proportional sei. Jameson und Hurvich wiesen darauf hin, dass durch die Annahme einer zusätzlichen subtraktiven Umkodierung des Farbsignals experimentelle Befunde erklärt werden können, die mit dem reinen von Kries-Modell unvereinbar wären. Umgekehrt sei die Annahme einer rein subtraktiven Umkodierung aber auch nicht mit den Befunden in Einklang zu bringen, so dass die Annahme eines Zusammenspiels von sowohl multiplikativer als auch subtraktiver Umkodierung naheliegend erscheine.

Eine noch modernere Klasse von Modellen mit tiefgreifenden theoretischen Implikationen werden in der vorliegenden Arbeit als Kontrastkodierungsmodelle bezeichnet. Diese Modelle können als Weiterentwicklungen, aber vor allem auch als noch genauer spezifizierte Submodelle des Zwei-Prozess-Modells aufgefasst werden. Wesentlich für diese Modelle ist die Annahme, dass der subtraktive Term im Zwei-Prozess-Modell nicht nur irgendwie proportional zu dem Farbsignal des Umfeldes, sondern sogar damit *identisch* ist (Whittle & Challands, 1969; Walraven, 1976). Die konzeptuellen Probleme, die sich durch die mit diesen Modellen verbundenen experimentellen Befunde und theoretischen Deutungen ergeben, bilden den unmittelbaren Ausgangspunkt für die vorliegenden Untersuchungen. Einerseits erscheint die radikale Annahme der Kontrastkodierungsmodelle, dass der Farbeindruck grundsätzlich von der Differenz zwischen den Farbsignalen von benachbarten Stellen des Gesichtsfeldes bestimmt werden, aus vielerlei Gründen unplausibel, andererseits ist die experimentelle Evidenz dafür so umfassend, dass sie fast unausweichlich erscheint (Whittle, 2003). Entsprechend unklar gestaltet sich der allgemeine theoretische Umgang mit dieser Vorstellung in der einschlägigen Literatur. Einerseits ist die grundlegende Idee hinter den Kontrastkodierungsmodellen in der Neurophysiologie und der Psychologie so etabliert, dass sie inzwischen fast als ein Gemeinplatz verstanden wird, andererseits wird sie bei konkreter Modellbildung oft völlig außen vor gelassen. Wie Whittle (2003) etwas sarkastisch, aber nicht gänzlich unbegründet bemängelt, könnte man leicht zu der Überzeugung gelangen, die Kontrastkodierungsmodelle würden nur dann herangezogen, wenn es einem genehm erscheint, und bei der Behandlung von Befunden, die damit nicht unmittelbar zu erklären sind, ohne weiteres außer Acht gelassen. Als Alternative zu einem unverbindlichen und opportunistisch-inkonsequenten Umgang mit der Hypothese der Kontrastkodierung steht in der Literatur ein eleganter theoretischer Rahmen zur Verfügung, innerhalb dessen man sowohl der empirischen Evidenz für Kontrastkodierung als auch einer Reihe damit scheinbar im Widerspruch stehenden Beobachtungen gerecht werden kann. Laut dieser, in allgemeiner Form verbreiteten und von Whittle (1994a) und Arend (1973) explizit formulierten Perspektive, ist Kontrastkodierung eine fundamentale Eigenschaft der initialen Farbkodierung in der Netzhaut, während in höheren visuellen Zentren und dem Gehirn Verarbeitungsschritte stattfinden, die diese initiale Kodierung unter gewissen Umständen rückgängig machen können. In enger Anlehnung an die komplementären mathematischen Operationen der Ableitung und der Integration spricht man davon, dass die Netzhaut Differenzen bildet, während das Gehirn diese Differenzsignale aufsummiert: „The retina differentiates; the brain can integrate" (Whittle & Challands, 1969, p. 1109). In Übereinstimmung mit den aus dieser Perspektive ableitbaren Erwartungen findet man unter speziellen experimentellen Bedingungen, bei denen die postulierten Integrationsprozesse als außer Kraft gesetzt betrachtet werden können, sehr starke Evidenz für reine Kontrastkodierung, während man unter anderen, alltäglicheren Bedingungen, bei denen die Integrationsprozesse unbeeinträchtigt am Werke sein sollten, feststellen kann, dass die mit der reinen Kontrastkodierung einhergehenden starken Simultankontrasteffekte in der Regel aufgehoben oder zumindest deutlich abgeschwächt sind.

In der vorliegenden Arbeit wird jedoch eine Schwäche dieser sonst sehr fundierten und befriedigenden Perspektive identifiziert und eingehenderen Untersuchungen unterzogen. Nach der Kontrastkodierung-plus-Integrations-Perspektive wird man nämlich dem von Walraven (1976) berichteten Befund nicht ohne weiteres gerecht. Walravens Experiment liefert auf der einen Seite starke Hinweise auf Kontrastkodierung, allerdings unter Bedingungen, bei denen man nicht erwarten würde, dass die Integration irgendwie beeinträchtigt sein sollte. Bei Kontrastkodierung ohne Integration ergibt sich jedoch eine paradoxe Vorhersage (Whittle, 1994a; Mausfeld & Niederée, 1993): Man würde nämlich auf der Basis der opponenten Farbtheorie erwarten, dass ein Infeld, das zu seinem Umfeld identisch ist, immer achromatisch erscheinen sollte, unabhängig davon, welche Farbe das Umfeld hat. Dass diese Vorhersage nicht zutreffen kann, ist offensichtlich, denn in diesem Falle erscheint das Infeld natürlich in derselben Farbe wie das Umfeld.

Vor diesem Hintergrund mag es beruhigend erscheinen, dass Walravens Modell nicht unangefochten blieb. Fast unmittelbar nach der Veröffentlichung von Walravens (1976) Befunden stellte Shevell (1978) Ergebnisse vor, die darauf hindeuteten, dass Walravens Modell zwar die Tatsachen in guter Näherung beschreibt, aber eben nicht ganz exakt. Sowohl Walravens Modell als auch das von Shevell vertretene entsprechen dem Zwei-Prozess-Modell, wobei der subtraktive Term nach Walraven genau den Umfeldkoordinaten entspricht, nach Shevell hingegen aber zwar einem großen aber eben nicht hundertprozentigen Anteil davon. Gemäß der theoretischen Bedeutung der Frage lösten die widersprüchlichen Befunde von Walraven und Shevell eine lebhafte Kontroverse aus, in deren Verlauf eine Vielzahl weiterer Befunde aus sorgfältig geplanten Experimenten vorgestellt wurden. Es erwies sich jedoch als äußerst schwierig, die Streitfrage eindeutig zu klären. Eine Schwierigkeit bestand darin, dass die sich die Vorhersagen der beiden Modelle in absoluten Zahlen nur sehr gering unterscheiden, so dass sehr exakte Messungen vonnöten waren. Augrund der formalen Struktur und der freien Parameter der Modelle kam noch hinzu, dass nur die Untersuchung von Reizen mit sehr geringen Kontrasten zwischen Infeld und Umfeld diagnostische Daten erhoffen ließen.

