Chromatic scene statistics as cues for the perception of surface and illumination colour

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

vorgelegt von

Jürgen Golz

Kiel
2004

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

Zweitgutachter: ..... PD Dr. Johannes Andres

Tag der Mündlichen Prüfung: 09.02.2005

Durch den zweiten Prodekan, Prof. Dr. Norbert Nübler
zum Druck genehmigt am: 09.02.2005

Contents

Chapter 1

Introduction

The present work will treat theoretically as well as experimentally the question how humans are able to perceive the colour of surfaces and illuminants. While this ability might seem trivial because of the ease and reliability with which it is accomplished by our visual system, the following problem becomes obvious as soon as one starts to think about the relation between the chromatic properties of a physical object and its illumination on the one hand and the colour of the two corresponding perceptual representations on the other hand: the light reflected from the object depends both on the spectral reflectance of the objects surface as well as on the spectral power distribution of the illuminant. These two chromatic components are merged in the retinal stimulation that forms the input of the visual system. How is the visual system able to disentangle these two components for the perception of surface and illumination colour? In doing so, humans (as well as many other species) fulfill the biologically important task to recognize objects by its colour despite changes in illumination. Our ability to tell, for instance, whether a fruit is ripe, whether a snake or a mushroom is poisonous, and which car is ours does usually not hinge on chromatic aspects of the lighting conditions. This (approximate) constancy of the perceived colour of surfaces despite changes of illumination is called colour constancy and is only one example of many perceptual constancies that the human visual system accomplishes.

Though this problem has been addressed scientificly since more than a century (the first systematic investigation of colour constancy was carried out by Helmholtz) it still lacks a satisfactory answer and can be regarded as one of the major mysteries of perceptual psychology. One reason for this surprisingly little progress of insight is the fact that, for most of the 20th century, a kind of research dominated the field that dedicated itself to the investigation or early peripheral processes of colour coding. To this end, isolated patches of light in otherwise dark surrounds were used as stimuli and the theoretical treatment was

mostly restricted to the local relation of the proximal stimulus and the resulting colour perception with little reference to more global aspects of complex stimuli and their relation to distal objects and illuminants. Only in recent decades it has been acknowledged again that the problem of colour constancy cannot be solved by focusing the analysis on processes of local nature (a point that has been stressed in earlier works for instance by Katz (1911), Gelb (1929), and Koffka (1932)) and hence the question what information of the entire retinal stimulus allows the visual system to solve the problem of colour constancy is increasingly attracting attention.

The approach taken here is, in a first step, to theoretically analyse what statistics of the chromatic distribution of reflected lights in the retinal image of natural scenes could allow the visual system to estimate the illuminant and therewith derive surface colours in a colour constant way. In a second step, the question whether the human visual system indeed exploits these statistics as cues for the perception of surface and illuminant colour will be addressed experimentally.

More precisely, the present work is organized as follows: chapter 2 provides a more detailed presentation of the above-mentioned problem of the perception of surface and illumination colour and the more specific problem on which I will mostly focus in this work, i.e. how the visual system can estimate the chromatic properties of the illuminant. The role that chromatic scene statistics can play in order to solve this problem are then considered on a general level.

In chapter 3, I will address the question how to obtain predictions what particular scene statistics could be helpful for estimating the illuminant by presenting four different idealized models of the chromatic environment that make different assumptions about the spectral functions of surfaces and illuminants. The usefulness of these idealized worlds for the purpose of deriving cues to the illuminant will be assessed and for one of these models, the so called "Gaussian World", I will analyse in detail the changes of retinal cone excitations that surfaces undergo when the illumination changes. One of the advantages of the Gaussian World over the other idealized worlds is that these changes are mathematically tractable in this model. The regularities of these changes are discussed with respect to the question what scene statistics depend systematically on the illuminant and are thus potential cues for estimating the chromatic properties of the illuminant.

In chapter 4, the regularities of the retinal changes of surfaces under changing illumination and the scene statistics that consequently could be useful cues to the illuminant are investigated again, but now by analysing (instead of an idealized world) images of real natural scenes for which different illuminations are simulated. As we will see, the analyse of this real world simulation yields results that are fairly consistent with the predictions of the Gaussian World.

In chapter 5, I will turn to the question whether the potential cues to the illuminant found in the two preceding theoretical analyses (namely the luminance-redness correlation in the retinal image, the luminance-blueness correlation, and the redness-blueness correlation) are indeed exploited by the visual system. To this end, a series of seven psychophysical experiments is presented and discussed.

Chapter 6 will finally summarize the theoretical and experimental findings of the previous three chapters and discuss how these results relate to each other as well as what questions remain open.

Chapter 2

The perception of surface and illumination colour

2.1 The problem

The nature of the problem underlying the perception of surface and illumination colour can be appreciated by considering the example of the spectral interplay of surfaces and illuminants in Fig. 2.1. In this example, a greenish piece of paper[1] is viewed by a human observer. Surfaces are characterized by their reflectance functions $S(\lambda)$, i.e. the ratio of the reflected radiant power to the radiant power incident on the surface for each wavelength in the visible range of the electromagnetic spectrum. As Fig. 2.1(c) shows, the surface in our example reflects medium wavelengths best while short and long wavelengths are less well reflected.

If this surface is lit in turn by two different illuminants (a bluish daylight of 16 200 K colour temperature whose spectral power distribution $I(\lambda)$ is plotted in Fig. 2.1(a) and a more orange daylight of 4500 K colour temperature plotted in Fig. 2.1(b)), two substantially different reflected lights result: the spectral power distribution of the light reflected from the surface under the first illuminant (Fig. 2.1(d)) has more energy in the short wavelengths and less energy in the long wavelengths than the light reflected under the second illuminant (Fig. 2.1(e)) because the two illuminants themselves differ in this respect. (The spectral power distribution of a reflected light $R(\lambda)$ is the wavelength-by-wavelength product of the spectral reflectance functions $S(\lambda)$ of the surface and the spectral power distribution $I(\lambda)$ of the illuminant.) Consequently, the retinal cone excitations

[1] More precisely, it is the greenish colour sample chip 5G 8/6 from the Munsell (1976) book of colour.

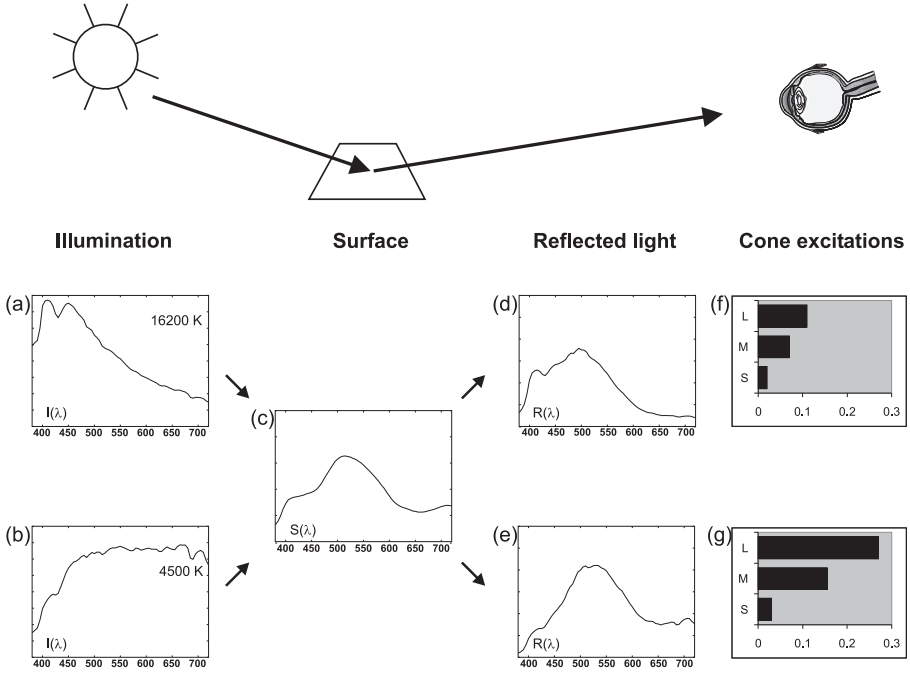

Figure 2.1: The problem of colour constancy (for explanation see text).

elicited by the surface under these two illuminants (Fig. 2.1(f) and (g)) differ, too.

This example shows that the proximal stimulus, i.e. the retinal image produced by the spectral interplay of the distal surfaces and illuminants, varies substantially with changes of illumination. But, as mentioned in the introduction, the human visual system has the ability to perceive the colour of surfaces to a large extent as constant despite changes of illumination. For instance, the surface of the above example when viewed under the two illuminants would appear in approximately the same greenish colour (at least under viewing conditions that are not highly artificial, see below). But despite this constancy of the perceived surface colour, the two cases might differ — with respect to the perception of the illuminant incident on the surface: when lit by the 16 200 K daylight, the surface would appear as a greenish surface illuminated by a cold, bluish light. But when lit by the 4500 K daylight, it would appear as the same greenish surface, but now illuminated by a warmer, more orange light. It is worth noting that colour constancy does not mean that a surface does look identical in all respects under changes of illumination (as e.g. pointed out already by Katz (1911) and Gelb (1929)). Mausfeld (2003) provides a theoretical framework that allows to formulate this insight in terms of modern cognitive science: the visual system creates (under suitable conditions) a representation for the surface as well as a representation for the illuminant, which both have a colour among their parameters.

The problem to determine these two parameters cannot be solved by analysing the retinal stimulus locally because a given cone excitation can be produced by many combinations of distal surfaces and illuminants. That the visual system relies for the perception of the colour of a surface on more global information than the corresponding local retinal stimulation becomes obvious e.g. from the fact that the perception is drastically altered when the same local information is presented in isolation. When for instance the greenish surface of the above example would be presented as the only object, illuminated by the 16 200 K light, in an otherwise completely dark room, it would not look the same green as when viewed in a richer scene with more objects, but rather turquoise. But moreover, this colour would not look like a surface colour at all but rather have the appearance of an "aperture colour" (Katz, 1911, he used the German term "Flächenfarbe"). Another example is the moon in the sky at night, which does not look like a rocky-grey surface that is lit by the sun, but rather like a light source itself. It seems that a perceptual segmentation in two representations, one for the surface and one for the illumination, does not occur in situations like these.

What is called traditionally "the problem of colour constancy" can thus be considered as two interrelated problems: first, how does the visual system estimate the chromatic properties of the illuminant in order to determine the colour parameter of the illumination representation and second, how does the visual system, taking into account this illumination estimate, determine the colour parameters of the surface representations? I will refer to the first problem as the "estimation problem" and address it by analysing, what global chromatic cues are available in the retinal image in order to estimate the illuminant (chapter 3 and 4), and by testing experimentally whether these cues are indeed used by the visual system (chapter 5). I will deal with the second problem only indirectly insofar as the analyses of chapter 3 and 4 yield results of how cone excitations of different surfaces are transformed by changes of illumination. These findings could be used to derive predictions for the way in which the visual system could determine the surface colours on the basis of the retinal image and the estimated illumination.[2] One major question in this context is whether under changes of illumination the retinal stimulation of all surfaces in a scene are altered by the same transformation or whether different surfaces undergo different illumination-induced shifts. In the former case, the visual system could determine the surface colours by compensating all surfaces of the scene in the same way when taking the estimated illuminant into account. I will refer to this second problem by the term "compensation problem".

[2]Depending on the broader theoretical perspective and the normative character in which such a consideration is pursued, it could lead to a valuable heuristic or right into the physicalistic trap (Mausfeld, 2002).

Before addressing these two problems in more detail in the remaining of this chapter, I will first introduce in the next section some colourimetric tools that we will need.

2.2 Colour coordinates

Throughout this work I will use cone excitation measures to specify the chromatic properties of the proximal stimulus. These colourimetric measures, instead of really corresponding to physiologically determined states of the retinal cones, are tristimulus codes (also called "Graßman codes") that are based on psychophysical colour matching experiments and then transformed on the basis of physiologically oriented considerations. It is not my aim here to provide an introduction to the field of colourimetry (this could e.g. be found in Koenderink and van Doorn (2003)).

The estimates of the spectral sensitivities of the human cones used in this work are those of Stockman, MacLeod, and Johnson (1993). The spectral sensitivity functions of the cone that is most sensitive to long wavelengths (L cone), the cone most sensitive to medium wavelengths (M cone), and the cone most sensitive to short wavelengths (S cone) are denoted by $C_L(\lambda)$, $C_M(\lambda)$, and $C_S(\lambda)$, respectively. The excitations elicited in these three cone types by a light with the spectral power distribution $R(\lambda)$ can then be determined as follows:

$$\begin{aligned} L &= \int R(\lambda)C_L(\lambda)d\lambda \\ M &= \int R(\lambda)C_S(\lambda)d\lambda \\ S &= \int R(\lambda)C_M(\lambda)d\lambda \end{aligned}$$

Instead of considering the three dimensional cone excitation vector (L, M, S), it is often more convenient to treat separately two chromaticity coordinates (that together form an intensity-invariant chromaticity plane) and luminance. The chromaticity coordinates used in this work are the measures l and s of MacLeod and Boynton (1979), which are defined as:

$$\begin{aligned} l &= L/(L+M) \\ s &= S/(L+M) \end{aligned}$$

It is generally assumed that S cone excitation does not contribute to luminance (Eisner & MacLeod, 1980), thus:

$$luminance = L + M$$

2.3 Estimating the illuminant

In this work, I will focus on one particular source of information for estimating the illuminant: the illumination-dependent chromatic properties of the retinal image. Illuminants affect the distribution of retinal excitations elicited from surfaces in the scene and thus may be recognized by statistics of this chromatic distribution in the retinal image, which I will referred to as "chromatic scene statistics". This leaves aside many other potential cues to the illuminant, particularly ones that depend on the three-dimensional geometry of the scene, like e.g. specular reflections (D'Zmura & Lennie, 1986; Lee, 1986; Tominaga & Wandell, 1989; Maloney & Yang, 2003) and mutual reflections (Funt, Drew, & Ho, 1991; Funt & Drew, 1993; Bloj, Kersten, & Hurlbert, 1999). Effectively, the situation considered is a world of co-planar diffusely reflecting surfaces, with a single illumination incident uniformly over each scene.

2.3.1 Scene-averaged chromaticity

When the illumination of a scene is changed, e.g. towards more energy at long wavelengths, all reflected lights will change correspondingly, and the chromaticity averaged over the entire retinal image becomes more reddish (higher l value). A simple approach to colour constancy could therefore take the space average chromaticity as an estimate for the chromatic properties of the illuminant and use this estimate to correct for shifts in chromaticity of objects due to changes of illumination. A similar consideration has already been proposed by Helmholtz (1896).

The Grey World assumption

Buchsbaum (1980) proposed an algorithm for colour constancy in line with the above consideration: it uses the spatial average of the receptor responses in order to estimate the illumination of the scene. His estimation of the illuminant is based on the assumption that for all scenes the scene-average reflectance is equal to an internal reflectance standard $S_0(\lambda)$. Such an assumption is often referred to as a "Grey World assumption" (though in the model of Buchsbaum, $S_0(\lambda)$ does not neccesarilly have to be chromatically neutral, but merely the same for all scenes). The illuminant is then estimated by determining which illuminant would have resulted in the actual obtained average receptor response, assuming that the scene is illuminated uniformly and the spatial average reflectance of the scene is $S_0(\lambda)$. This estimate is then used together with the responses for

the subfields in the image in order to obtain illuminant independent reflectance descriptors for each subfield.