Um eine Klärung herbeizuführen, wurden entsprechend Experimente durchgeführt, in denen diese Randbedingungen erfüllt waren (Walraven, 1979; Drum, 1981). Aus diesen Arbeiten geht jedoch hervor, dass die Versuchspersonen Schwierigkeiten hatten, ihre Urteile gemäß der Instruktion zu fällen. Ihre Aufgabe bestand darin, die Farbkoordinaten des Infeldes so einzustellen, dass es Urgelb, d.h. weder rötlich noch grünlich erscheint, was nach den Standardannahmen der opponenten Farbtheorie problemlos möglich sein sollte. Dieser Standardannahme widersprechend, deuten jedoch die Beobachtungen einiger Autoren (Helmholtz, 1911; Brenner & Cornelissen, 1991; Mausfeld, 1998) darauf hin, dass unter ähnlichen Bedingungen zwei opponente Farbkomponenten gleichzeitig wahrgenommen werden: Das Infeld wird unter Umständen sowohl rötlich als auch grünlich wahrgenommen. Dass es den Versuchspersonen Schwierigkeiten bereiten sollte, Urgelbeinstellungen zu machen, liegt vor dem Hintergrund dieser Beobachtungen auf der Hand.

In der vorliegenden Arbeit wurde der Hypothese nachgegangen, dass diese Schwierigkeiten auf der Ungültigkeit einer Grundannahme beruhen, die klassische Modelle des Simultankontrastes ihrer formalen Struktur nach implizit als gegeben voraussetzen. Quantitative Modelle des Simultankontrastes gehen traditionellerweise von der naheliegenden und scheinbar selbstverständlichen Annahme aus, dass die durch das Umfeld bewirkten Farbveränderungen im Prinzip von derselben Art sind wie die Farbveränderungen, die durch Änderungen der Infeldkoordinaten hervorgerufen werden. Entsprechend wird davon ausgegangen, dass beide Arten von Farbveränderungen sich gegenseitig *kompensieren* können. Obwohl diese „Kompensationsannahme” in der Literatur selten Gegenstand expliziter Diskussion war, ist sie von grundlegender theoretischer Bedeutung. Die harten Tatsachen der Young-Helmholtzschen Dreifarbentheorie zeigen, dass das menschliche Farbensehen dreidimensional ist, sowohl in einem wohldefinierten psychophysikalischen (Krantz, 1975b) als auch in einem entsprechenden physiologischen Sinne (Sharpe & Stockman, 1999). Wäre der durch einen lokalen Lichtreiz ausgelöste Farbeindruck unabhängig von dem Umfeld, in das er eingebettet ist, so könnte man folgerichtig aus den Tatsachen der Dreifarbentheorie schließen, dass der dem lokalen Lichtreiz entsprechende Farbeindruck auch nur entlang dreier Dimensionen variieren kann. Die allgemeine Kontextabhängigkeit des Farbeindruckes, die unter anderem durch das Phänomen des Simultankontrastes belegt wird, eröffnet jedoch die theoretische Möglichkeit, dass der Farbeindruck entlang mehr als dreier Dimensionen variieren könnte, ganz einfach weil die Kontextabhängigkeit nichts anderes heißt, als dass der Farbeindruck von mehr Rezeptorerregungstripeln abhängt als dem einen Rezeptorerregungstripel des lokalen Reizes (Evans, 1974). Tatsache ist, dass auf der Basis der Dreifarbentheorie allein keineswegs auf eine Trivarianz des Farbeindruckes geschlossen werden kann. Stattdessen folgt dies nur unter der zusätzlichen Annahme, dass die Kompensationsannahme gerechtfertigt ist (W. S. Stiles, 1961).

Die Tatsache, dass klassische Modelle des Simultankontrastes den Farbeindruck durch Zahlen*tripel* zu beschreiben versuchen, zeigt, dass sie die Gültigkeit der Kompensationsannahme voraussetzen. Dass dies jedoch ungerechtfertigt sein könnte, wird durch die Arbeiten mehrerer Autoren nahegelegt. Basierend auf einer Reihe unterschiedlicher Argumente und Beobachtungen wurden etliche Behauptungen aufgestellt, denen zufolge mehr als drei Zahlen notwendig seien, um den Farbeindruck eines Zielreizes angemessen zu charakterisieren (e.g. Katz, 1911; Gelb, 1929; Kanizsa, 1966; Evans, 1964, 1974; Mausfeld, 1998; Niederée, 1998).

Experimente zur Nullpunktproblematik

In Kapitel 4 werden Ergebnisse einer Versuchsreihe vorgestellt, die deutlich für die Richtigkeit dieser Behauptungen sprechen und umgekehrt auf die Ungültigkeit der Kompensationsannahme schließen lassen. In diesen Experimenten werden neu entwickelte, aber dennoch recht einfache Techniken eingesetzt, um die Lage des achromatischen Punktes in der Chromatizitätsebene zu bestimmen. Diese Verfahren können als Abwandlungen einer klassischen Methode zur quantitativen Bestimmung des Simultankontrasteffektes verstanden werden, nämlich der Methode der Graueinstellungen (Helson & Michels, 1948; Speigle & Brainard, 1999). Bei vorgegebener Luminanz (Helligkeit) eines Reizes entsprechen die möglichen Einstellungen für seine Farbkoordinaten einer Ebene (Chromatizitätsebene) im dreidimensionalen Rezeptorerregungsraum, in der es allgemein anerkannten theoretischen Vorstellungen zufolge *einen* ausgezeichneten Punkt gibt, der zu einem achromatischen Farbeindruck führt. Die Lage dieses Punktes hängt von der Farbe des Umfeldes ab, und wird traditionellerweise als Maß für den Simultankontrasteffekt verwendet. Der achromatische Punkt ist jedoch nicht nur dadurch gegenüber allen anderen Punkten in der Chromatizitätsebene ausgezeichnet, dass er achromatisch (weisß, grau oder schwarz, je nach Luminanzstufe) erscheint. Allgemein akzeptierten theoretischen Vorstellungen über den Zusammenhang zwischen Punkten im Rezeptorerregungsraum und Aspekten des Farbeindrucks zufolge, konvergieren Geraden konstanten Farbtons in der Chromatizitätsebene auf den achromatischen Punkt (Konvergenzkriterium). Des weiteren müsste nach diesen Vorstellungen jede geschlossene Kurve in der Chromatizitätsebene, die alle Farbtöne des Farbenkreises enthält, den achromatischen Punkt umschließen (Umschließungskriterium). Die grundlegende Idee der in diesem Abschnitt beschriebenen Experimente besteht darin, den Graupunkt indirekt über diese alternativen Kriterien zu bestimmen und die so erhaltenen Ergebnisse mit dem auf traditionellem Wege ermittelten Graupunkt zu vergleichen. Die Vorstellung, dass diese alternativen Definitionen des achromatischen Punktes equivalent sind, beruht vor allem auf Beobachtungen, die mit isolierten, im Dunkelkontext dargebotenen Reizen gemacht wurden, und dass sie in diesem Falle auch zutreffend ist, kann zweifelsfrei vorausgesetzt werden. Dass sie auch für Reize, die in einem farbigen Umfeld dargeboten werden, zutrifft, wurde bisher nie in Frage gestellt, denn diese Verallgemeinerung folgt unmittelbar aus der traditionellerweise als selbstverständlich erachteten Kompensationsannahme.

Die Ergebnisse der vorliegenden Experimentalreihe deuten jedoch stark darauf hin, dass diese Verallgemeinerung nicht gemacht werden kann, und dass somit die Gültigkeit der Kompensationsannahme in Frage gestellt werden muss. Als durchgängiger Befund stellt sich heraus, das der mit den neuen indirekten Kriterien bestimmte Neutralpunkt im Rahmen der Messgenauigkeit immer exakt mit der Umfeldchromatizität übereinstimmt, während der unter vergleichbaren Bedingungen mit der konventionellen Methode der direkten Graueinstellungen ermittelte Neutralpunkt deutlich und in gesetzmäßiger Weise davon abweicht. In weiteren Varianten des Basisexperiments konnten mögliche Alternativerklärungen dieses Befundes, etwa durch den Abney-Effekt (der bekannte Befund, dass Geraden konstanten Farbtons tatsächlich leicht gekrümmt sind, siehe Burns, Elsner, Pokorny & Smith, 1984) oder durch eine unterschiedliche Komplexität der Reize in den zu vergleichenden Versuchsbedingungen, ausgeschlossen werden.

Der Befund, dass Geraden gleichen Farbtons auf die Chromatizität des farbigen Umfeldes zulaufen statt auf den direkt bestimmten achromatischen Punkt, muss aus der traditionellen Perspektive paradox anmuten, denn ersteres heißt, dass der Konvergenzpunkt für Farben konstanten Farbtons in der bunten Farbe des Umfeldes erscheint, und damit der achromatisch erscheinende Punkt auf einer Geraden gleichen Farbtons liegen muss. Wie in der Arbeit aufgezeigt wird, kann man diesem Befund jedoch in ganz rationaler Weise gerecht werden, wenn man die Kompensationsannahme aufgibt und stattdessen davon ausgeht, dass bei der Betrachtung von einfachen Infeld-Umfeldreizen das sonst nur aus komplexeren Reizkonfigurationen bekannte Phänomen der perzeptuellen Transparenz auftritt, bei dem per Definition zwei Farben an der gleichen Stelle des Gesichtsfeldes wahrgenommen werden (Metelli, 1970; D'Zmura et al., 1997; Adelson, 1993, 2000; Anderson, 1997; Faul & Ekroll, 2002).

Da sich die Kompensationsannahme als unberechtigt herausstellt, erscheint es klar, dass eine darauf basierende Modellbildung die Entwicklung von besseren und adäquateren Modellen nur verzögern kann. Vor diesem Hintergrund erscheint es wenig überraschend, dass frühere Forschungsanstrengungen, die auf der Kompensationsannahme aufgebauten, nur mäßig erfolgreich dabei waren, ein allgemein akzeptiertes und befriedigendes quantitatives Modell zu entwickeln. Es ist zu hoffen, dass es leichter sein wird, Fortschritte zu erzielen, wenn die auf der Kompensationsannahme basierende Perspektive durch einen angemesseneren Verständnisrahmen ersetzt wird. Auch eine Newbewertung bisheriger Modelle und Befunde unter Berücksichtigung der hier gewonnenen Erkenntnisse erscheint notwendig. Der auf der Annahme einer Beteiligung perzeptueller Transparenz aufbauende Verständnisrahmen bietet eine natürliche Erklärung für die inkonsistenten und widersprüchlichen Befunde, die im Laufe der Walraven-Shevell-Kontroverse berichtet wurden (Walraven, 1976, 1979; Shevell, 1978, 1980; Drum, 1981; Adelson, 1981; Davies et al., 1983). Des weiteren erkennt man vor dem Hintergrund dieser Perspektive, dass die schon angesprochene Nullpunktproblematik nicht als ein Reductio-Ad-Absurdum-Argument gegen die Kontrastkodierungsmodelle verstanden werden kann. Vielmehr erscheinen diese Modelle ihrer Grundidee nach in sehr guter Übereinstimmung mit den Tatsachen zu sein, vorausgesetzt, dass man die Kontrastkodierungsmodelle im Sinne der Transparenzperspektive interpretiert statt wie ursprünglich konzipiert im Sinne der Kompensationshypothese. Perzeptuelle Transparenz heißt, dass zwei Farben gleichzeitig an der gleichen Stelle des Gesichtfeldes wahrgenommen werden, und die Kontrastkodierungsmodelle scheinen *eine* von diesen beiden Farbkomponenten sehr zutreffend zu beschreiben.

Auf der Basis der in dieser Versuchsreihe gewonnenen Erkenntnisse wurde zu Demonstrationszwecken eine Reizkonfiguration entwickelt (siehe die Farbtafeln XI und XII auf S. 143 ff.)

bei deren Betrachtung wesentliche Aspekte der experimentellen Befunde auf einem Blick deutlich werden. Zusätzlich verdeutlichen diese Demonstrationen auch einige rein informelle Beobachtungen, die im Laufe der Experimente gemacht wurden. Zu erkennen ist dabei, dass bei geringem Kontrast zum Hintergrund transparenzähnliche duale Farbeindrücke auftreten, wobei die Farbe einer der Komponenten der Farbe des Umfeldes entspricht, und dass in diesen Fällen der Simultankontrast deutlich stärker ist als sonst.

Experimente zur Farbengleichheit

Nicht nur klassische Modelle des Simultankontrastes setzen die Gültigkeit der Kompensationsannahme voraus, sondern auch klassische psychophysikalische Methoden zur quantitativen Bestimmung des Simultankontrasteffektes, wie etwa die Methode der asymmetrischen Farbabgleiche (e.g. Wyszecki & Stiles, 1982; Smith & Pokorny, 1996). Da die Ergebnisse der oben beschribenen Versuchsreihe stark darauf hindeuten, dass man die Kompensationsannahme fallen lassen muss, ergibt sich die Vermutung, dass Daten aus asymmetrischen Farbabgleichexperimenten mit Artefakten behaftet sein können und somit nicht in kanonischer Weise zu interpretieren wären.