The ambiguity of the mean image chromaticity

For scenes for which the mean reflectance function differs from $S_0(\lambda)$, the algorithm wrongly attributes a non-neutral mean receptor response to a spectrally biased illuminant. Every change in the mean receptor response is interpreted as due to a change in the illumination. Thus, this algorithm, like many others relying on the Grey World assumption, is not able to discriminate variations of colour balance inherent to the surfaces of scenes from those due to the illuminantion. This weakness is inherent to all models that use a space averaged chromaticity of the scene to estimate the illuminant. The reason is that this measure is ambiguous: for a given retinal image of a scene the mean chromaticity could be e.g. reddish due to a predominance of reddish surfaces within this scene, or as well due to a reddish illuminant.

This ambiguity of the mean image chromaticity is further illustrated in Plate I(a) by considering a sample of scenes under two illuminants (a white and a reddish one). The same mean chromaticity could result either from a reddish scene under white illumination or from a white scene under reddish illumination, since these two illuminants, applied to the given sample of scenes, generate overlapping distributions of mean image redness. In section 4.3.2 I will show that this ambiguity of the mean image chromaticity is not a mere hypothetical problem, but that the mean image chromaticities of natural scenes lit by the same illuminant can vary substantially, leading to overlapping distributions even for substantially different illuminants.

2.3.2 The use of higher order scene statistics

Nevertheless, human observers are somehow able to discriminate a white scene under reddish illumination from a red scene under white illumination (Gilchrist & Ramachandran, 1992; Kraft & Brainard, 1999). Which information could help them to resolve the ambiguity of the mean image chromaticity? The use that higher order scene statistics (standard deviations, correlations, skewnesses of the chromatic distribution in the retinal image) could yield for estimating the illuminant, when these measures are taken into account additionally to the mean image chromaticity, is illustrated in the schematic constellations (b)-(d) in Plate I.

In the example (b) of Plate I a higher order scene statistic depends systematically on the illumination: it is increased for all scenes under the reddish illumination. Thus the overlap of the two probability distributions of the sample of scenes under the two illuminants is reduced and by taking into account the higher order scene statistics the visual system could resolve the ambiguity encountered in considering mean chromaticity alone. A high value of this statistic within the image of a scene would be evidence that the illuminant is reddish.

The constellation (c) in Plate I shows another way in that higher order scene statistics could reduce the ambiguity of the mean chromaticity. Though, in this case the higher order scene statistic does not, for a given scene, differ for the two illuminants, together with the mean image redness it separates the two distributions of images from the two different illuminants. Given a certain mean redness, a higher value of the higher order scene statistic would indicate that the illumination is more reddish.

But not every statistic, even if it systematically varies under changes of illumination, is useful for resolving the ambiguity of the mean chromaticity, as illustrated in constellation (d) of Plate I. Here the two statistics are highly correlated both within and across the clusters of different illuminants. In such a constellation, the combination of the two statistics is only as helpful for estimating the illuminant as each statistic is alone.

Experimental evidence that the visual system exploits higher order scene statistics for estimating the illuminant

The first study that has explicitly addressed the question whether the visual system takes into account a higher order scene statistic as a cue for the illumination was carried out by Mausfeld and Andres (Mausfeld, 1998; Mausfeld & Andres, 2002). Based on the consideration that a chromatically biased illuminant with a sufficiently narrow bandwidth would reduce the chromatic variance in the retinal image, they investigated whether the tendency of the visual system to attribute a chromatically biased average in the retinal image to a non-neutral illumination increases when the chromatic variance in the retinal image decreases.

To this end, they created a new type of stimulus that makes it possible to vary certain statistics of the chromaticity distribution of the image independently. (A very similar type of stimuli will be used in the experiments in chapter 5. For examples of these stimuli see Plate IV.) By varying the modulation of the colour of the circles in the surround along the luminance axis and along the red–green axis, Mausfeld and Andres obtained a set of stimuli, each with a different combination of variances for chromaticity and luminance. The spatial mean

chromaticity and luminance of the surround were equal for all stimuli. Subjects made red–green equilibrium settings ('unique yellow' settings) by a double random staircase procedure. These subjective yellow settings point toward the mean surround chromaticity, as expected by standard findings on colour induction, but the shift was greater for surrounds of reduced chromatic variance than for very heterogeneous surrounds. Mausfeld and Andres conclude from these results that a reduced chromatic variance increases the tendency of the visual system to interpret a scene with a biased average chromaticity as chromatically illuminated. This results in a stronger correction for the appearance of all patches including the test spot. Thus, the effect of a chromatically non-neutral surround on unique yellow settings is larger if the surround has low chromatic variance than if it has high chromatic variance. That is, these surrounds are not functionally equivalent even though the space-averaged chromaticity of both surrounds is equal. The dependence of the yellow settings on the chromatic variance in the surround can be taken as evidence that the visual system takes into account this statistic for resolving the ambiguity of the mean chromaticity when estimating the colour of the illuminant.

Lotto and Purves (1999) also investigates the role of the chromatic variance in the retinal image for the estimation of the illuminant by the visual system. From experiments in which subjects have to match a test field in perceived lightness to a standard that is embedded in variegated isoluminant surrounds of different chromatic variances, they conclude that the visual system infers a brighter illumination for retinal images with higher chromatic variances. Their consideration about the underlying regularity is that an image with a certain chromatic variance (and a given luminance) can only result from illuminants with at least a certain minimum intensity and this required minimum intensity increases for images with higher chromatic variance. (The narrower bandwidth of surfaces that are chromatically more biased require a more intense illumination to yield the same luminance in the retinal image than surfaces with broader bandwidths that are chromatically more neutral.)

As a third experimental evidence that the visual system takes into account higher order scene statistics for the perception of surface and illumination colour, I have found in cooperation with Donald I. A. MacLeod (Golz & MacLeod, 2002) that the luminance-redness correlation within the retinal image is exploited as a cue for the illumination. Experiment 1 in chapter 5 is a slightly modified replication of that work. The underlying regularity, namely that surfaces that are chromatically similar to the illuminant are rendered as lighter than surfaces that are dissimilar from the illuminant, is presented and discussed in detail in chapter 3 and 4.

2.4 Determining surface colours

Given an estimate for the illumination incident on a scene, the question is now, how the visual system can determine the perceived colours of the surfaces within that scene such that they remain more similar to the corresponding colours perceived under a different illuminant than could be expected on the basis of the illumination-dependent retinal image. To this end, the visual system has to take into account the estimated colour of the illumination when deriving the colours of the surfaces from the cone excitations of the corresponding elements in the retinal image. For this dependence of the colour parameters of the surface representations on the colour parameter of the illumination representation, the visual system has to make assumptions about how different surfaces are affected by the illuminant when forming the retinal stimulus. Of special interest in this context is the question whether cone responses of different surfaces all undergo the same change under changes of illumination.

As we will see in section 3.2, the fact that there is at all some degree of order in the changes of the cone responses of surfaces under changes of illumination is only possible because the chromatic world we live in is suitably constrained to smooth and broad spectral functions. In a more malign world, these changes could be chaotic and unpredictable in that the cone responses of a set of surfaces could be transformed in almost any imaginable way when the illumination changes.

One particular illumination-dependent transformation that has recurrently been considered (e.g. Ives, 1912; West & Brill, 1982; Worthey & Brill, 1986; Foster & Nascimento, 1994) is a multiplicative change of the cone excitations under changes of illumination, where the excitations of each cone type are scaled for all surfaces by the same factor. Let (L_i, M_i, S_i) denote the cone excitation triple elicited from a particular surface i under one illuminant and (L'_i, M'_i, S'_i) the corresponding cone excitation triple under a second illuminant. Than this type of transformation can be written as

$$\begin{aligned} L'_i &= k_L * L_i \\ M'_i &= k_M * M_i \\ S'_i &= k_S * S_i \end{aligned}$$

Note that the multiplicative scaling factors are the same for all surfaces. Thus, when considering a logarithmic cone excitation space in which the cone excitation triples are represented on a log scale, all surface are shifted by the same additive translation vector whose components are the logs of the scaling factors:

$$\begin{aligned} log(L'_i) &= log(L_i) + log(k_L) \\ log(M'_i) &= log(M_i) + log(k_M) \\ log(S'_i) &= log(S_i) + log(k_S) \end{aligned}$$

Geometrically, the constellation of points representing the surfaces would move rigidly in log cone excitation space under changes of illumination. To the degree that this transformation correctly describes the illumination-induced changes of cone excitations, the visual system could (at least in the restricted flat matte world considered here) determine the colours of surfaces by some kind of normalization of the corresponding cone excitations, where the normalization factors would depend on the estimated colour of the illuminant. I will therefore refer to this kind of transformation of the changes of cone excitations under changes of illumination as normalization-compatible.

It has often been proposed that the visual system accomplishes colour constancy by some kind of re-adjustment of the sensitivity of each cone type, called von Kries adaptation, in order to compensate for the illumination-induced transformation of the retinal input discussed above (see e.g. Ives, 1912; Webster, 2003). There are many objections, mainly of a non-computational character, to the adoption of reciprocal adjustments of photoreceptor sensitivity as a model for colour constancy (see e.g. MacLeod, 1985; Mausfeld, 1998). One serious objections is that such adaptational models normalize out variations of lightness and colour balance inherent to the surfaces of a scene, as well as those due to the illuminant. Thus, such models face the same problems that any Grey World algorithm faces (cf. page 12). Moreover, Mausfeld (2003) points out that such adaptational models of colour constancy are connected to an elementaristic perspective on colour perception that neglects illumination-related issues and has not produced much insight into the internal conceptual architecture of the perceptual system that underlies the perception of surface and illumination colour.

Chapter 3

Theoretical worlds

3.1 Which heuristic to use for discovering chromatic cues?

In the previous chapter it has been suggested that the visual system could use chromatic scene statistics as cues for the perception of surface and illumination colour in order to resolve the ambiguity that is inherent in the mean image chromaticity. But which chromatic scene statistics depend systematically on the chromatic properties of the illumination and hence are useful cues to the illumination? For an approach to that question that is guided by computational considerations in the sense of Marr (1982), an idealized world model would be desirable that describes how the chromatic properties of the proximal stimulus depends on the chromatic properties of distal surfaces and illuminants. Based on such a model, one can analyse how the proximal chromaticities change when the illumination changes. Any realistic model that predicts systematic changes of the proximal chromaticities under changes of illumination, constitutes a promising search heuristic for chromatic cues for the illumination: a statistic of the chromatic distribution in the proximal stimulus that catches an illumination-dependent regularity predicted by the model can be taken as a candidate cue that might be used by the visual system for estimating the illumination.

In general, it is not possible to determine the chromatic properties of a reflected light without knowledge of the spectral functions of the corresponding surface and illuminant. Different from the results of additive light mixtures, the interplay of surfaces and illuminants cannot be described lawfully on the level of cone excitations (or other tristimulus values), but, like the results of subtractive light mixtures, only on the spectral level. Therefore, any candidate idealized world

model has to make explicit assumptions about the spectral properties of surfaces and illuminants. The spectral assumptions made in such an idealized world model then determine whether (and which) illumination-dependent regularities will be predicted.

In this chapter I present the following four idealized worlds and discuss their respective usefulness as a heuristic for discovering chromatic cues for the illumination:

1. In the *Chaotic World*, no constraints at all are made about the spectral properties of surfaces and illuminants. A scenario is presented (using a malign choice of surface and illumination spectra) in order to show that in such an unconstrained world the interplay of surfaces and illuminants can become chaotic: there is, in general, no order in the changes of cone excitations under changes of illumination.

2. Any such order in changes of cone excitations under changes of illumination will only exist given suitable constraints about the spectral properties of surfaces and illuminants. The constraints of the *Three-band World* yield one of the highest degrees of order one could think of: normalization compatible changes of cone excitations under changes of illumination. But the constraints that this idealized world use are very unrealistic as descriptions of the spectral properties of the natural world.

3. A more realistic approximation of naturally occurring spectra is intended in the *Linear World* by modelling surface reflectance functions and illumination power distributions by weighted sums of spectral basis functions. While it is obvious that such a model could perfectly describe any set of spectra by means of a sufficiently large number of suitably chosen basis functions, it is still hotly debated whether a satisfactory approximation could be obtained with a number of basis functions small enough to make the model a plausible heuristic for human colour constancy. Despite extensive work employing the linear world as an analytical tool, it does not seem to have produced much insight into the actual principles of the chromatic interplay of surfaces and illuminants in the natural environment and it has not led to predictions of chromatic scene statistics that the visual system could exploit as cues for the illuminant.

4. In the *Gaussian World*, all spectral functions (surface reflectances, spectral power densities of illuminants, cone sensitivities) are described by Gaussian functions. Besides several other advantages of this way of modelling, the Gaussian World gives us something not available in the other idealized worlds: a simple algebraic solution for the retinal cone excitations as resulting from the interplay of surfaces and illuminants. With this tool at

disposal, it will be possible to derive several illumination-dependent regularities and thus predict potential cues for the illumination. These cues will then be analysed in chapter 4 with regard to their usefulness in natural scenes, and in chapter 5 it will be tested experimentally whether the human visual system exploits these cues for the perception of surface and illumination colour.

3.2 The Chaotic World

In a world without any constraints on the spectral functions of the illuminants and surface reflectances, the interplay of these two components, i.e. the physical process that yields the spectral properties of the proximal stimuli, could be chaotic and unpredictable: the reflected lights and thus the cone excitations would in general not change systematically with changes in illumination.

The nature of this problem can be appreciated by considering the following example (Fig. 3.1): Suppose that two (almost) metameric yellow illuminants, visually similar but physically different, are each supplying light in only two narrow spectral bands, one green and one red, but that the bands supplied by the two light sources are slightly offset from one another, by enough that they do not overlap (left column in Fig. 3.1). Denote the bands by G1 and R1 for illuminant 1 and G2 and R2 for illuminant 2. Imagine each of these sources in turn lighting the following two yellow surfaces. One surface, (which we may call surface G1R2), has a spectral reflectance that takes nonzero values only in the bands G1 and R2. The other yellow surface, surface G2R1, reflects only in the other two bands (middle column in Fig. 3.1).

Consider now the resulting reflected lights of these two surfaces under each illuminant, respectively (right column in Fig. 3.1): Surface G1R2 is physically green under illuminant 1, it only reflects the green band G1 of radiations to the eye, and not the second, red band R1 supplied by illuminant 1. Surface G2R1, on the other hand, turns red under illuminant 1, because only its red reflection band R1, and not the green one G2, is present in the spectrum of illuminant 1. The situation is completely reversed under illuminant 2. Now surface G2R1 reflects green light only, and surface G1R2 reflects only red.

This example is enough to show how there are, in general, no constraints at all on how a change of illumination will change the reflected chromaticity of a surface: any set of chromaticities can be transformed into any other set. In this example the two illuminants, as well as the two surfaces (when lit by a white illuminant), can be indistinguishable yellows (because the reflection bands

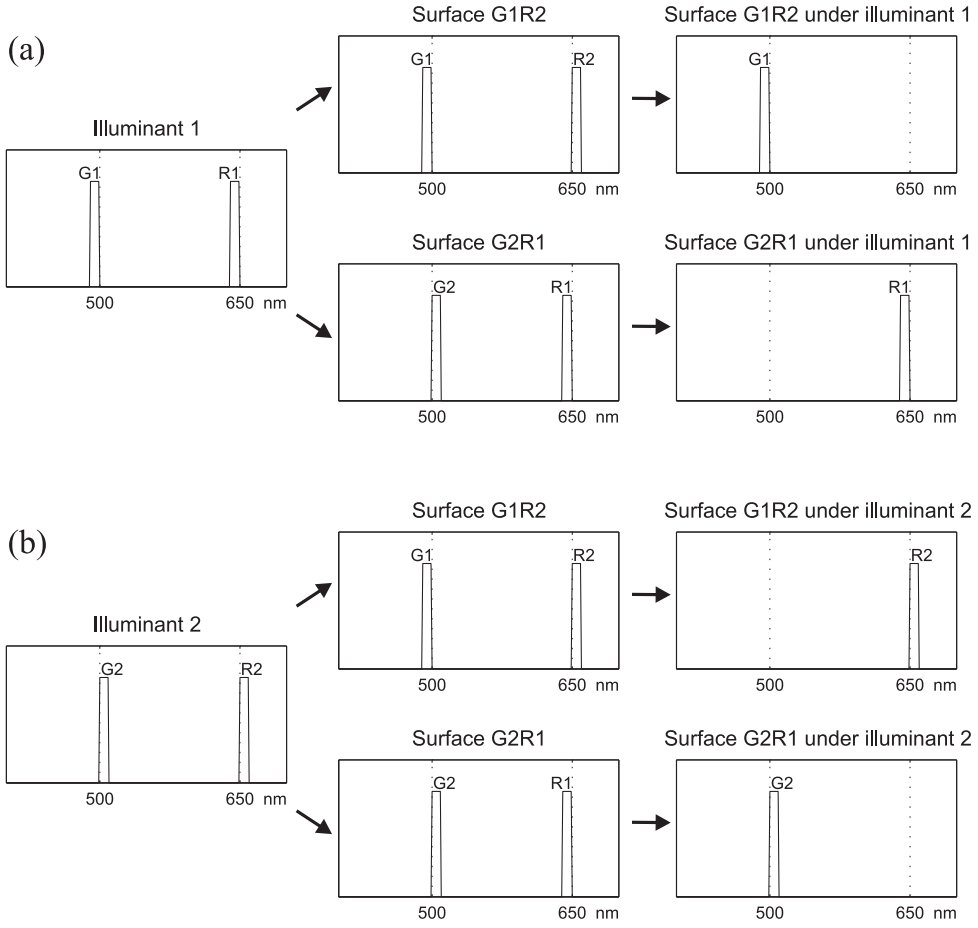

Figure 3.1: A possible scenario in the Chaotic World. (a) Illuminant 1 lighting two very similar surfaces yields very different reflected lights, a green band for surface G1R2 and a red band for surface G2R1. (b) When the same surfaces are instead illuminated by illuminant 2 (which is very similar to illuminant 1), the pattern of reflected lights is reversed.

can be arbitrarily narrow and correspondingly closely spaced). Nevertheless, the reflected chromaticities of these two (under white light indistinguishable) surfaces become dramatically different under a yellow illuminant and this difference is even complete reversed for a second yellow illuminant, which itself is visually indistinguishable from the first yellow illuminant. In the more general case when the two narrow spectral bands are not restricted to the red and green part of the spectrum, such surfaces can assume any two chromaticities, limited only by the chromaticities of their reflection bands, when a suitably contrived illuminant is switched on to light them. The chromaticities they adopt can span the physically realizable gamut, and the changes due to changing the illuminant are independent for different surfaces, in the sense that any set of reflected chromaticities can be transformed into any other set. This lack of illumination-dependent regularities in the distribution of proximal chromaticities makes it impossible to infer the current illuminant by use of any chromatic scene statistic, i.e. the estimation problem (cf. section 2.3) is unsolvable in the Chaotic World.

Of course, the compensation problem (cf. section 2.4) is unsolvable in the Chaotic World, too: the chaotic proximal changes would defeat any attempt to achieve colour constancy by any systematic remapping of retinal stimuli onto subjective appearances. If a change of illumination could be relied on to transform the proximal chromaticities in a more or less orderly way, the visual system might be able to keep the surface colours constant by reversing that orderly transformation or mapping. But you can't create order by trying to unscramble total chaos. So colour constancy is not in general possible! But of course, colour constancy does occur, because the chaotic world of this scenario is not the real world we live in. This leads to the question investigated in the next section: under what spectral constraints would a simple compensating process, such as normalization through sensitivity adjustments, give us colour constancy?

3.3 The Three-band World

For normalization-compatible changes of proximal chromaticities under changes of illumination, some extremely artificial conditions must be satisfied (Worthey & Brill, 1986). First, the three cone types must respond to non-overlapping bands in the illuminant spectra. This will happen if the cone spectral sensitivities have no overlap, or if the illuminant energy is always confined to discrete non-overlapping bands, with one cone type sensitive to each such band. Second, illuminants must never vary significantly in their relative distribution of power across any one such band; they can vary only in the amounts of power they radiate in the three bands. (Alternatively but still less plausibly, the surface reflectances, rather than the illuminant power spectral densities, could be uniform in this sense for each band.)

Under this "three-band" scenario, the cone excitations for each cone type vary in proportion to the total illuminant power within the corresponding spectral band, and this variation can be compensated — for all surfaces at once — by a reciprocal adjustment of sensitivity appropriate to the scene as a whole. Thus the Three-band World exemplifies normalization-compatible mapping (cf. section 2.4): the possible illumination-dependent mappings from surface reflectance to cone excitations differ only in the scaling factors applied to the three cone excitations.

A convenient way of picturing the effects of changes in intensity and in colour balance of the illuminant is to consider a logarithmic cone excitation space in which three coordinates represent the excitations of the long-, mid- and short-wavelength cones on a log scale (see section 2.4). The position of the point representing any surface moves parallel to the positive diagonal under variations in illuminant intensity. If the colour of the illuminant changes, then surfaces move in the two orthogonal, "chromatic" directions as well. But in the Three-band World, the constellation of points representing a set of surfaces in a scene moves rigidly under all changes in illumination. This happens because a change of illuminant changes the excitation of, say, the long-wavelength cones by exactly the same factor for all surfaces in the scene — the factor by which the illuminant power in the red band is changed. Each cone excitation is scaled by the power being supplied within its associated spectral band. The logs of these three cone-specific scaling factors are added to the corresponding coordinate in the logarithmic plot.

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

3.4 The Linear World

In the search for order in the chromatic universe, the approach usually taken is to approximate naturally occurring spectral reflectance functions and spectral power distributions by a weighted sum of suitably chosen basis functions, the latter forming a fixed set common to all spectra (e.g. Sällström, 1973; Brill, 1978; Buchsbaum, 1980; Maloney, 1986, 2003). To account for trichromatic perception of surface colour, we need, minimally, three basis functions for surface

reflectance. Schemes like this that use linear models of reflectance and power distributions have been investigated quite a lot. And one of the first things to emerge is that the kind of normalization that worked for constancy in the three-band scenario won't do at all, in principle, for the linear world. This is not surprising, since it is only in the hopelessly artificial Three-band World that the cone excitations from different surfaces all change by the same factor when the light changes. Yet illumination-dependent mapping in the linear world remains formally fairly simple. Each of the three surface reflectance components generates an illumination-dependent triple of cone excitations (a cone excitation vector). Each such cone excitation vector is the weighted sum of N_i fixed vectors contributed by the N_i possible illuminant components: the fixed vectors are the cone excitation triples resulting from the interaction of the illuminant component (in unit quantity) with the surface component under consideration, and the weights are determined by the particular illuminant's composition. The visual system can recover an illumination-invariant representation of surface colour (its composition in terms of the three basis functions for reflectance) if it multiplies the cone excitation triple by the inverse of an illuminant-specific 3-by-3 matrix that contains the cone excitation vectors that would be generated by unit values of each surface component under the prevailing illumination.

The approximation of natural reflectance functions with a combination of three basis functions might not be considered satisfactory (it for instance accounts for only 60% of the chromatic variance in a haphazard sample of natural surface reflectance functions spectro-radiometrically measured by Brown (2003)), but the linear world is at least a great improvement over the Three-band World in this respect, and as a result it has come to dominate discussions of colour constancy from a computational perspective. Yet the linear world has serious shortcomings as a theoretical framework.

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

this is more than is necessary or desirable for characterizing them if three are sufficient for the functions that generate them (unless, of course, the linear model happens to be an accurate idealization of reality in this respect — a point that has not been established).

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

The theoretical value of the linear world scenario is also limited. The formal simplicity of the compensation operation (matrix division) is not supported by any known or plausible basis in visual processing. In particular, there is a lack of evidence that the visual system does in any sense internalize the matrices of coefficients specifying cone excitations for different combinations of surface and illumination components, or that its colour corrections are consistent with such a computational scheme. In general, despite extensive work employing the linear world as an analytical tool, it does not seem to have produced much insight either into the actual principles of illumination-dependent mapping in the environment, or into the visual processes subserving colour constancy.

The fondness of theorists for the linear idealization of the chromatic universe may in part reflect the lack of an alternative. In the ideal idealized world, the appropriate illumination-dependent or compensatory transformations would be more intuitively understandable and mechanistically plausible, as is the case with simple normalization, and regularity would be guaranteed by restricting the model to smooth (and not arbitrary) functions throughout. Such an alternative is introduced next.

3.5 The Gaussian World

In MacLeod and Golz (2003) Donald I. A. MacLeod and I have presented a model in which natural spectral reflectance functions, the spectral power density of illuminants, as well as the cone sensitivity functions are approximated by Gaussian functions:

- Surface spectral reflectance function for surface S:

$$S(\lambda) = S_{max} exp[-k_S(\lambda - \lambda_S)^2] \tag{3.1}$$

- Illuminant spectral power density for illuminant I:

$$I(\lambda) = I_{max} exp[-k_I(\lambda - \lambda_I)^2] \tag{3.2}$$

- Cone sensitivity function for cone C:

$$C(\lambda) = C_{max} exp[-k_C(\lambda - \lambda_C)^2] \tag{3.3}$$

In this model, surfaces, illuminants and cone sensitivities are completely characterized by three parameters respectively: the centroids λ_S, λ_I, λ_C (where λ_C has a different value for each cone type), the bandwidths (or more precisely, signed spectral curvatures) k_S, k_I, k_C, and the values at the centroid S_{max}, I_{max}, and C_{max}. Depending on the sign of the bandwidth parameter, the spectral functions are either Gaussians, (with exponentials as a special case) or the reciprocals of Gaussians. With integration over a finite spectral range, one needs not be troubled by the fact that the latter functions approach infinity in the limit of infinite wavelength.

Like the linear world with three degrees of freedom for surfaces and illuminants, the Gaussian World is a fully (but minimally) trichromatic world. Special cases of the linear world may involve a principled choice of basis functions and guarantee smooth spectra. One such scheme (cf. Moon & Spencer, 1945) approximates natural spectral reflectance functions and spectral power distributions by polynomials in wavelength or some other spectral variable. Trichromacy requires at least second-order polynomials for the surfaces, so a second-order polynomial constraint on the surfaces and illuminants yields a minimal trichromatic world. The Gaussian World differs from this in just the following way: it is the logarithms of the spectral distributions that are built up from additive components. That is, the log of the radiant power or of the surface reflectance is always describable as a second order polynomial. This representation is an equally crude approximation to reality as the Linear World. But the benefits of the Gaussian World are many:

- Any spectral function in the Gaussian World is physically realizable. One does not have to worry about negative light emissions or negative reflectances, since the numbers represented by their logarithms cannot become negative. And because of their more natural approach to zero the Gaussian descriptions fit natural spectra better than do the comparable linear ones.

- Monochromatic radiations can be included in the chromatic universe as a limiting case, something not possible in the Linear World.

- There is an important gain in mathematical simplicity and tractability: the proliferation of degrees of freedom that occurs in the Linear World in the progression from light and surface to the visual stimulus is here entirely avoided. When incident light is multiplied by surface reflectance, the second-order polynomials descriptive of the log power and the log reflectance simply add to form the log of the retinal stimulus, making this, too, a second-order polynomial (in non-logarithmic terms, the product of the Gaussians for surface and illuminant yields another Gaussian.) Thus the function specifying the proximal stimulus still has the same form as the surface spectral reflectance and the illuminant spectral power distribution, and the same three degrees of freedom (each of its three parameters being a function of the surface and illuminant parameters), whereas in the Linear World it has six.

- Each stimulus chromaticity is associated with the same spectral energy distribution no matter what the light and surface reflectance that combined to generate it, so indistinguishable stimuli are now consistently represented as indistinguishable (i.e. they have the same parameter values).

- Most important, the cone spectral sensitivities themselves can be approximated by Gaussians, and this gives us something not available in the other theoretical worlds: a simple algebraic solution (cf. Equation 3.5 on page 28) for the cone excitations generated by any surface and illuminant, where the only variables are the parameters specifying cone sensitivity, reflectance and illuminant.

I will now present this algebraic solution and derive specific predictions for illumination-dependent regularities in the Gaussian World.

3.5.1 Algebraic solution for retinal stimulation

Using the Gaussian World representations for the spectral functions of cones, surfaces, and illuminants (Equations 3.1 – 3.3) introduced on page 25, the cone excitation elicited when cone C samples surface S under illuminant I is given by:

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

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

$$\begin{aligned} E(I,S,C) &= I_{max}S_{max}C_{max}exp[-(k_I\lambda_I^2 + k_S\lambda_S^2 + k_C\lambda_C^2)] \\ &\times \int_{-\infty}^{\infty} exp[-(k_I + k_S + k_C)\lambda^2 + 2(k_I\lambda_I + k_S\lambda_S + k_C\lambda_C)\lambda]d\lambda \\ &= I_{max}S_{max}C_{max}exp\left[-(k_I\lambda_I^2 + k_S\lambda_S^2 + k_C\lambda_C^2) + \frac{(k_I\lambda_I + k_S\lambda_S + k_C\lambda_C)^2}{k_I + k_S + k_C}\right] \\ &\times \int_{-\infty}^{\infty} exp\left[-\left(\sqrt{k_I + k_S + k_C}\lambda - \frac{k_I\lambda_I + k_S\lambda_S + k_C\lambda_C}{\sqrt{k_I + k_S + k_C}}\right)^2\right] d\lambda \qquad (3.4) \\ &= k_0 exp\left[-(k_I\lambda_I^2 + k_S\lambda_S^2 + k_C\lambda_C^2) + \frac{(k_I\lambda_I + k_S\lambda_S + k_C\lambda_C)^2}{k_I + k_S + k_C}\right] \end{aligned}$$

where $k_0 = I_{max}S_{max}C_{max}\sqrt{\pi/(k_I + k_S + k_C)}$. The square root factor here is the value of the Gaussian integral in Equation 3.4. $k_I + k_S + k_C$ is generally positive since the cone sensitivity functions are typically narrower than $I(\lambda)$ or $S(\lambda)$.

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

$$\begin{aligned} E(I,S,C) &= k_0 exp\left[-\frac{k_Ik_S(\lambda_I - \lambda_S)^2 + k_Ik_C(\lambda_I - \lambda_C)^2 + k_Sk_C(\lambda_S - \lambda_C)^2}{k_I + k_S + k_C}\right] \\ &= k_0 exp\left[-\frac{k_Ik_S}{k_I + k_S + k_C}(\lambda_I - \lambda_S)^2\right] exp\left[-\frac{k_Ik_C}{k_I + k_S + k_C}(\lambda_I - \lambda_C)^2\right] \\ &\times exp\left[-\frac{k_Sk_C}{k_I + k_S + k_C}(\lambda_S - \lambda_C)^2\right] \qquad (3.5) \end{aligned}$$

Equation 3.5 yields, in closed form, an expression for cone excitations in dependence of surface and illumination parameters and thus allows to analyse the chromatic changes in the proximal stimulus under changes of illumination.