In Kapitel 5 der vorliegenden Arbeit werden die Ergebnisse zweier Versuchsreihen vorgestellt, die durchgeführt wurden, um diese Problematik näher zu beleuchten. In beiden Versuchsreihen wurden asymmetrische Farbabgleiche mit gleich-hellen, oder genauer gesagt, äquiluminanten, In- und Umfeldern durchgeführt. Als Chromatizitäten der zwei Umfelder und des vorgegeben Infeldes wurde eine Reihe unterschiedlicher Kombinationen von je einer der zwei sog. Kardinalachsen (Krauskopf et al., 1982) der Chromatizitätsebene gewählt, d.h. von je einer der Hauptachsen der MacLeod-Boynton (1979) Chromatizitätsebene. Die Chromatizität des zweiten Infeldes konnte jedoch von den Versuchspersonen frei eingestellt werden, um einen Farbeindruck herzustellen, der möglichst identisch zu demjenigen des vorgegebenen Infeldes sein sollte. In der ersten Versuchsreihe wurden ausschließlich homogen eingefärbe Umfelder verwendet. Dabei wiesen die resultierenden Datenkurven, in denen die S-, bzw. L-Zapfenerregung des Abgleichsinfeldes in Abhängigkeit von der entsprechenden Erregung des Vorgabefeldes aufgetragen wurde, deutliche und charakteristische Nonlinearitäten auf. Die Ergebnisse dieser Experimente wichen somit deutlich von den Erwartungen ab, die sich aus klassischen Modellen des Simultankontrastes ergeben, denn diese Modelle würden typischerweise ein lineares oder affin-lineares Datenmuster erwarten lassen (Jameson & Huvich, 1972; Walraven, 1976; Shevell, 1978; Mausfeld & Niederée, 1993; Chichilnisky & Wandell, 1995). Der charakteristische Verlauf der Daten stimmt jedoch grob mit einigen in der Literatur schon berichteten Ergebnissen aus ähnlichen Experimenten überein (Smith & Pokorny, 1996; Miyahara et al., 2001).

Ein auffälliges Muster in den Datenkurven war eine treppenförmige Stufe, deren horizontale und vertikale Seiten immer mit den Chromatizitäten der jeweils verwendeten Umfelder übereinstimmten. Außerhalb der Region der Datenkurven, in der die Stufe auftritt, wurden nur geringfügige Simultankontrasteffekte festgestellt. Diese waren mit dem von Kries-Modell verträglic, wichen aber aufgrund ihrer geringen Ausprägung sehr deutlich von den auf Basis der Kontrastkodierungsmodelle zu erwartenden, sehr starken Effekten ab. Der charakteristische Verlauf der Datenkurven ließ vermuten, dass der gemessene Simultankontrasteffekt auf zwei unterschiedliche Mechanismen zurückzuführen ist, nämlich auf einen mit geringen Simultankontrasteffekten einhergehenden und einfachen von Kries-artigen Mechanismus, und einen zweiten Mechanismus, der nach Kontrastkodierungsprinzipien funktioniert. Letzterer schien für die charakteristische treppenförmige Stufe in den Datenkurven verantwortlich zu sein, denn inhaltlich konnte diese Stufe so verstanden werden, dass Abgleiche zwischen Inkrementen und Dekrementen[2] vermieden wurden, was einer

---

[2] Mit Inkrement bzw. Dekrement ist ein Infeld gemeint, dass einen höheren bzw. niedriegeren ’Farbwert’ hat als

zentralen Vorhersage der Kontrastkodierungsmodelle entspricht (Whittle, 1994b, 1994a; Kingdom, 2003).

Dabei fiel auf, dass diese Vermeidung von Inkrement-Dekrement-Abgleichen subjektiv sowohl mit ausgeprägten Abgleichsschwierigkeiten als auch mit dem Vorhandensein von Transparenzeindrücken einherzugehen schien. Diese Beobachtung führte auf die Hypothese, dass der eine der zwei postulierten Mechanismen, nämlich derjenige, der für die Stufe in den Datenkurven verantwortlich gemacht wurde, in enger Verbindung mit der Wahrnehmung von Transparenz steht. Diese Hypothese legt es nahe, die zwei postulierten Mechanismen dadurch zu isolieren, dass man die Wahrnehmung von Transparenz unterbindet, während man die sonstigen Versuchsbedingungen möglichst unverändert läßt. Gelänge die Umsetzung dieser Randbedingungen, so würde man erwarten, dass der von Kries-artige Effekt unverändert bestehen bleibt, während die mit Kontrastkodierung und Transparenzwahrnehmung assozierte Stufe in den Datenkurven verschwindet.

Basierend auf diesen Überlegungen wurde eine zweite Versuchsreihe durchgeführt, bei der – anders als in der ersten – neben homogen eingefärbten Umfeldern auch solche verwendet wurden, die mit einem zufälligen Farbmuster versehen waren. Diese variierten Umfelder, die auch als ’Seurat-Umfelder’ bezeichnet werden, da sie eine gewisse Ähnlichkeit zu einigen der pointillistischen Gemälden Georges Seurats (1859-1891) aufweisen, wurden bisher in einer Reihe von einschlägigen Untersuchungen eingesetzt (Andres, 1997; Mausfeld & Andres, 2002; Golz & MacLeod, 2002; Webster et al., 2002). Für die Zwecke der vorliegenden Untersuchungen erschien sie aus zwei Gründen besonders geeignet: Einerseits lag es nahe, dass durch den strukturelle Unterschied zwischen einem farblich variierten Umfeld und einem homogenen Infeld eine perzeptuelle Interpretation im Sinne einer gemeinsamen transparenten Schicht verhindert werden sollte. Dass dies in der Tat der Fall ist, wurde durch informelle Beobachtungen bestätigt. Andererseits konnte zu jedem homogenen Umfeld ein variiertes Umfeld hergestellt werden, das insofern damit vergleichbar war, dass es im räumlichen Mittel die gleichen Farbkoordinaten hatte. Entsprechend dieser Logik wurde in den asymmetrischen Farbabgleichexperimenten zu jedem Paar von homogenen Umfeldern auch ein entsprechendes Paar von variierten Umfeldern untersucht und die resultierende Paare von Datenkurven für homogene und variierte Umfelder direkt verglichen. Bei der Planung dieser zweiten Experimentalreihe wurde zusätzlich die in der ersten Versuchsreihe gemachte informelle Beobachtung berücksichtigt, dass die mit der Stufe in den Datenkurven assozierten Farbabgleiche nicht als subjektiv befriedigend empfunden wurden. Um diese Beobachtungen quantitativ fassbar zu machen, wurden Versuchsdurchgänge eingebaut, die es im nachhinein ermöglichten, zu prüfen, ob die von den Versuchspersonen vorgenommenen Abgleiche den Symmetrie- und Transitivitätsbedingungen einer Äquivalenzrelation genügen.