3.5.2 Illumination-dependent regularities in the Gaussian World

Conditions with broad bandwidths

In Equation 3.5, cone excitation is a product of three Gaussian factors. The *third Gaussian factor* represents the surface colour independently of the illuminant. It depends on the spectral distance of the centroid of the surface reflectance to the centroid of the cone sensitivity ($\lambda_S - \lambda_C$) but not on the centroid of the illuminant (λ_I). This term is what a perfect colour constancy mechanism would have to recover in order to construct surface colour designators that are independent of illumination.

The *second Gaussian factor* depends on the illuminant colour. If surface bandwidths are not too narrow, this factor is approximately the same for all surfaces, as it depends only on the difference between the illuminant spectral centroid and the centroid of the cone sensitivity ($\lambda_I - \lambda_C$). Therefore, this factor could be successfully compensated — for all surfaces at once — by a normalization-compatible remapping of cone excitations appropriate to the scene as a whole.

The *first Gaussian factor* in the equation for cone excitation depends both on the illuminant and the surface: it increases as the distance between the centroid of the surface reflectance and the centroid of the illuminant ($\lambda_S - \lambda_I$) decreases. Thus, if the illumination changes by a shift in λ_I, those surfaces with peak reflectance closer to the spectral centroid of the new illumination will yield stronger cone excitations compared to other surfaces. In a sense this factor precludes full colour constancy based on normalization, since the normalization factor required for its compensation is different for different surfaces within a scene. Luckily though, this factor is independent of the cone type, which means that it generates variations in effective intensity only, not chromaticity. What this factor means is simply that, for instance, *when the lighting gets red, the reds get lighter*, relative to the shorter-wavelengths surfaces. Surfaces similar to the illuminant in colour are rendered as lighter than surfaces dissimilar from the illuminant in colour.

Similarly, bluish surfaces (with high s values) increase in luminance more than e.g. greyish or yellowish surfaces when the illumination becomes bluish. The fact that S cone excitation does not contribute to luminance (Eisner & MacLeod, 1980) is of no concern here: the intensity benefit for surfaces similar to the illuminant affects all three cone classes (and thus luminance, which is modelled by $L + M$) equally, since the first Gaussian factor is independent of λ_C).

Conditions with narrow bandwidths

While the above analysis assumes more or less broad bandwidths (which natural surfaces and illuminants usually do have), I will consider now the chromaticities of retinal stimulation in the Gaussian World when surfaces or illuminants become more and more spectrally selective, with monochromatic spectra as a limiting case. For this, a separate version of Equation 3.5 for each cone type will be used. Instead of the universal equation for excitation E of an abstract cone C, the cone excitations will be labelled L, M, B for the long-wavelength (L), mid-wavelength (M) and blue (B) cone respectively[1]. λ_L, λ_M, and λ_B replaces λ_C, but the same cone spectral bandwidths k_C is assumed in each case (both for simplicity, and because the bandwidths can be made equal by adopting a more suitable function of wavelength — approximately its logarithm — as the spectral variable at the outset), as well as the same value for peak cone sensitivity C_{max} (this can always be arranged by appropriate choice of the units for the cone excitations). These substitutions yield:

$$\begin{aligned} L &= k_0 exp\left[-\frac{k_I k_S}{k_I+k_S+k_C}(\lambda_I-\lambda_S)^2\right] exp\left[-\frac{k_I k_C}{k_I+k_S+k_C}(\lambda_I-\lambda_L)^2\right] \\ &\times exp\left[-\frac{k_S k_C}{k_I+k_S+k_C}(\lambda_S-\lambda_L)^2\right] \end{aligned} \tag{3.6}$$

$$\begin{aligned} M &= k_0 exp\left[-\frac{k_I k_S}{k_I+k_S+k_C}(\lambda_I-\lambda_S)^2\right] exp\left[-\frac{k_I k_C}{k_I+k_S+k_C}(\lambda_I-\lambda_M)^2\right] \\ &\times exp\left[-\frac{k_S k_C}{k_I+k_S+k_C}(\lambda_S-\lambda_M)^2\right] \end{aligned} \tag{3.7}$$

$$\begin{aligned} B &= k_0 exp\left[-\frac{k_I k_S}{k_I+k_S+k_C}(\lambda_I-\lambda_S)^2\right] exp\left[-\frac{k_I k_C}{k_I+k_S+k_C}(\lambda_I-\lambda_B)^2\right] \\ &\times exp\left[-\frac{k_S k_C}{k_I+k_S+k_C}(\lambda_S-\lambda_B)^2\right] \end{aligned} \tag{3.8}$$

Illuminant-dependent chromatic shifts in the Gaussian world are most easily analysed by considering intensity-invariant chromatic coordinates in log cone excitation space. One such coordinate is the difference in the logs of L and M cone excitation for which one obtains by use of Equation 3.6 and 3.7:

$$log(L) - log(M) = log(L/M)$$

[1]The latter is referred to elsewhere as the S cone, but here, to avoid confusion with the surface identifier, as the B cone with excitation B.

$$= \frac{2(\lambda_L - \lambda_M)k_C}{k_I + k_S + k_C}\{k_I[\lambda_I - (\lambda_L + \lambda_M)/2] + k_S[\lambda_S - (\lambda_L + \lambda_M)/2]\} \tag{3.9}$$

This expression gives a particularly simple mathematical embodiment to the trade-off between illuminant colour and surface colour in the determination of the retinal stimulus, since illumination and surface colour here contribute two symmetrical terms, derived respectively from the second and the third Gaussian factors in Equation 3.6 and 3.7. (The k_0 factor and the first of the three Gaussian factors have been cancelled, being equal for all cone types). Similarly, the second chromatic coordinate, the difference in the logs of B and L cone excitation results in:

$$\begin{aligned} log(B) - log(L) &= log(B/L) \\ &= \frac{2(\lambda_B - \lambda_L)k_C}{k_I + k_S + k_C}\{k_I[\lambda_I - (\lambda_B + \lambda_L)/2] + k_S[\lambda_S - (\lambda_B + \lambda_L)/2]\} \end{aligned} \tag{3.10}$$

Equations 3.9 and 3.10 allow to locate stimuli according to the bandwidth of the corresponding surface and illuminant in the chromatic plane ($log(L) - log(M)$, $log(B) - log(L)$):

1. Under spectrally flat illumination ($k_I = 0$), spectrally non-selective surfaces ($k_S = 0$) are at the origin $(0, 0)$ of this chromatic plane since the illuminant term and the surface colour term involving λ_I and λ_S in the parentheses of Equation 3.9 (and analogously for Equation 3.10) are both zero.

2. For selective surfaces, the nonzero surface colour term generates the stimulus coordinates

$$\left(\frac{2(\lambda_L - \lambda_M)k_S k_C}{k_S + k_C}[\lambda_S - (\lambda_L + \lambda_M)/2],\ \frac{2(\lambda_B - \lambda_L)k_S k_C}{k_S + k_C}[\lambda_S - (\lambda_B + \lambda_L)/2]\right)$$

3. Under chromatic illumination, the illuminant term becomes nonzero and locate spectrally non-selective surface at point

$$\left(\frac{2(\lambda_L - \lambda_M)k_I k_C}{k_I + k_C}[\lambda_I - (\lambda_L + \lambda_M)/2],\ \frac{2(\lambda_B - \lambda_L)k_I k_C}{k_I + k_C}[\lambda_I - (\lambda_B + \lambda_L)/2]\right)$$

Thus, if the illumination changes from a non-selective to a selective one, non-selective surfaces shift from the origin to the coordinates given above in 3. Selective surfaces undergo a comparable displacement. In the important range

of conditions where the illuminant bandwidth and surface bandwidth are both large enough, relative to the cone bandwidth, to keep the factor $k_C/(k_I+k_S+k_C)$ close to 1, the displacement is approximately the same for all surfaces. Under such conditions, normalization achieves approximate colour constancy.

Deviations from this normalization-compatible translation take two forms: gamut compression due to illuminant bandwidth restriction and shift resistance due to surface-reflectance bandwidth restriction.

Gamut compression. If illuminant bandwidths differ, the narrower band illuminants (with high k_I) reduce chromatic variance, compressing the constellation of stimulus chromaticities by the factor $k_C/(k_I + k_S + k_C)$ toward the locus of a neutral (non-selective) surface under the prevailing illuminant. In the limit of narrow illuminant bandwidth, surface stimuli all assume the chromaticity of the monochromatic illuminant.

Shift resistance. Likewise if surface bandwidths differ, variation in k_S reduces the shift of the narrower band surfaces toward the illuminant chromaticity, by this same factor. In the limit of narrow surface bandwidth (monochromatic reflectance), changing illumination makes no difference to the chromaticity of the stimulus.

3.5.3 Summary

Illumination-dependent mapping in the Gaussian World has the following characteristics:

1. *Normalization-compatible mapping for chromaticity.* Independent normalization of the cone photoreceptor excitations can recover the chromatic aspects of surface reflectance accurately for sufficiently broad-band surfaces and illuminants. In that sense, normalization could yield colour constancy in this idealized world. Normalization-compatible mapping of chromatic values has a simple geometric expression: normalization-compatible mapping of both chromaticity and intensity implies rigid translation in $log(L, M, S)$ cone excitation space under changes of illumination. But normalization-compatibility is only approximate, and the analysis in this section has also shown how the Gaussian World violates normalization-compatibility in three distinct and quantitatively specifiable (but generally minor) ways:

2. *When the lighting gets red, the reds get lighter.* There remains (after normalization) a deviation from constancy, occurring for both broad-band and narrow-band colours, that affects stimulus intensity alone: surfaces similar to the illuminant in colour are rendered as lighter than surfaces dissimilar from the illuminant in colour. While it precludes rigid translation in log cone excitation space, this does not affect chromatic values and thus preserves rigid translation in the plane of intensity-invariant chromatic coordinates used in Section 3.5.2. For narrow-band conditions, normalization fails to restore even the chromatic values correctly:

3. *Resistance of narrow-band surfaces to chromaticity shift.* Stimuli from narrow-band surfaces are more resistant to shifts in log cone excitation space than stimuli from broad-band surfaces, with illumination-invariance as a limit approached for monochromatic reflectances (these remain at the spectrum locus).

4. *Gamut compression with narrow-band illuminants.* For narrow-band illuminants, the spectrum locus is again approached as a limit, but this time for all surfaces, so the gamut is compressed.

Since discussions framed in the linear world have emphasized that normalization is in principle inadequate, the approximate validity of normalization in the Gaussian World is somewhat surprising. In the next chapter, I will consider the illuminant-dependent regularities found in images of real natural scenes and compare them with the predictions of the Gaussian World.

Unlike the latter two deviations from strict normalization compatibility (3. and 4. above), the increased intensity for surfaces similar to the illuminant in colour (2. above) is potentially significant even in the important case of broad-band illuminants and surfaces and thus represents a promising candidate for a cue of the illuminant under natural conditions. In chapter 4 it will be shown that this regularity is found to hold also in images of natural scenes and in chapter 5 evidence will be presented that this regularity is indeed exploited by the visual system. As we will see, the relation between the the shift resistance of narrow-band surfaces found in the Gaussian world (3. above) and corresponding results of the real world simulation as well as the experiments are not that straight forward.

In summary, the Gaussian World idealization has been proved as a promising heuristic for chromatic cues for the perception of surface and illumination colour since that model of the environment is tightly constrained, allows an algebraic analysis of the chromatic interplay of surfaces and illuminants and provides a rich set of easily testable predictions.

Chapter 4

The real world

4.1 Which chromatic cues are available in the natural environment?

In this chapter, I will address the following question: which statistics of the chromatic distribution in images of natural scenes depend systematically on the illumination and are thus potential cues for the illumination that the visual system could exploit? Different from the previous chapter where a similar question has been addressed by analysing an idealized model of the natural environment (the Gaussian World), I will here analyse images taken from real natural scenes for which different illuminations are simulated. In doing so, particular reference will be made to the regularities of chromatic shifts under changes of illumination predicted by the Gaussian World. This analysis can thus be regarded as a test whether the findings of the Gaussian World analysis can be confirmed for the real world. But as we will see, the real world analysis accomplishes more: the way in which a theoretical regularity emerges in the real chromatic world can depend on details of the distribution of chromatic properties in the natural environment, contingencies that cannot fully be taken into account in a theoretical analysis. Especially the size of illumination-dependent chromatic shifts cannot be predicted by the theoretical analysis alone. And finally, the real world analysis will yield important insight into the way in which these regularities relate to each other in the natural environment.

The presentation of the results of the real world analysis are divided into two parts:

- In the first part (section 4.3.1), the chromatic distribution of individual surfaces and the chromatic shifts that they undergo when the illumination

changes are analysed (i.e. the results are presented on the level of individual pixels of the images).

- In the second part (section 4.3.2), statistics of the chromatic distribution of entire images that could capture these regularities are analysed with respect to their usefulness for estimating the illuminant (i.e. the results are presented on the level of scene statistics).

4.2 Methods

4.2.1 Hyperspectral images of natural scenes

In order to gather samples of real world images for analysing chromatic properties of the environment, the approach usually taken nowadays by colour scientists is to use a hyperspectral camera. Different from a common digital camera that captures images by means of three types of light sensors with different but broad and overlapping spectral sensitivities, a hyperspectral camera takes a large series of images for different spectral bands of the same scene, all narrow bands together covering the whole visible spectrum. This series of images, each capturing the radiation intensity of the scene in one spectral band, is called a hyperspectral image and represents the proximal spectral stimulation elicited from the scene under the illuminant that was present when the image was taken. If one discounts the spectrum of that illuminant from the hyperspectral image, one can derive an estimate of the spectral surface reflectance at each pixel of the scene. For this, usually a standard white surface with known reflectance function is included in the scene, and the illuminant that can be inferred from the reflected light of this standard white surface is assumed to light all patches of the scene equally. The estimated surface reflectance function of each pixel is then obtained by dividing, wavelength by wavelength, the measured radiance spectrum of that pixel by the inferred illuminant spectrum.[1]

For the following analysis of chromatic regularities in natural scenes, I have used the data of Ruderman, Cronin, and Chiao (1998) who obtained hyperspectral images for twelve views of foliage-dominated natural environments in several locations, including the vicinity of Baltimore, Maryland (temperate woodland), and Brisbane, Australia (sclerophyll forest, subtropical rainforest, and mangrove

[1]The validity of the so corrected hyperspectral image as an estimate of the reflectance spectra of the surfaces in the scene depends on the degree to which the illumination in the scene is indeed homogeneous, and decreases e.g. with occurrences of shadows, mutual illumination among surfaces that are oriented towards each other, and filtering of the illumination by leaves. So far, this issue has not been investigated quantitatively.

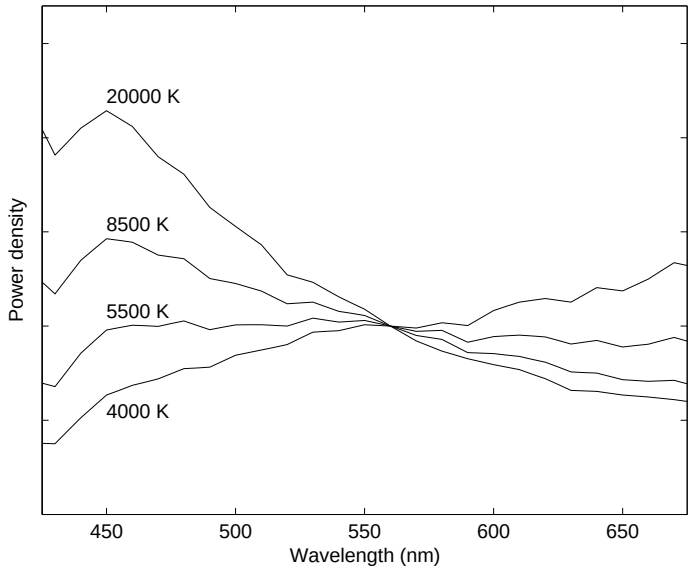

Figure 4.1: Spectra of the four CIE daylights applied to the natural images.

swamp). The scenes contained numerous natural objects, including leaf foliage, bark, rocks, herbs, streams, bare soil, etc. Unlike many other sets of hyperspectral natural images (e.g. Webster & Mollon, 1997; Páraga, Brelstaff, Troscianko, & Moorehead, 1998), the images of Ruderman et al. (1998) do not include areas with direct view of the sky, which make this set a good choice for the purpose of the present real world simulation. Their hyperspectral images have a resolution of 128 × 128 pixels (with angular resolution of 0.047 × 0.055 degrees per pixel, horizontal × vertical), yielding a data set comprising nearly 200 000 pixels. Each images was taken in 43 successive wavelength steps at 7-8 nm intervals from 403 to 719 nm trough a variable interference filter with a half bandpass typically of 15 nm at each wavelength. Plate II shows the twelve scene contained in this hyperspectral image set of Ruderman et al. (1998).