Die Ergebnisse dieser Versuchsreihe bestätigten sowohl die Hypothese, dass sich der bei den homogenen Umfeldern gemessene Simultankontrasteffekt aus zwei isolierbaren Teileffekten zusammensetzt, als auch die Vermutung, dass einige der von den Versuchspersonen vorgenommenen Abgleiche nicht dahingehend interpretiert werden können, dass es den Versuchspersonen tatsächlich gelungen wäre, die beiden Infelder ihrer Farbe nach perzeptuell gleich zu machen. Bei den variierten Umfelder traten durchgängig nur schwache Simultankontrasteffekte auf, die sich gut durch ein einfaches von Kries-Modell beschreiben ließen, und es wurden weder Transparenzeindrücke noch etwaige Abgleichprobleme beobachtet. Bei den homogenen Umfeldern hingegen trat zusätzlich zu diesem schwachen von Kries-artigen Effekt die treppenförmige Stufe in den Datenkurven auf, und damit assoziert auch wieder subjektive Abgleichprobleme.

Diese Ergebnisse legen nahe, dass homogene Umfelder insofern speziell sind, als dass sie eine Art von Simultankontrasteffekt hervorrufen, der bei variierten Umfeldern nicht auftritt. Dieser

---

sein Umfeld, wobei unter ’Farbwert’ entweder Luminanz oder die Erregung einer der drei Rezeptortypen zu verstehen ist (Wir betrachten hier jeweils nur eine dieser Größen, so dass uns die möglichen Mischfälle nicht zu beschäftigen brauchen).

Befund spricht stark gegen die in der experimentellen Farbenforschung verbreitete theoretische Vorstellung der sogenannten Funktionaläquivalenz, die besagt, dass man zu jedem variierten Umfeld ein entsprechendes homogenes Umfeld finden kann, das dengleichen Einfluss auf die wahrgenommene Farbe von darin eingebetteten Infeldern hat (Valberg & Lange-Malecki, 1990; Brill, 2000).

Den vorliegenden Befunden nach liegt die Besonderheit von homogenen Umfeldern auch darin, dass nur bei ihnen subjektive Abgleichprobleme auftraten. In den Abschnitten 5.3.1 und 5.3.2 der Arbeit wird gezeigt, wie die treppenförmigen Stufen in den Datenkurven direkt auf solche Abgleichprobleme zurückgeführt werden können, was bedeutet, dass diese Anteile der Datenkurven keinesfalls im Sinne einer perzeptuellen Gleichheit der beiden Infelder interpretiert werden sollten. Basierend auf informellen qualitativen Beobachtungen wird zunächst davon ausgegangen, dass es für jedes homogene Umfeld eine charakteristische Menge von Farbeindrücken gibt, die durch ein darin eingebettetes Infeld grundsätzlich nicht hervorgerufen werden können, gleichgültig welche Farkoordinaten man dafür wählt. Interessanterweise handelt es sich hierbei nicht etwa – wie man nach klassischen Modellen des Simultankontrastes erwarten könnte – darum, dass je nach Farbe des Umfeldes bestimmte Farbeindrücke vom oberen Ende der Sättigungsskala durch ein darin eingebettetes Infeld nicht hervorgerufen werden können, sondern darum, dass vor allem Farbeindrücke *vom unteren Ende der Sättigungsskala nicht realiserbar sind*. Der Zusammenhang zwischen Farbe des Umfeldes und der Menge der darin nicht realisierbaren Farbeindrücke kann grob wie folgt beschrieben werden: Während in einem *variierten* Umfeld alle Farbeindrücke herstellbar zu sein scheinen, fehlen in einem achromatischen *homogenen* Umfeld schon niedergesättigte Farbeindrücke jeden Farbtons. Nimmt man aber stattdessen ein farbiges Umfeld, so können noch weniger von den Farbeindrücken am unteren Ende der Sättigungsskala für Farbtöne, die dem Farbton des Umfeldes ähnlich sind, realisiert werden, dafür aber werden mehr Farbeindrücke vom unteren Ende der Sättigungsskala für dazu komplementäre Farbtöne wieder realisierbar. Bei der Betrachtung der in unterschiedliche Umfelder eingebetteten Farbskalen in den Farbtafeln XIV und XV, auf S. 145 ff. kann man eine ungefähre Vorstellung von diesen Phänomenen gewinnen.

Diese Beobachtungen, die die Vorstellung von perzeptuellen Sättigungsskalen nahelegen, die im unteren Bereich je nach Farbe des homogenen Umfeldes entweder abgeschnitten oder erweitert werden, werden in einem deskriptiven Modell zusammengefasst. Es wird dann aufgezeigt, dass aufgrund der nach diesem Modell zu erwartenden umfeldabhängigen Mengen von nichtherstellbaren Farbeindrücken die beobachteten treppenförmigen Stufen in den Datenkurven zu erwarten wären. Hierbei wird vorausgesetzt, dass die Versuchspersonen in den Fällen, bei denen ein exakter Abgleich unmöglich ist, stattdessen die ähnlichste Alternative mit demselben Farbton wählen.

Wie in den weiteren Teilen der Arbeit dargelegt wird, stellt es sich heraus, dass dieses deskriptive Modell bei der integrativen Beschreibung von vielen schon bekannten Aspekten des Simultankontrastes ein sehr hilfreiches Werkzeug darstellt und auch Vorhersagen erlaubt, die in weiteren Experimenten bestätigt werden konnten. Im Abschnitt 5.4 der Arbeit werden die Ergebnisse eines Experiments zum sogenannten ’Gamut-Expansion-Effekt’ von Brown und MacLeod (1997) dargestellt. Dieses Phänomen besteht darin, dass Infelder beliebigen Farbtons gesättigter erscheinen, wenn sie in einem homogenen grauen (neutralen) Umfeld eingebettet sind, als wenn sie in einem variierten Umfeld eingebettet sind, das dem räumlichen Mittelwert ihrer Farbwerte nach ebenso als neutral bezeichnet werden kann (siehe Farbtafel XVIII auf S. 148). In der Untersuchung von Brown und MacLeod (1997) wurde für jeden Farbton der Infelder nur eine einzige (geringe) Sättigungsstufe untersucht. Dabei stellte sich heraus, dass um einen Abgleich zu erzielen, die kolorimetrische Sättigung des in dem variierten Umfeld eingebetteten Infeldes im allgemeinen etwa um den Faktor 3.5 größer gewählt werden musste als die des im homogenen Umfeld vorgegebenen Infeldes. Dieser Faktor, der als ’relative Sättigung’ (’relative richness’) bezeichnet wurde, wurde von Brown und MacLeod als Maß für den Effekt verwendet.