4.2.2 Simulation

In order to simulate the proximal stimuli reflected from the natural scenes under different illuminants, four daylight spectra with correlated colour temperatures 4000 K, 5500 K, 8500 K and 20 000 K, covering the extreme range of unimpeded daylight illuminants, were applied to each hyperspectral reflectance image. The spectral functions of these illuminants (Fig. 4.1) were generated using a CIE standard procedure and set of basis functions (Wyszecki & Stiles, 1982, p. 145).

After applying the illuminants to each surface reflectance image, the cone responses to each pixel in these simulated images of reflected lights were derived using the energy basis cone sensitivity functions of Stockman et al. (1993). For these calculations, the format of all discretized spectral functions (for surface reflectances, illuminant power densities, and cone sensitivities) consisted in steps of 1 nm and ranged from 380 nm to 780 nm. All functions that originally had a coarser wavelength resolution were interpolated to this format by spline interpolation.

The vector containing the power density values of an illuminant in ascending order from 380 nm to 780 nm (i.e. the discrete counterpart of the continuous function $I(\lambda)$) is labeled $\vec{I}$. Similarly, $\vec{S}$ denotes the vector of the surface reflectance values and $\vec{l}$, $\vec{m}$, and $\vec{s}$ denote the vectors containing the spectral sensitivities for cone L, M, and S. $\vec{S}_{x,y}$ is subjoined with an index to denote the surface reflectance vector of the pixel at location (x, y) and $\vec{I}_i$ denotes the power density vector of illuminant i (with $i = 1, \ldots, 4$ for the illuminants of 4000 K, 5500 K, 8500 K, and 20 000 K colour temperature).

The proximal light spectrum $\vec{P}_i^{x,y}$ reflected from pixel (x, y) under illuminant i is computed by:

$$\vec{P}_i^{x,y} = \vec{S}_{x,y} * \vec{I}_i$$

where the operator $*$ denotes component by component (i.e. wavelength by wavelength) multiplication. The cone excitations elicited by this proximal light from pixel (x, y) under illuminant i are then obtained by computing the dot products (also called scalar or inner products) of the proximal light vector with the respective cone sensitivity vector, which yields the discrete counterpart of integration:

$$\begin{aligned} L_i^{x,y} &= \vec{P}_i^{x,y} \bullet \vec{l} \\ M_i^{x,y} &= \vec{P}_i^{x,y} \bullet \vec{m} \\ S_i^{x,y} &= \vec{P}_i^{x,y} \bullet \vec{s} \end{aligned}$$

The colour codes I mainly consider in the following analysis are the log of luminance and the logs of the intensity-invariant chromaticity coordinates (l, s) of MacLeod and Boynton (1979). The latter are defined by $l = L/(L + M)$ and $s = S/(L + M)$, with units for s chosen to make it unity for equal energy white.

The resulting distributions of colour codes for individual pixels under the four different illuminants are analysed in section 4.3.1. In a second step (cf. 4.3.2),

this information is analysed on the level of statistics calculated across the entire area of each of the twelve scenes. For instance, the mean redness of a scene under illuminant i was calculated by space averaging across all $128 * 128$ pixels:

$$mean(log(l))_i = \frac{\sum_{x=1}^{128}\sum_{y=1}^{128}\log(l_i^{x,y})}{128*128} = \frac{\sum_{x=1}^{128}\sum_{y=1}^{128}\log(\frac{L_i^{x,y}}{L_i^{x,y}+M_i^{x,y}})}{128*128}$$

The following chromatic scene statistics were calculated for each of the twelve scenes under each of the four illuminants:

- Mean redness, blueness and luminance:
  - $mean(log(l))$
  - $mean(log(s))$
  - $mean(log(luminance))$
- Standard deviations:
  - $std(log(l))$
  - $std(log(s))$
  - $std(log(luminance))$
- Correlations:
  - $correlation(log(l), log(luminance))$
  - $correlation(log(s), log(luminance))$
  - $correlation(log(l), log(s))$
- Fischer skewnesses:
  - $skewness(log(l))$
  - $skewness(log(s))$
  - $skewness(log(luminance))$

Tables providing the values of these scene statistics of all twelve scenes under the four illuminants used are presented in appendix C.2.

4.3 Results

In this section, I will present the chromatic regularities found in the simulated real world scenario, first by analysing the chromatic shifts of individual pixels under changes of illumination and then by considering the level of scene statistics for the chromatic distribution of entire scenes. For this, particular reference is made to the predictions of the Gaussian World (cf. section 3.5.2).

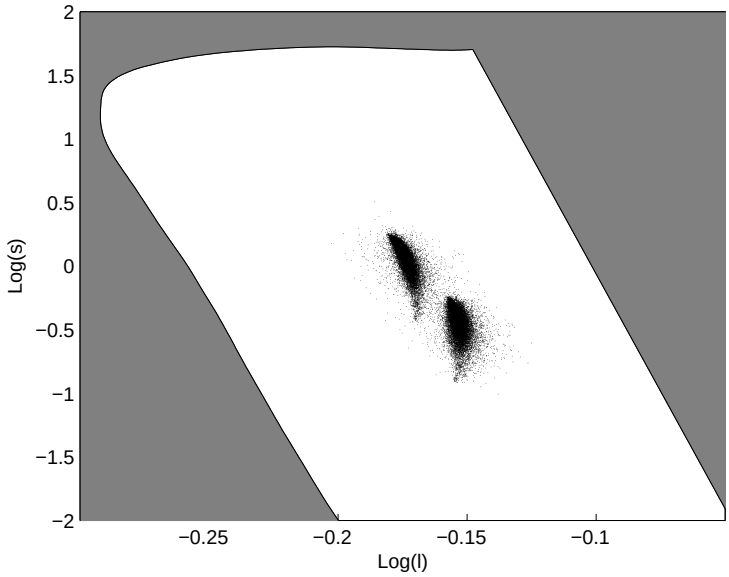

Figure 4.2: Distribution of pixel chromaticities for one of the twelve natural scenes (Park4) of Ruderman et al. (1998) under two of the simulated illuminants: 20 000 K daylight (upper left cluster of data points) and 4000 K (lower right cluster).

4.3.1 Chromatic distributions of individual pixels

As a first orientation for the distribution of chromaticities in the natural images of Ruderman et al. (1998) and for their shifts under changes of illumination, Fig. 4.2 shows the distribution of all pixel chromaticities of one typical scene (their Park4) under two illuminants (4000 K and 20 000 K colour temperature).[2] The chromaticities reflected from this scene move in the ($log(l)$,$log(s)$) chromaticity plane from the lower right cluster to the upper left cluster when the illumination changes from 4000 K daylight to 20 000 K daylight. Thereby, all pixels undergo a comparable shift, yielding a rigid (but not exactly rigid) translation of the entire scene cluster. This raises the question, to what degree are the illumination-dependent shifts in the natural images normalization-compatible?

Normalization-compatible chromatic translation

A perfectly normalization compatible mapping implies an exactly rigid translation in $log(L, M, S)$ space under changes of illumination (cf. section 2.4). In Fig. 4.3, these logarithmic coordinates are plotted for all pixels of one typical

[2]Corresponding figures for all twelve scenes are presented in Appendix C.1.

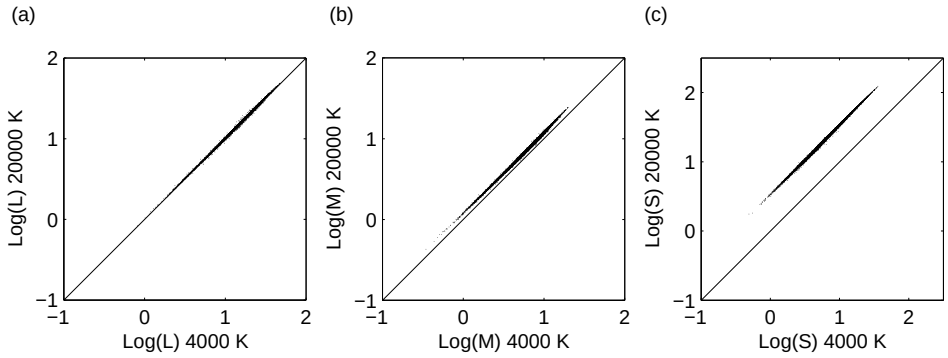

Figure 4.3: Cone excitation values for individual pixels of the Park4 scene of Ruderman et al. (1998), under 4000 K illumination (horizontal axis) and under 20 000 K illumination (vertical axis). Normalization-compatible mapping requires proportionality between the values, or clustering of the points around lines of slope 1 in these double-logarithmic plots.

scene (Park4) under 4000 K daylight against the corresponding values under 20 000 K daylight. The translation due to changing the illumination is close to normalization-compatibility: the pixels cluster closely around lines of slope 1, which indicates that the pixels all undergo very similar shifts. These shifts in log cone excitations are largest for the S cone: averaged over twelve scenes, the mean factor (± 1 standard deviation) of the linear scaling of the cone excitations under an illumination change from 4000 K to 20 000 K is 1.003 (± 0.016), 1.19 (± 0.023), and 3.39 (± 0.055) for the L, M, and S cone respectively.

These results of nearly normalization-compatible shifts in the real world simulation support those of Dannemiller (1993) based on Krinov's average terrain data, and of Foster and Nascimento (1994) based on Munsell papers and extend them to individual elements of natural scenes. But different from these works, here I additionally analyse the $log(l)$, $log(s)$ and $log(luminance)$ measures (Fig. 4.4). This separate treatment of the chromatic axes is indicated because approximately rigid displacements measured on the coarse scale required for displaying the individual cone excitations, or luminance, do not imply (perceptually) rigid displacement along the much more sensitive chromatic axes. While the validity of the rigid translation approximation appears better for the luminance axis (Fig. 4.4c) than for the chromatic ones, it should be noted that the axis for $log(l)$ has a very different scale from the other two, due to the very small variance of this quantity in natural images (Ruderman et al., 1998) which is accompanied by a commensurate sensitivity of the visual system to this variable (MacLeod & von der Twer, 2003).

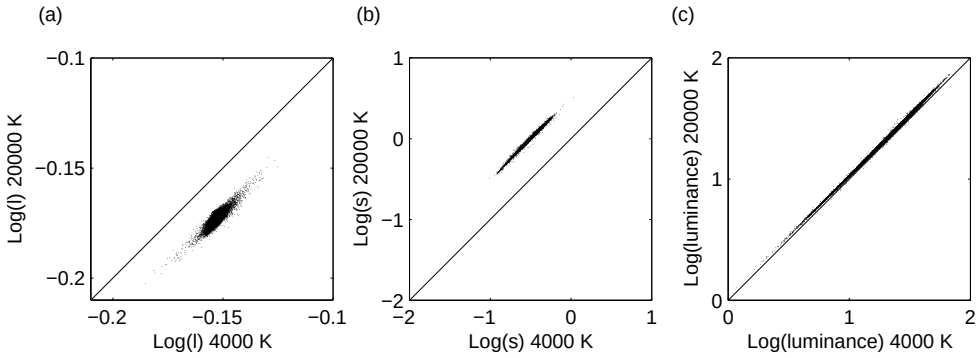

Figure 4.4: Chromaticity and luminance values for individual pixels of the Park4 scene of Ruderman et al. (1998), under 4000 K illumination (horizontal axis) and under 20 000 K illumination (vertical axis). For these double-logarithmic plots, too, normalization-compatible mapping requires clustering of the points around lines of slope 1. Note the very different scale for the different plots.

But even for these highly sensitive chromatic axes (as well as for luminance) the rigidity principle appears to be a useful approximation. Averaged over twelve scenes, the standard deviation of the changes in the (decimal) log of l (among the pixels of one scene) is only 0.0024; this means that the factor by which l changes is stable across pixels with a standard deviation of 0.55% of its mean. For s and for luminance, the standard deviation is 1.8%; all three deviations from rigid translation are small enough to be undetectable or barely detectable, hence perceptually inconsequential. The real world, then, is far from chaotic; and while it is not perfectly orderly like the three-band world, on this evidence it exhibits normalization-compatible mapping to a perceptually acceptable approximation. The Gaussian idealization therefore appears to be a viable model in this respect.

When the light gets red, the reds get lighter

Fig. 4.5 shows for the same scene the change in luminance of individual pixels in going from the bluish 20 000 K illuminant to the reddish 4000 K illuminant, plotted versus the pixel's l coordinate (under a neutral 7000 K illumination). The chromaticity-dependent intensity shift predicted in the Gaussian World (cf. section 3.5.2) is present (in this scene as well as in all others of the twelve scenes): the reds (mostly) get lighter and the blues and greens dimmer. As Fig. 4.5 shows, the higher the l value of a pixel, the more it increases in luminance when the illumination takes on a higher l value (i.e. it changes from 20 000 K to 4000 K).

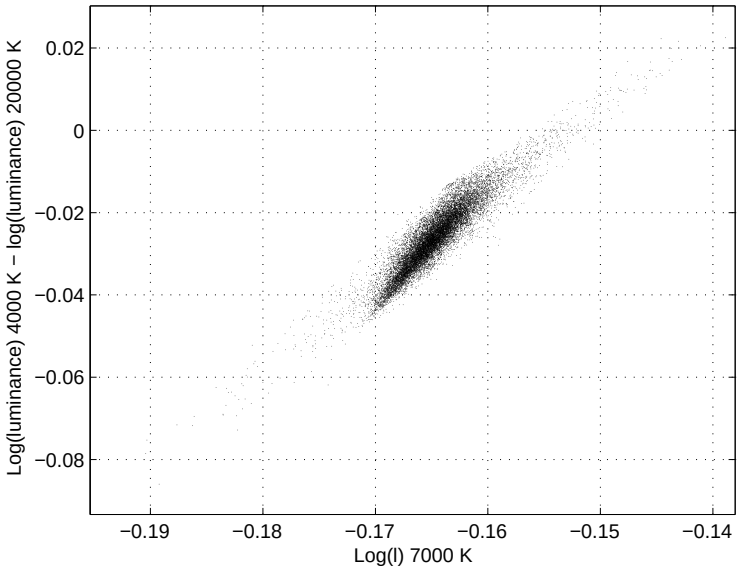

Figure 4.5: Change in log(luminance) for individual pixels of scene Park4 when the 4000 K illuminant replaces the 20 000 K illuminant, plotted versus redness of the pixels (as measured by $\log(l)$ under 7000 K).

Similarly, when the bluish 20 000 K illuminant is replaced by the reddish 4000 K illuminant, i.e. the illumination takes on a lower s value, the pixels with high s values are most reduced in luminance (cf. Fig. 4.6). The magnitude of these effects is only a few percent, but it is a major source of all departure from normalization-compatible mapping for the luminance axis (since the departures are themselves small). Again, the Gaussian idealization appears to be a viable model, as it has predicted these two effects (cf. page 29).