In dem in dieser Arbeit vorgestellten Experiment zum Gamut-Expansion Effekt werden für jeden Farbton mehrere Sättigungsstufen untersucht. Es zeigte sich, dass die Größe des Effektes mit zunehmender Sättigungsstufe des vorgegebenen Infeldes rapide abnimmt. Der allgemeine Verlauf der Datenkurven stimmte dabei sehr gut mit der Erwartung überein, die sich auf der Basis des oben genannten deskriptiven Modells ergibt (Vergleiche die Datendiagramme in Abbildung 5.15 auf S. 111 mit der in Farbtafel XIX auf S. 149 dargestellten Vorhersage). Während dieses Modell die Daten gut beschreiben konnten, konnte ein einfaches, auf der Vorstellung einer multiplikativen Sensitivitätsregulierung ('gain-control') von opponenten Mechanismen aufbauendes Modell ausgeschlossen werden. Dieses in der Literatur verbreitete gain-control Modell (Webster, 2003) kann als eine Art von Kries-Modell höherer Ordnung aufgefasst werden. Es basiert auf der Idee, dass Farbkodierungsmechanismen, die Farbkontraste zwischen benachbarten Flächen kodieren, durch den hohen Kontrast im variierten Umfeld ermüden, und somit für den Kontrast zwischen dem Infeld und den Umfeldflächen weniger empfindlich seien. Dieser Vorstellung zufolge würde die im homogenen Umfeld stattfindende Farbwahrnehmung den Normalfall darstellen, während ein variiertes Umfeld zu einer modifizierten Farbwahrnehmung führt. Vor dem Hintergrund der in dieser Arbeit gewonnenen Einsichten scheint es aber umgekehrt der Fall zu sein, dass die im variierten Umfeld stattfindende Farbwahrnehmung den Normalfall darstellt, während homogene Umfelder zu einem spezielleren Modus der Farbwahrnehmung führen. Als wesentliche Hinweise für diese Deutung kann einerseits angeführt werden, dass in homogenen Umfeldern duale, transparenzähnliche Farbeindrücke auftreten, während die in dem variierten Umfeld eingebetteten Infelder eher die Erscheinungsweise von opaken Oberflächen haben, und andererseits die Tatsache, dass ein Infeld, das in einem variierten Umfeld eingebettet ist, Farbeindrücke aller möglichen Sättigungsstufen annehmen kann, während im Falle eines homogenen achromatischen Umfeldes Farbeindrücke niederen Sättigungsgrades nicht realiserbar sind.

Allgemeine Diskussion

Im Kapitel 6 der Arbeit werden die experimentellen Befunde mit anderen aus der Literatur bekannten Beobachtungen und aktuellen theoretischen Vorstellungen in Beziehung gesetzt. Dabei werden die folgenden Themen aufgegriffen:

- Deutung des Florkontrastes (Meyers Effekt)
- Die Rolle von Kontrastkodierung in der Farbwahrnehmung
- Funktionaläquivalenz von variierten und homogenen Umfeldern
- Die Validität von Graueinstellungen und asymmetrischen Farbabgleichen
- Die Repräsentativität von Graueinstellungen
- Der Zusammenhang zwischen Simultankontrast und Farbkonstanz
- Der Zusammenhang zwischen Simultankontrast und perzeptueller Transparenz

Im folgenden werden die zu diesen Themen angestellten Überlegungen kurz zusammengefasst.

Deutung des Florkontrastes Ein Phänomen, das in der älteren deutschsprachigen Literatur zum Simultankontrast eine zentrale Rolle spielte (siehe z.B. Helmholtz, 1911; Hering, 1887b; Perls, 1932), in moderneren englischsprachigen Arbeiten aber aus unklaren Gründen nur in seltenen

Fällen (Walls, 1960; Mausfeld, 2003a; Brown, 2003) Berücksichtigung findet, ist der sogenannte *Florkontrast*, der auch unter dem Namen *Meyers effekt* bekannt ist (Meyer, 1855). Die ursprünglich als bedeutsam erachtete Beobachtung besteht darin, dass der Simultankontrasteffekt nicht dadurch beeinträchtigt wird, dass man die Reizvorlage durch ein graues, durchsichtiges Medium wie etwa Transparenzpapier beobachtet (für eine Demonstration siehe Farbtafel XX auf S. 150). Da man damals wie heute oft davon ausgeht, dass der von einem farbigen Umfeld auf ein nominell graues Infeld ausgeübte Simultankontrasteffekt um so stärker ausfällt, je gesättigter das farbige Umfeld ist (alle klassischen Modelle des Simultankontrastes würden dies erwarten lassen), wurde dies als erklärungsbedürftig angesehen, denn ein graues transparentes Medium hat einen entsättigenden Einfluss auf die Umfeldfarbe und müsste somit zu einer Abnahme des Simultankontrasteffektes führen. Das in Abschnitt 5.3.2 der vorliegenden Arbeit vorgestellte deskriptive Modell sagt aber entgegen dieser verbreiteten Vermutung, und in guter Übereinstimmung mit empirischen Befunden (Kinney, 1962), vorher, dass der in ein graues Infeld induzierte Simultankontrasteffekt für alle Sättigungsstufen des farbigen Umfeldes genau gleich stark sein sollte (siehe Abbildung 6.2, S. 116).

Ein wesentlicher Grund für die Verwirrung in diesem Punkt dürfte der Umstand sein, dass nach klassischen Modellen des Simultankontrastes zu erwarten wäre, dass eine Änderung der Umfeldfarbe zu einer Farbverschiebung im Infeld führt, die für alle möglichen Infeldfarben in etwa gleich groß wäre (wegen ihrer linearen oder affin-linearen Struktur). Dies entspricht den Tatsachen aber keineswegs. Während es unbestritten bleibt, dass eine Änderung der Umfeldsättigung zu dramatischen Änderungen in der wahrgenommenen Farbe von einigen Infeldfarben führen kann, trifft dies für nominell graue Infelder nicht zu. Im Unterschied zu klassischen Modellen des Simultankontrastes trägt das hier vorgestellte deskriptive Modell sowohl der ersteren als auch der letzteren Beobachtung Rechnung.

Die Rolle von Kontrastkodierung in der Farbwahrnehmung Im traditionellen Verständnis der Farbwahrnehmung ging man davon aus, dass der Farbeindruck primär durch die lokale Stimulation der Netzhaut bestimmt wird. Dementsprechend wären Kontexteffekte auf durch das Umfeld hervorgerufene Perturbationen des lokalen Farbsignals zurückzuführen. Als prototypisch für diese Vorstellung kann das von Kries-Modell aufgefasst werden, das sich seinem Geiste nach auf theoretische Vorstellungen wie Rezeptorumstimmung, Adaptation und Ermüdung bezieht. Aus dieser Perspektive heraus ist es klar, dass das lokale Farbsignal eine grundlegendere Rolle spielt als die Farbsignale von benachbarten Regionen, denn die Rolle der letzteren besteht lediglich darin, das lokale Farbsignal zu *modifizieren*.

Nach der neueren theoretischen Perspektive, die durch die Kontrastkodierungsmodelle nahegelegt wird, spielen dagegen die benachbarten Farbsignale eine ebenso grundlegende Rolle wie das lokale Farbsignal. Da der Farbeindruck durch die *Differenz* zwischen benachbarten Regionen bestimmt wird und die Rezeptorerregungen auf beiden Seiten der dazwischenliegenden Kante gleichermaßen zu dieser Differenz beitragen, wäre es unangebracht, einem der beiden Farbsignal eine grundlegendere Bedeutung zuzuschreiben als dem anderen.