Shift resistance

While it is inevitable on physical grounds that the light reflected from surfaces with sufficiently narrow-band reflectance will undergo a smaller chromaticity shift toward the illuminant, this fact cannot be exploited in the perception of illumination-independent surface colour unless the observers can recognize the differences in surface bandwidth. The only information available for that purpose is the triple of cone excitations generated by the surface. This means that only those bandwidth effects that are correlated with surface colour are relevant. Three such surface bandwidth effects emerged in the real world simulations.

The first is illustrated for the Park4 scene in Fig. 4.7. In this and in the other scenes, surface-reflectance bandwidth (or more precisely, signed spectral curva-

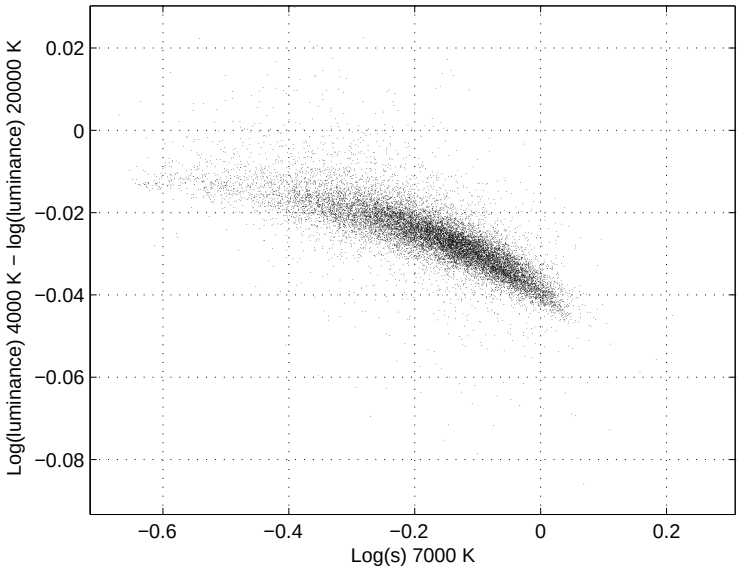

Figure 4.6: Change in log(luminance) for individual pixels of scene Park4 when the 4000 K illuminant replaces the 20 000 K illuminant, plotted versus blueness of the pixels (as measured by log(s) under 7000 K).

ture, as defined in section 3.5) tends to vary markedly along the s axis, with narrower bandwidth for the abundant yellowish colours than for whites, or for the relatively infrequent purplish colours, which tend to have concave-upward spectra (negative spectral curvature). Accordingly, Fig. 4.7 shows a result that would be surprising, if not the Gaussian World had predicted it (cf. page 32): the illumination-induced shift in log(l) is greater for surfaces with high log(s). While not large, this deviation from rigid logarithmic translation is enough, in many scenes, to skew the distribution of pixels in the ($log(l)$,$log(s)$) plane in just the manner visible for scene Park4 by comparing the two pixel clusters in Fig. 4.2.

Secondly, the simulations revealed for nine of the twelve scenes a weak tendency for the lighter surfaces to be more stable in chromaticity under a change of illumination. Fig. 4.8 shows this effect for scene Treetop4, where it is stronger than in all other scenes, as Table 4.1 shows. This effect was largely mediated by the fact that the lighter surfaces are more yellowish in these scenes and hence undergo a smaller illumination-induced shift as a result of the first bandwidth effect.

Thirdly, the redward shift in stimulus chromaticity under a reddish illuminant is ultimately limited by the spectrum locus for the reddest surfaces in any scene. One might therefore suppose that a shift to reddish illumination would shift the already reddish pixels less toward red than greenish pixels. The distribution of chromaticities would in this way get skewed away from the chromaticity of the

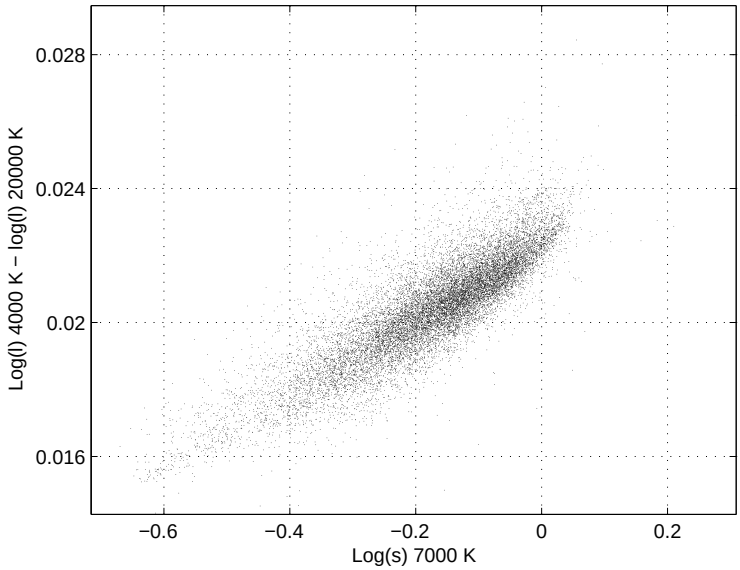

Figure 4.7: Change in redness ($\log(l)$) for individual pixels of scene Park4 when the 4000 K illuminant replaces the 20 000 K illuminant, plotted versus blueness of the pixels (as measured by $\log(s)$ under 7000 K).

Scene	*Correlation*	*Scene*	*Correlation*
Bank	0.028	Park2	-0.367
Barrine2	-0.053	Park4	0.242
Cootth12	0.239	Rainfor5	-0.214
Coottha8	-0.038	Riverrd2	-0.519
Glade	-0.133	Treetop4	-0.575
Hillside	-0.225	Wynnam4	-0.075

Table 4.1: Size of the second bandwidth effect (as measured by the correlation between the shift in $\log(l)$ in going from 20 000 K illumination to 4000 K illumination and log(luminance) under 7000 K) for each of the twelve scenes of Ruderman et al. (1998).

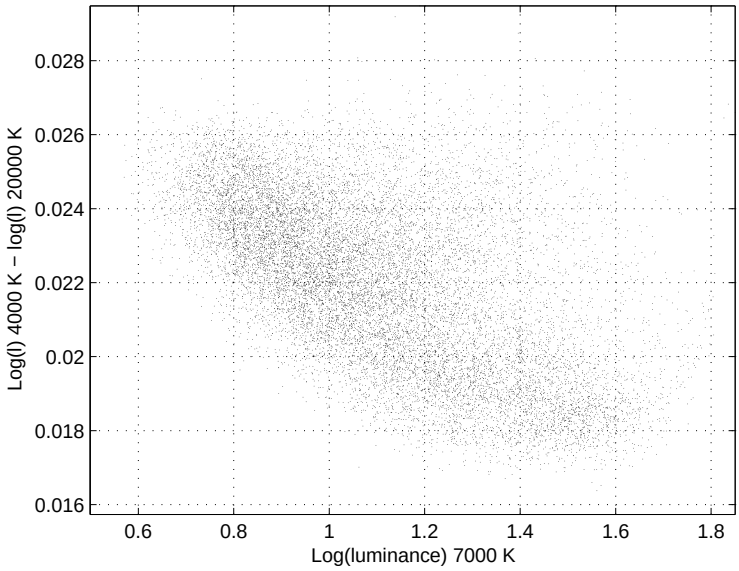

Figure 4.8: Change in redness ($\log(l)$) for individual pixels of scene Treetop4 when the 4000 K illuminant replaces the 20 000 K illuminant, plotted versus luminance of the pixels (as measured by log(luminance) under 7000 K).

illuminant.

The results of the simulation actually reveals an opposite trend. For each of the twelve natural scenes the shift in $\log(l)$ in going from 20 000 K illumination to 4000 K illumination is higher for pixels with high redness (as measured by $\log(l)$ values under 7000 K illumination) — if one considers separately subsets of pixels of similar blueness (similar values of the chromaticity coordinate $\log(s)$), as Fig. 4.9 shows for one slice of pixels with medium blueness of scene Park4. This may be due in part to the abundance in the natural scenes of greenish vegetation, which tends to have a relatively narrow reflectance band, and consequently to be resistant to chromaticity shift. In any case, this relation does not emerge clearly in whole set of pixels for a scene because of the dependence of the shift in redness on the blueness of the pixels (the first surface bandwidth effect above); the latter effect counterbalances the former if one pools all slices of constant blueness values (cf. Fig. 4.10).

4.3.2 Chromatic scene statistics

The above analysis of illumination-dependent regularities in the chromatic distributions of natural scenes focuses on the translations that the colour codes of

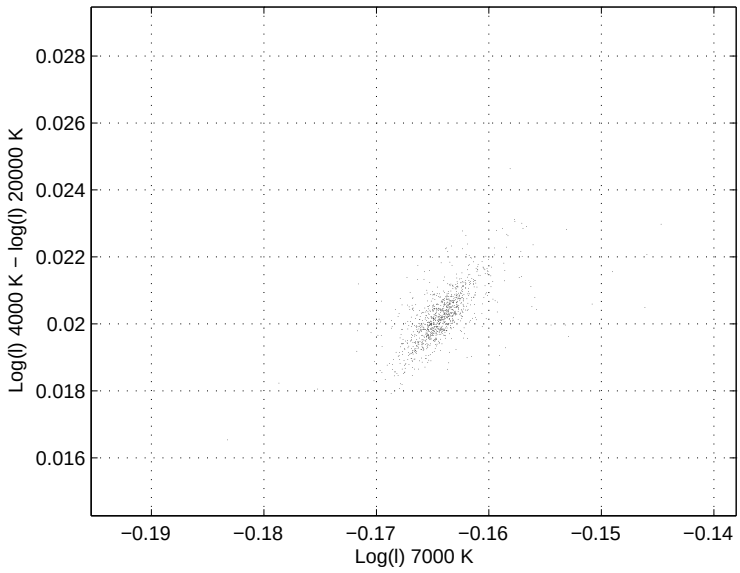

Figure 4.9: Change in redness ($\log(l)$) for a subset of pixels of scene Park4 when the 4000 K illuminant replaces the 20 000 K illuminant, plotted versus redness of the pixels (as measured by $\log(l)$ under 7000 K). The subset of pixels that are plotted comprises only the pixels of medium blueness, i.e. their $\log(s)$ value under 7000 K lies between -0.20 and -0.18. For this slice as well as for any other slice of pixels with similar blueness, the shift in $\log(l)$ under changing illumination increases with redness of the pixels.

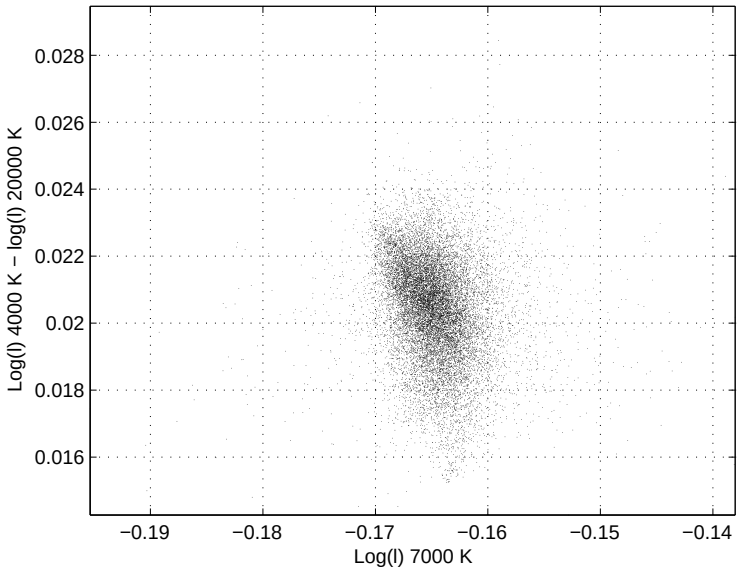

Figure 4.10: Change in redness ($\log(l)$) for all pixels of scene Park4 when the 4000 K illuminant replaces the 20 000 K illuminant, plotted versus redness of the pixels (as measured by $\log(l)$ under 7000 K). Please note that the cluster of the subset in Fig. 4.9 is part of the cluster for all pixels in this figure. The shift in $\log(l)$ under changing illumination is independent of redness of the pixels if one considers the whole set of pixels.

individual pixels undergo when the illumination changes. This information is generally not available to the visual system (e.g. when a new scene has been entered or the direction of view has been shifted), so in order to estimate the illuminant or infer illumination-independent surface colours, the visual system must rely on other measures. Here I assume that for these accomplishments the visual system takes into account (among other things) chromatic statistics of the whole scene as it is currently viewed, without necessarily analysing proximal chromaticity shifts by monitoring chromatic properties of individual surfaces over time. Thus, of special interest for a computational analysis of colour constancy is the level of chromatic scene statistics which is considered now.

Normalization-compatible chromatic translation

Not surprisingly, the space averaged chromaticities of each of the twelve scenes in the simulation change very systematically with changes of illumination (cf. Plate III). The mean redness of all pixels within an image of a scene under a particular illuminant is highest for the 4000 K daylight and decreases for the illuminants with higher colour temperature, whereas the mean blueness is lowest for the 4000 K daylight and then increases. These changes are very similar for all of the twelve scenes and resemble the chromaticity changes of the illuminants themselves, such that the four clusters of these scenes under the different illuminants line up (almost equally spaced in ($log(l)$,$log(s)$) plane) parallel to the daylight loci that are indicated by the four asterisks in Plate III.

However, the centre of each cluster does not coincide with the corresponding daylight locus. Instead, within each cluster all scenes have lower $log(l)$ as well as lower $log(s)$ values than the corresponding illuminant. This means that the reflectances are not spectrally neutral when averaged over all pixels within a scene but are biased toward greenish-yellowish surfaces. As to be expected for this sample of foliage-dominated natural scenes, the world does not seem to be grey. But this by itself does not imply that the Grey World assumption (cf. section 2.3.1) is inappropriate, since that assumption really only requires that the space-averaged reflectances of all possible scenes are rather similar, but not necessarily that they are chromatically neutral. If the averages were all similar and close to a greenish-yellowish standard surface, the Grey World assumption would be an appropriate simplification of the chromatic world, notwithstanding its misleading name.

But as Plate III shows, the Grey World assumption is inappropriate for a world similar to the natural scenes of Ruderman et al. (1998), since the space averaged chromaticities of the scenes under the same illumination do differ substantially.

Even with this small sample of very similar scenes, the chromatic difference between two scenes under the same illuminant can be as large as the difference between one scene under one illuminant compared to the same scene under the illuminant with the next higher (or lower) colour temperature. This is surprising because in the simulation the extreme range of unimpeded daylight illuminants was sampled by only four rather different illuminants (with colour temperatures of 4000 K, 5500 K, 8500 K, and 20 000 K). But even in these conditions, a substantial amount of overlap between the distributions of mean scene chromaticities for different illuminants exists.

This indicates that the ambiguity of the mean chromaticity is not only a hypothetical problem (cf. section 2.3.1). The image of a scene under a particular illuminant changes systematically with changes of illumination, but simply to invert this regularity does not suffice to solve the problem of colour constancy. The scenes themselves differ substantially in average chromaticity and therefore, the visual system cannot simply use the mean chromaticity of the proximal image as the only cue to the illuminant colour. Thus, in natural environments the visual system must rely on additional cues, candidates of which are considered next.