Diese beiden Hypothesen werden in der Literatur oft so behandelt, als wären sie sich gegenseitig ausschließende Alternativen, und je nachdem welche der beiden Prinzipien man als grundlegend betrachtet – lokale Absolutkodierung oder Kontrastkodierung – werden eine Reihe von empirischen Tatsachen grundlegend anders bewertet. Geht man zum Beispiel von der Kontrastkodierungsperspektive aus, so gerät man in die etwas sonderbare Position, dass man Tatsachen erklären muss, die dem unvoreingenommenen Laien völlig unspektakulär erscheinen. Zum Beispiel stellt man wohl im Alltag selten fest, dass die wahrgenommene Farbe eines Gegenstandes sich nennenswert ändert, wenn er vor einen anders gefärbten Hintergrund bewegt wird. Das Phänomen, dass die Farbe des Gegenstandes hierbei relativ konstant bleibt, überrascht den Laien wohl kaum, aus der Kontrastkodierungsperspektive betrachtet stellt es jedoch ein Rätsel dar, das als so wesentlich erachtet wird, dass man einen eigenen Namen dafür reserviert hat: Man spricht von ’hintergrundunabhängiger Farbkonstanz’ oder ’Typ-2 Farbkonstanz’ (Gilchrist et al., 1999; Whittle, 2003). Dieses Phänomen, das aus Sicht der Absolutkodierungsperspektive keiner weiteren Erklärung bedarf, wird im Rahmen der Kontrastkodierungsperspektive durch hypothetische Integrationsprozesse erklärt, die in höheren Gehirnzentren stattfinden sollen und die in der Netzhaut immer stattfindende Kontrastkodierung gegebenenfalls rückgängig machen. Innerhalb der Kontrastkodierungsperspektive werden die hypothetischen Integrationsprozesse entsprechend dafür verantwortlich gemacht, dass Simultankontrasteffekte manchmal sehr schwach oder sogar ganz fehlen, bisweilen aber auch dramatisch sein können: Beobachtet man starke Simultankontrasteffekte, so sei dies darauf zurückzuführen, dass die Integrationsprozesse irgendwie unterbunden oder umgangen werden, wohingegen die Abwesenheit von Simultankontrasteffekten darauf zurückgehe, dass die Integrationsprozesse normal funktionierten.

In der vorliegenden Arbeit wird aufgezeigt, dass die dramatischen Änderungen in der Stärke des Simultankontrastes, die bei einem fest gewählten Umfeld dadurch hervorgerufen werden können, dass man unterschiedliche Infeldfarben wählt, nicht durch unterschiedliche Aktivität der postulierten Intergrationsprozessen erklärbar sind. Es bleibt also aus dieser Perspektive unerklärt, warum man bei einem festen Umfeld für einige Infeldfarben fast perfekte Absolutkodierung feststellt und bei anderen Infeldfarben wiederum so starke Effekte beobachtet, dass die Annahme von Kontrastkodierung unvermeidlich scheint.

Als eine tragfähige theoretische Auflösung dieser Problematik wird in der Arbeit dafür argumentiert, dass höheren Gehirnzentren sowohl absolutkodierte als auch kontrastbasierte Signale zur Verfügung stehen, statt nur eines von beiden. Dabei scheint Kontrastkodierung nicht ein allgemeines Kodierungsprinzip zu sein, sondern eines, das bei speziellen Leistungen des Wahrnehmungssystems herangezogen wird, wie etwa bei der Wahrnehmung von Durchsichtigkeit.

Funktionaläquivalenz von variierten und homogenen Umfeldern Wie schon angesprochen, geht man nach klassischer Vorstellung von der Funktionaläquivalenz homogener und variierter Umfelder aus, d.h. man nimmt an, dass man zu jedem variierten Umfeld, ein homogenes finden kann, das denselben Effekt auf die wahrgenommene Farbe darin eingebetteter Infelder ausübt (Valberg & Lange-Malecki, 1990; Brill, 2000). In Übereinstimmung mit einer Reihe von neueren Untersuchungen zeigen die vorliegenden Befunde eindeutig, dass diese Vorstellung nicht zutrifft (e.g. Schirillo & Shevell, 1996; Brown & MacLeod, 1997; Shevell & Wei, 1998; Barnes et al., 1999). Es ist aber auch möglich und eigentlich auch sinnvoller, die Idee der Funktionaläquivalenz auf visuelle Mechanismen zu beziehen, statt auf psychophysikalisch direkt zu beobachtende Phänomene. Vor dem Hintergrund der vorliegenden Befunde erscheint es durchaus möglich zu sein, dass homogene und variierte Umfelder bzgl. desjenigen Mechanismus’ funktionaläquivalent sind, der hier als von Kries-Mechanismus bezeichnet wurde und scheinbar sowohl in homogenen als auch in variierten Umfeldern dieselbe Rolle spielt.

Die Validität von Graueinstellungen und asymmetrischen Farbabgleichen Asymmetrischen Farbabgleiche und Graueinstellungen stellen die wichtigsten traditionellen Verfahren zur quantitativen Messung des Simultankontrastes dar. Sie werden auch häufig in Untersuchungen zur Farbkonstanz eingesetzt (Brainard et al., 2003; Brainard, 1998; Golz & MacLeod, 2002). Diese Methoden basieren unmittelbar auf der in dieser Arbeit als Kompensationsannahme bezeichneten Grundannahme. Da sowohl die in Kapitel 4 als auch die in Kapitel 5 vorgestellte Versuchsergebnisse zeigen, dass diese Grundannahme im Falle von homogenen Umfeldern nicht zutrifft, müssen Ergebnisse von entsprechenden Untersuchungen mit äußerster Vorsicht interpretiert werden. In der vorliegenden Arbeit wird am Beispiel der Untersuchung von Smith und Pokorny (1996) aufgezeigt, wie eine Interpretation, bei der fälschlicherweise die Gültigkeit der Kompensationsannahme vorausgesetzt wird, zu Schlussfolgerungen führen kann, die nach den vorliegenden Befunden zu urteilen als nicht zutreffend betrachtet werden müssen. Im Falle von variierten Umfeldern scheint diese Problematik jedoch keine Rolle zu spielen.