When the light gets red, the reds get lighter

While the Gaussian World analysis (section 3.5.2) as well as the results of the real world simulation for individual pixels (section 4.3.1) both suggest that the luminance-redness correlation of a scene should increase when the illumination becomes more reddish, Fig. 4.11 shows that, for a given scene, the correlations are almost independent of illumination. Thus, the luminance-redness correlation alone is not an effective cue in determining the illuminant colour.

Nevertheless this measure can resolve ambiguity, but rather like in the case of constellation (c) of Plate I than in constellation (b): when used together with the mean image redness it separates the clusters of images from different illuminations (the four diagonally oriented clusters in Fig. 4.11) and thus permits the estimation of illumination colour. For example, an image with a mean $log(l)$-value of -0.155 could be a reddish scene under neutral light or a neutral (or even greenish) scene under reddish light (Fig. 4.11). But if the visual system takes into account the correlation between redness and luminance, it can distinguish between illumination redness and scene redness. The more positive the correlation, the greater the evidence for redness of the illumination as opposed to redness of the scene. In the above example, image A is with high probability a greenish scene under the reddish 4000 K illuminant whereas image B is likely to be a reddish scene that belongs to the cluster of the more neutral 5500 K illumination.

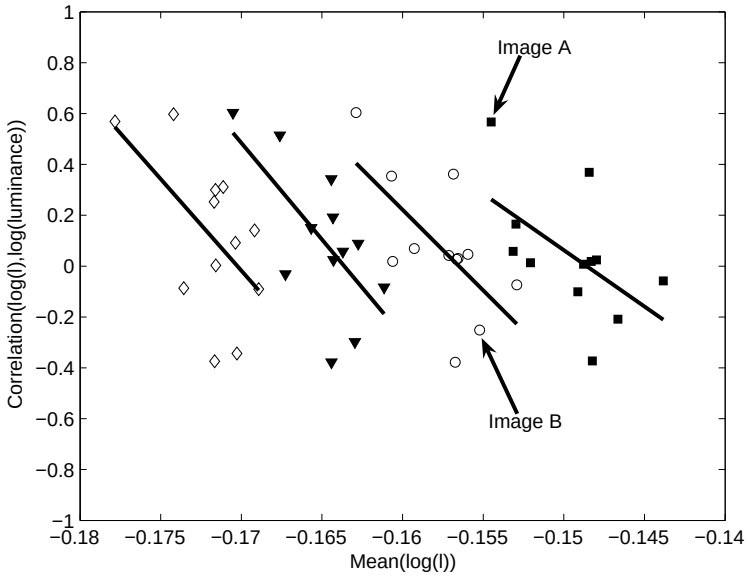

Figure 4.11: Correlation between pixel redness and pixel luminance within an image of a natural scene under a particular illuminant, plotted versus space average image redness. The four clusters show data for twelve scenes under four different illuminants (■: colour temperature 4000 K, ○: 5500 K, ▼: 8500 K, ◇: 20 000 K).

The reason why the luminance-redness correlation can help in this way is that scene redness and illuminant redness affect the correlation differently. Because of the eye's spectral sensitivity, middle wavelengths contribute more to luminance than do short or long wavelengths. Under neutral lighting, it is the neutral or moderately reddish and greenish pixels that reflect most middle-wavelength light and which therefore have the highest luminance. In neutral or greenish scenes surfaces that reflect predominantly long wavelengths and thus have lower luminance are rare. But in reddish scenes under neutral lighting the eye's diminishing sensitivity at long wavelengths will discriminate against the most reddish pixels of a reddish scene, and these will therefore tend to be of much lower luminance. The luminance-redness correlation thus becomes more negative the more reddish the scene is, as the sloped regression lines in Fig. 4.11 show. (The correlations between correlation(log(l),log(luminance)) and mean(log(l)) for the four illuminant clusters are -0.555 (4000 K), -0.654 (5500 K), -0.653 (8500 K), and -0.556 (20 000 K), in each case significantly different from zero ($P<0.05$).)

No comparable effect occurs if it is not the scene, but the light that is reddish. In this case, the low luminosity of reds is counterbalanced by the illuminant's greater energy at long wavelengths. Reddish surfaces reflect a larger proportion of long-wavelength than of short-wavelength light. Thus when the illuminant

becomes reddish they will be more luminous relative to other colours within the scene due to the better overlap of their spectral reflectances with the spectral range in which the illuminant has its highest power.

Thus reddish scenes, but not reddish illuminants, generate images with negative luminance-redness correlations. By evaluating both mean and correlation, two independent quantities, an observer can estimate two unknowns - the predominant colour (in this case, the degree of redness) inherent in the objects making up the scene, and the redness of the light source that illuminates the scene. In this way the luminance-redness correlation can resolve the ambiguity encountered in considering mean chromaticity alone and thus improve the estimation of the illuminant, even though this statistic is almost unaffected by changes in illumination.

The Gaussian World analysis predicts, as another deviation from normalization compatible mapping that the bluish surfaces (with high s values) increase in luminance more than e.g. greyish or yellowish surfaces when the illumination becomes bluish. This regularity has indeed been found to hold for all scenes in the analysis of the natural images on the level of individual pixels (page 41). Thus, it is surprising that the correlation between luminance and blueness within a scene is almost independent of illumination in the real world simulation (Fig. 4.12). Unlike the correlation between luminance and redness, the correlation between luminance and blueness varies widely and unsystematically within each illumination cluster such that the resulting constellation does not seem to allow a useful disambiguating role for this scene statistic in a chromatic world that resembles the used sample of natural scenes.

Shift resistance

The treatment of shift resistance regularities on the level of individual pixels (cf. section 4.3.1) has brought to light three surface bandwidth effects, of which only the first yields a candidate illumination cue on the level of scene statistics that might be worth to be taken into account by the visual system. But even this is questionable, as will be explained below.

The first surface bandwidth effect consists in greater illumination-induced shifts in redness for surfaces with high blueness (cf. page 42). In going from the bluish 20 000 K illuminant to the reddish 4000 K illuminant, the greyish or bluish pixels with higher $\log(s)$ values increase more in $\log(l)$ than the yellowish pixels with lower $\log(s)$ values (which are more shift resistant due to their smaller bandwidth). Thus, the clusters of pixels in the ($\log(l)$,$\log(s)$) plane that are negatively sloped for all scenes under 20 000 K rotate clockwise and become

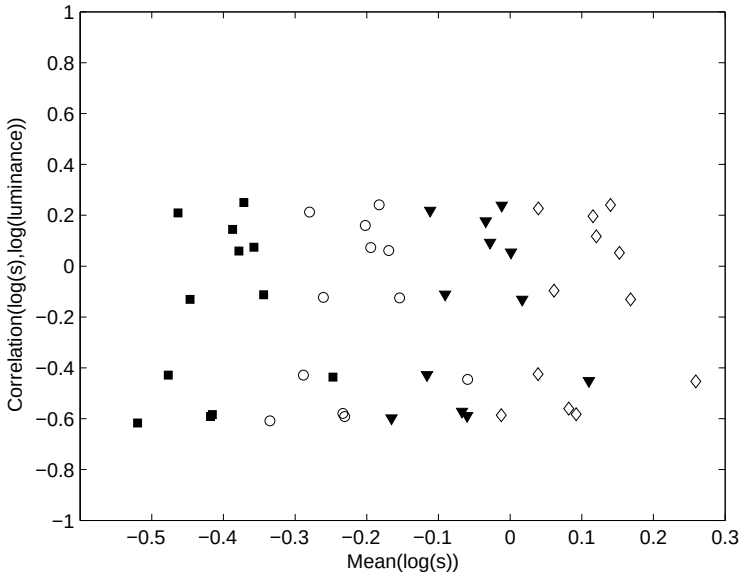

Figure 4.12: Correlation between pixel blueness and pixel luminance within an image of a natural scene under a particular illuminant, plotted versus space average image blueness. The four clusters show data for twelve scenes under four different illuminants (■: colour temperature 4000 K, ○: 5500 K, ▼: 8500 K, ◇: 20 000 K).

upright or even positively sloped under 4000 K (as can be seen in Fig. C.1 in the Appendix C.1).

The result of this first bandwidth effect is therefore that, for a given scene, the correlation between redness ($\log(l)$) and blueness ($\log(s)$) increases when the illumination changes from 20 000 K toward 4000 K (Fig. 4.13). This illumination dependence of the redness-blueness correlation could make this scene statistic a helpful cue for estimating the illuminant. But as Fig. 4.13 also shows, within each illuminant cluster this statistic depends on mean scene redness, too; the higher the mean $\log(l)$ value, the higher the redness-blueness correlation, as the positively sloped regression lines indicate. This illumination-independent relation between correlation and mean redness lets this statistic appear less promising as a candidate cue for the visual system to exploit. As explained in section 2.3.2, in a constellation where two cues are highly correlated both within and across illumination clusters, the combination of the two cues is only as helpful as each cue is alone (cf. constellation (d) in Plate I). The situation in Fig. 4.13 is, of course, more heterogeneous than the schematic constellation (d), e.g. it is unclear, whether the slight vertical offset between the regression lines of the 20 000 K cluster and the 8500 K cluster separates these two clusters enough to make the redness-blueness correlation useful at least for images with neutral to

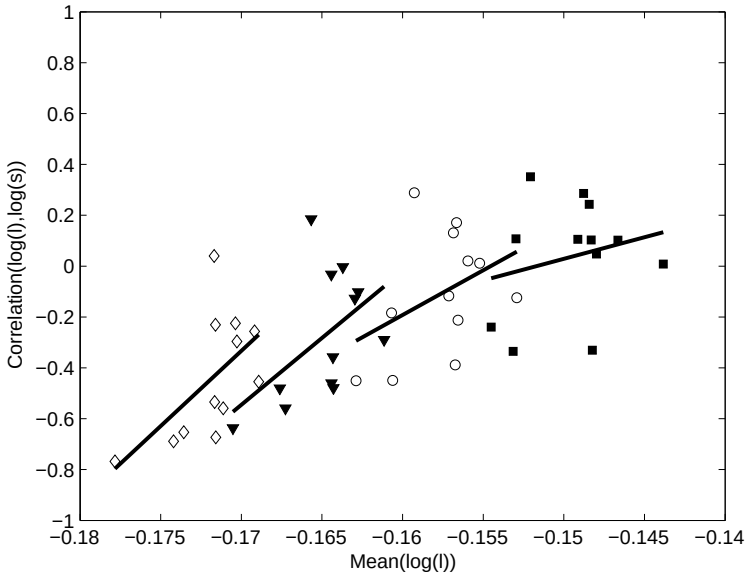

Figure 4.13: Correlation between pixel redness and pixel blueness within an image of a natural scene under a particular illuminant, plotted versus space average image redness. The four clusters show data for twelve scenes under four different illuminants (■: colour temperature 4000 K, ○: 5500 K, ▼: 8500 K, ◇: 20 000 K).

bluish space averages.

In summary, on the basis of the real world simulation no conclusions can be drawn whether the redness-blueness correlation is diagnostic for the illuminant in the scenes of Ruderman et al. (1998). It remains an open question, whether this cue is helpful for estimating the illuminant in the environment to which our visual system has adapted. Nevertheless, it seems worth investigating whether the visual system takes this cue into account (which will be done in section 5.7).

The second surface bandwidth effect that has been found on the level of individual pixels (cf. page 43), viz. the tendency for the lighter surfaces to be more stable in chromaticity under changing illumination, has only for a minority of the scenes a strength worth mentioning, such that it does not play any significant role on the level of scene statistics.

The third bandwidth effect is only evident if one separately considers within each scene subsets of pixels of similar blueness (cf. page 43). Within these slices with similar blueness values ($\log(s)$), the shift in redness ($\log(l)$) in going from 20 000 K illumination to 4000 K illumination is higher for pixels with high redness. Since this effect relates for each pixel the illumination-induced shift in one coordinate ($\log(l)$) to the starting value in the same coordinate (unlike the two bandwidth

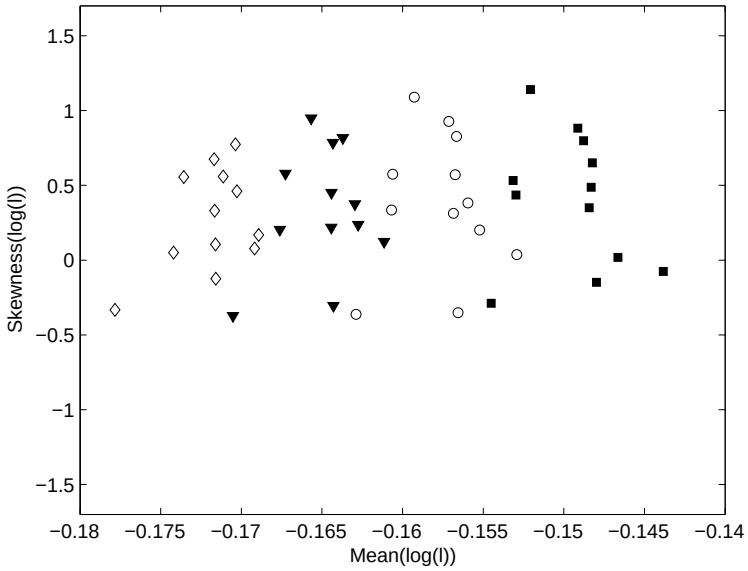

Figure 4.14: Skewness of pixel redness within a scene, plotted versus space average image redness for twelve natural scenes under four different illuminants (■: colour temperature 4000 K, ○: 5500 K, ▼: 8500 K, ◇: 20 000 K).

effects above, where the shift depends on the starting value in a different coordinate), it would on the level of scene statistics affect the skewness of $\log(l)$ — if this effect emerged clearly for each scene taken as a whole. But if one pools for each scene all slices of similar blueness values and considers the whole set of all pixels, the third bandwidth effect gets counterbalanced by the first bandwidth effect (cf. Fig. 4.10). Thus, the skewness of $\log(l)$ is independent of illumination, as Fig. 4.14 shows. As well is neither the skewness of $\log(s)$ (Fig. 4.15) nor the skewness of log(luminance) (Fig. 4.16) useful for estimating the illuminant.

Gamut compression

Gamut compression by narrow-band illumination is not clearly evident in the real world simulation. Variances among the scene elements in $\log(l)$, $\log(s)$, as well as log(luminance) do not differ systematically for the four illuminants (cf. Fig. 4.17, Fig. 4.18, Fig. 4.19 respectively). But the chosen daylight illuminants are not well suited to reveal such an effect. As Fig. 4.1 shows, they are roughly exponential in the visible range, and in the algebraic analysis of the Gaussian World (section 3.5) they would be assigned spectral curvatures close to zero, and therefore would not be expected to restrict the range of $\log(l)$ or $\log(s)$ much. Gamut compression is physically inevitable when the illumination is of sufficiently

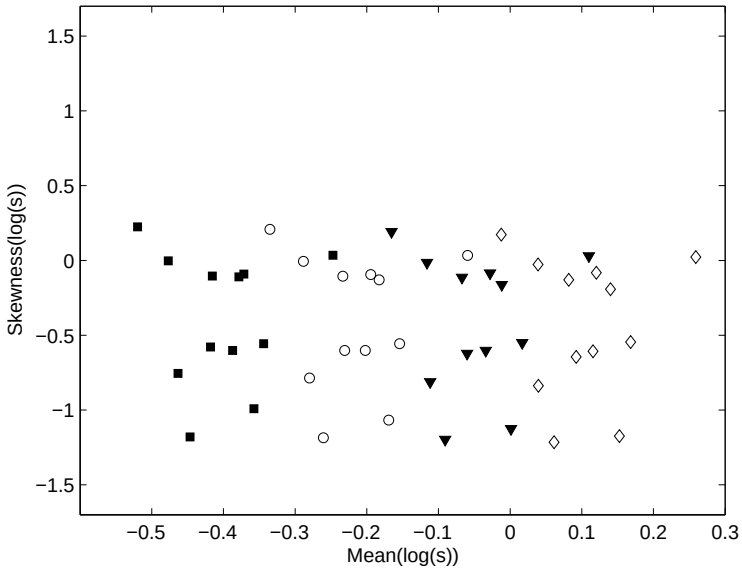

Figure 4.15: Skewness of pixel blueness within a scene, plotted versus space average image blueness for twelve natural scenes under four different illuminants (■: colour temperature 4000 K, ○: 5500 K, ▼: 8500 K, ◇: 20 000 K).