Die Repräsentativität von Graueinstellungen Auch wenn man die mit der Validität von Graueinstellungen und asymmetrischen Farbabgleichen verbundene Problematik beiseite läßt, haben die vorliegenden Befunde interessante Implikationen für eine kürzlich von Speigle und Brainard (1999) behandelte Fragestellung. Diese Autoren gingen der praktisch und theoretisch interessanten Frage nach, ob man auf der Basis von relativ einfach durchzuführenden Graueinstellungen die Ergebnisse von wesentlich aufwendiger zu erhebenden asymmetrischen Farbabgleichen vorhersagen kann. Geht man von der Gültigkeit des von Kries-Modells aus, so müsste dies in der Tat der Fall sein, mit der Einschränkung, dass man nur die Chromatizität der Abgleiche vorhersagen könnte, während die Luminanz unbestimmt bliebe. In ihrer Untersuchung, die mit Kontextreizen durchgeführt wurde, die eher einer natürlichen Szene entprechen, konnte diese Frage bejaht werden. Die vorliegenden Ergebnisse deuten jedoch darauf hin, dass die Verwendung von homogenen Umfeldern, die in einschlägigen Untersuchungen weitaus üblicher ist, zu einem anderen Ergebnis führen würde. Insbesondere erscheint es naheliegend, dass man auf der Basis von Graueinstellungen den allgemeinen Simultankontrasteffekt deutlich überschätzen würde.

Der Zusammenhang zwischen Simultankontrast und Farbkonstanz Wie in der Einleitung schon angedeutet wurde, geht man im Allgemeinen davon aus, dass der Simultankontrast das Resultat eines fehlgeleiteten Versuchs des Wahrnehmungssystems darstellt, eine beleuchtungsunabhängige (d.h. farbkonstante) Repräsentation von Objektfarben zu gewinnen (Helmholtz, 1911; Hering, 1920; Walraven et al., 1987). Auch wenn diese Vorstellung prinzipiell durchaus vernünftig erscheint, so legen die vorliegenden Ergebnisse doch nahe, dass es irreführend sein könnte, wenn man die Art von Simultankontrasteffekten, die in *homogenen* Umfeldern auftreten, ohne weiteres im Sinne eines Farbkonstanzmechanismus' interpretiert.

Der Zusammenhang zwischen Simultankontrast und perzeptueller Transparenz Die vorliegenden Befunde deuten in Übereinstimmung mit einer Reihe früherer Beobachtungen (Masin & Idone, 1981; Brenner & Cornelissen, 1991; Brown & MacLeod, 1997; Mausfeld, 1998; Niederée, 1998) stark darauf hin, dass einfache Infeld-Umfeld-Reize transparenzähnliche, zweiwertige Farbeindrücke hervorrufen können. Mehrere Hinweise deuten darauf hin, dass starke Simultankontrasteffekte, wie sie in den vorliegenden Experimenten bei homogenen Umfeldern beobachtet wurden, stets von einer transparenz-ähnlichen Farbspaltung begleitet werden.

Eine Frage von zentraler theoretischer Bedeutung ist, ob die beobachtete Korrelation zwischen Stärke des Simultankontrasteffektes und dem Vorhandensein von deutlichen Transparenzeindrücken im Sinne eines kausalen Zusammenhangs gedeutet werden kann. Es ist nicht einfach, diese Frage experimentell zu prüfen, da die Wahrnehmung von Transparenz als potentielle kausale Variable nur auf indirektem Wege dadurch manipuliert werden kann, dass man den Reiz irgendwie ändert, was theoretisch auch den Simultankontrasteffekt direkt beeinflussen könnte. In den vorliegenden Experimenten wurde festgestellt, dass sowohl die Stärke des Simultankontrasteffektes als auch der Transparenzeindruck mit der Homogenität des Umfeldes kovariiert, so dass unklar bleibt, ob es die Homogenität des Umfeldes oder der Transparenzeindruck ist, der den Simultankontrasteffekt verursacht. Berücksichtigt man jedoch weitere Befunde, so scheint der Transparenzeindruck die relevantere Variable zu sein. In den vorliegenden Experimenten wurde festgestellt, dass der

Transparenzeindruck dadurch eliminiert werden kann, dass man das homogene Umfeld durch ein entsprechendes variiertes ersetzt, und dass dies auch zu einer Elimination der starken Simultankontrasteffekte führt. Geht man davon aus, dass der Transparenzeindruck die relevante Variable ist, so sollte es auch möglich sein, starke Simultankontrasteffekte in variierten Umfeldern zu erzeugen, indem man dafür sorgt, dass auch hier Transparenzeindrücke auftreten. Tatsächlich kann man hier den Eindruck von Transparenz entstehen lassen, indem man auch das Infeld entsprechend inhomogen macht, wie es etwa bei der Chubb-Täuschung der Fall ist (Chubb et al., 1989; Lotto & Purves, 2001). Hurlbert und Wolf (2004) untersuchten diesen Fall und stellten in Übereinstimmung mit dieser Hypothese starke Simultankontrasteffekte fest. Dadurch wird nahegelegt, dass der Transparenzeindruck eine wichtigere Variable darstellt als die Homogenität des Umfeldes an sich.

Fazit Die Ergebnisse der in dieser Arbeit vorgestellten Experimente zeigen, dass der Simultankontrast ein weitaus komplizierteres und vielschichtigeres Phänomen ist als allgemein angenommen wird. Während der Simultankontrasteffekt, der in variierten Umfeldern auftritt, sehr einfach zu bescheiben ist, scheinen einfache homogene Umfelder komplexere perzeptuelle Mechanismen anzustoßen, die zu transparenzähnlichen zweiwertigen Farbeindrücken führen.

Anhang C

Lebenslauf (Curriculum Vitae)

Persönliche Daten

- Name: Vebjørn Ekroll
- Geburtsdatum: 24.09.71
- Geburtsort: Bergen, Norwegen
- Familienstand: Ledig

Schulausbildung

- 1978 – 1987: Neunjährige Grundschule, Lynghaug Skole, Bergen, Norwegen
- 1987 – 1988: Berufschule, Grundausbildung zum Zimmermann, Bergen Yrkesskole, Bergen
- 1990 – 1993: Gymnasium, Tanks Skole, Bergen
- September 1993: Erlangung der allgemeinen Hochschulreife

Hochschulstudium

- 1993 – 1994: Universität Oslo, Norwegen, Examen Philosophicum im Fernunterricht
- 1994 – 1995: Universität Bergen, Norwegen, Grundkurs Psychologie
- 1995 – 2000: Christian-Albrechts-Universität zu Kiel, Diplomstudiengang Psychologie
- Oktober 1997: Vordiplom
- September 2000: Diplom

Berufliche Erfahrungen

- 1988 – 1990: Lehre zum Zimmermann, Fa. Thor A. Marthinsen A/S
- 1993 – 1994: Zivildienst, Rogaland Psykiatriske Sykehus, Stavanger, Norwegen
- 1997 – 2000: Wissenschaftliche Hilfskraft für Prof. Dr. Mausfeld, Institut für Psychologie, Christian-Albrechts-Universität zu Kiel
- Seit Dezember 2000: Wissenschaftlicher Mitarbeiter von Prof. Dr. Mausfeld, Institut für Psychologie, Christian-Albrechts-Universität zu Kiel