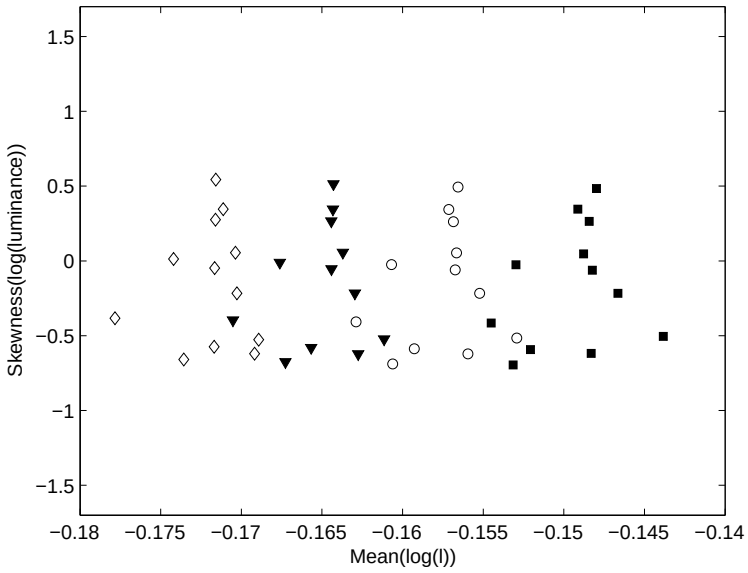

Figure 4.16: Skewness of pixel luminance within a scene, plotted versus space average image redness for twelve natural scenes under four different illuminants (■: colour temperature 4000 K, ○: 5500 K, ▼: 8500 K, ◇: 20 000 K).

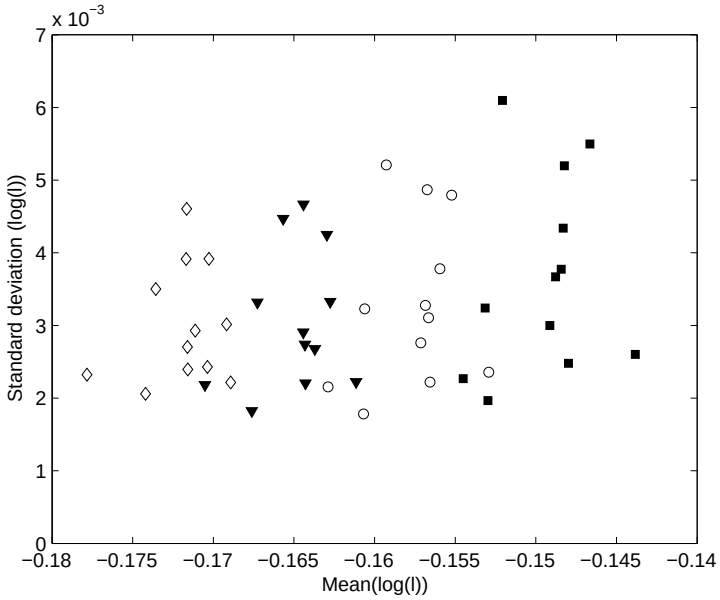

Figure 4.17: Standard deviation of pixel redness within a scene, plotted versus space average image redness for twelve natural scenes under four different illuminants (■: colour temperature 4000 K, ○: 5500 K, ▼: 8500 K, ◇: 20 000 K).

restricted bandwidth, but in the natural world it may become pronounced (if at all) only under restricted or indirect illumination (for instance in forest scenes).

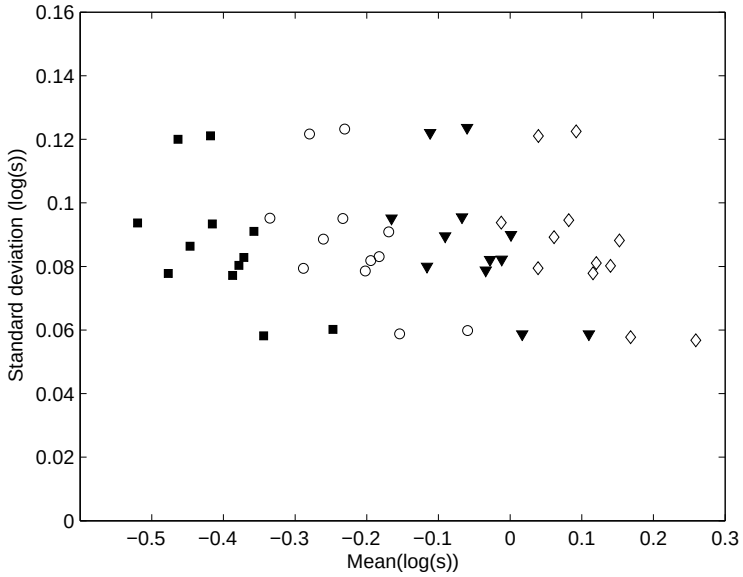

Figure 4.18: Standard deviation of pixel blueness within a scene, plotted versus space average image blueness for twelve natural scenes under four different illuminants (■: colour temperature 4000 K, ○: 5500 K, ▼: 8500 K, ◇: 20 000 K).

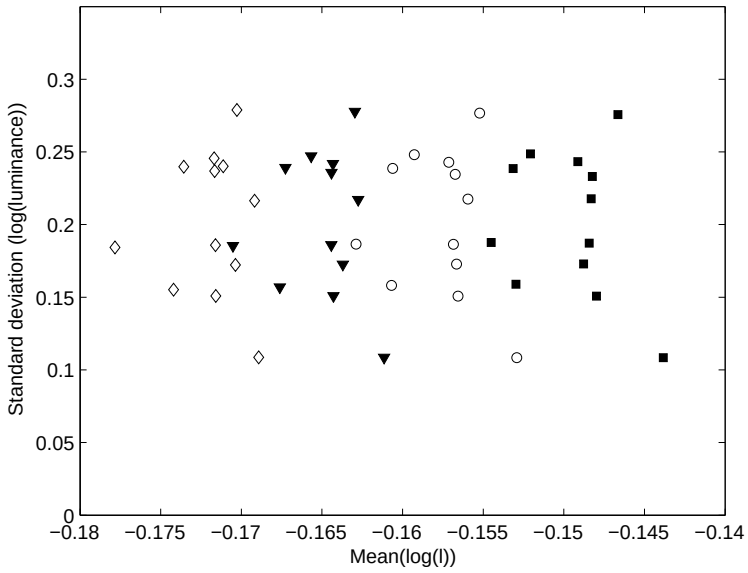

Figure 4.19: Standard deviation of pixel luminance within a scene, plotted versus space average image redness for twelve natural scenes under four different illuminants (■: colour temperature 4000 K, ○: 5500 K, ▼: 8500 K, ◇: 20 000 K).

Chapter 5

Experiments

5.1 Which chromatic cues are used by the visual system?

The theoretical analysis of the Gaussian World idealization (section 3.5) and the analysis of the real world simulation (section 4.3) have yielded independent but concordant evidence that several higher order chromatic scene statistics can (besides the mean chromaticity) be helpful for separating the retinal image into chromatic components for surfaces and illumination in natural scenes. Thus, the question arises whether the visual system exploits these proposed scene statistics as cues for the perception of surface and illumination colour. This question shall be addressed experimentally in the present chapter. The principal idea for this is to vary systematically, in a separate experiment for each candidate scene statistic, the value of that statistic in a variegated stimulus while keeping all other statistics of the chromatic distribution constant. If the manipulation of the scene statistic tested in a particular experiment leads to the expected systematic shifts in subjects' grey settings for a test field in these variegated stimuli, this will be taken as evidence that the human visual system takes into account this statistic for the perception of surface and illumination colour and that it thus indeed exploits this candidate cue.

The following chromatic scene statistics have been proposed as candidate cues by the Gaussian World idealization and the real world simulation and will be tested experimentally:

1. *Luminance-redness correlation.* Several theoretical findings propose that the correlation between luminance and redness in the retinal image is helpful

for estimating the illuminant colour. Firstly, the algebraic solution for cone excitations in the Gaussian World predicts that surfaces similar to the illuminant in colour are rendered as lighter than surfaces dissimilar from the illuminant in colour (cf. page 29). Thus, when the lighting gets red, the reddish surfaces (or more precisely, surfaces with higher l values) should get lighter, relative to e.g. greyish or greenish surfaces (surfaces with lower l values). Secondly, this illumination-dependent regularity has in deed been found to hold for the pixels of natural scenes in the real world simulation (cf. page 41). Thirdly, the analysis of the natural scenes on the level of chromatic scene statistics has shown that this illumination-dependent regularity makes the luminance-redness correlation a useful cue for the illumination, even though this illumination-dependent regularity interacts with a second regularity, such that the correlation value of a given scene is almost unaffected by changes in illumination. Nevertheless, in the real world simulation this statistic is diagnostic for the illuminant, in that, for instance, a high luminance-redness correlation indicates a reddish illumination (cf. page 49).

In accordance with these theoretical findings, I have found in cooperation with Donald I. A. MacLeod (Golz & MacLeod, 2002) that the luminance-redness correlation is exploited as a cue for the illumination by the human visual system. *Experiment 1* presented here is a slightly modified replication of that work. In *experiment 2*, I will survey the bearing of this basic effect and investigate whether it holds also for stimuli with non-neutral mean chromaticities. In *experiment 3*, I will address a possible objection to experiment 1 and 2. In these experiments, the space-averaged redness and blueness (mean(l) and mean(s)) is kept constant for all conditions of luminance-redness correlation. But a different possible measure of space-averaged redness (the luminance-weighted mean of l) then covaries with the luminance-redness correlation. One can conclude based on estimations for the shifts of grey settings induced by changes in luminance-weighted mean chromaticity that the effect of the luminance-redness correlation is not an artifact of differences in the luminance-weighted mean (MacLeod & Golz, 2003), but here I will address and refute this objection more directly in experiment 3 by keeping the luminance-weighted mean constant for all conditions of luminance-redness correlation.

2. *Luminance-blueness correlation.* The theoretical evidence that the correlation between luminance and blueness in the retinal image is helpful for estimating the illuminant colour is more ambiguous than in the case of luminance-redness correlation. Similar to the illumination-dependent regularity underlying the latter scene statistic, the Gaussian World analysis predicts that the bluish surfaces (with high s values) increase in luminance more than e.g. greyish or yellowish surfaces when the illumination becomes

bluish (cf. page 29). This regularity has indeed been found to hold for all scenes in the analysis of the real world simulation on the level of individual pixels (cf. page 41). But surprisingly, in the analysis of the real world simulation on the level of scene statistics the correlation between luminance and blueness within a scene is almost independent of illumination (cf. page 51). In a chromatic world that resembles the used sample of natural scenes, the luminance-blueness correlation thus seems to be not diagnostic for the illumination, though, of course it cannot be ruled out that the used scenes are not representative in this respect for the chromatic environment to which the human visual system has adapted. Based on this mixed theoretical evidence, I nevertheless include the luminance-blueness correlation in the set of candidate cues that are tested experimentally.

Thus, *experiment 4* tests whether the luminance-blueness correlation is taken into account by the human visual system for the perception of surface and illumination colour. The design of this experiment is comparable to the investigation of the luminance-redness correlation in experiment 1, so the non-luminance-weighted mean chromaticity is kept constant for the different correlation conditions and again, one could be worried by the fact that the luminance-weighted mean chromaticity co-varies with the manipulated correlation. Thus, in *experiment 5* now the luminance-weighted mean blueness is kept constant for all conditions of luminance-blueness correlation.

3. *Redness-blueness correlation.* While the potential usefulness for the two correlation statistics above result from a chromaticity-dependent increase in intensity for surfaces similar to the illuminant due to their better spectral overlap with the illumination, the third tested scene statistic, viz. the correlation between redness and blueness in the retinal image, rests on a illumination-dependent regularity that is mediated by differences in surface bandwidth. As the Gaussian World idealization proposes (cf. page 32), the light reflected from surfaces with sufficiently narrow-band reflectance will undergo a smaller chromaticity shift under changes of illumination than the light from broad-band surfaces. Accordingly, the illumination-induced shift in the real world simulation is greater for bluish and greyish surfaces than for yellowish surfaces which are more shift resistant due to their smaller bandwidth (cf. page 42). This leads, in the analysis of the real world simulation on the level of scene statistics, to an increase in the redness-blueness correlation when the illumination changes toward the more reddish 4000 K illuminant. Unfortunately, this statistic depends also systematically on the (illumination-independent) scene-averaged surface redness, in a way that weakens the usefulness of this statistic as a cue to the illumination: the resulting constellation of this statistic to the mean image redness (Fig. 4.13) resembles to some degree the schematic constellation (d) in Plate I, where two cues are correlated both within and across illumination clusters such

that the combination of the two cues is only as useful as each cue is alone (cf. section 2.3.2). On the basis of the real world simulation it thus remains an open question whether the redness-blueness correlation is helpful for estimating the illuminant in the natural environment.

Therefore, in *experiment 6* I investigate whether the human visual system takes the redness-blueness correlation within the retinal image into account for the perception of surface and illumination colour. In *experiment 7*, this question is addressed for stimuli with non-neutral mean chromaticities. Different from the situations for the two chromaticity-luminance correlations above, in the case of the redness-blueness correlation, both the luminance-weighted as well as the non-luminance-weighted mean chromaticity can be kept constant simultaneously (because the luminance-weighted mean is not affected by varying a correlation between two chromaticity measures). Thus, for this scene statistic, no additional experiment to controll for the luminance-weighted mean chromaticity is needed.

5.2 Experiment 1

5.2.1 Luminance-redness correlation: basic experiment

The aim of this experiment is to investigate whether the human visual system takes into account for the perception of surface and illumination colour the correlation between redness (l value) and luminance within the retinal image. More precisely, the hypothesis is — based on the results of the Gaussian World analysis and the real world simulation — that the higher the luminance-redness correlation in the stimulus is the more reddish becomes the illumination inferred by the visual system.Consequently, the visual system would have to take this estimated illumination into account when determining perceived surface colours: if the estimated illumination becomes more reddish, a retinal area of constant chromaticity would look like a less reddish surface. For the task of experiment 1 the following prediction results: for different conditions of luminance-redness correlation in a variegated surround, subjects have to adjust a central test field to look grey. For higher correlation values, the visual system infers a more reddish illumination and in order to keep the appearance of the test field constantly grey subjects adjust the test field to more reddish chromaticities. (A test field with a chromaticity that looks grey in an a condition with low luminance-redness correlation would look greenish in a condition with high luminance-redness correlation because, loosely speaking, more of the constant retinal input for the test field is attributed to the more reddish illumination and thus less redness can be attributed to the surface. In order to make the test field appear grey they thus would have to pick

a more reddish chromaticity.) Thus, the precise hypothesis is that the higher the luminance-redness correlation is in the surround, the higher is the l-value of subjects' grey settings for the test field.

In order to test this hypothesis, it is of course necessary to isolate the assumed effect of the luminance-redness correlation from all other known or potential influences of other statistics of the chromatic distribution in the surround on the perceived colour on the test field. (For instance, it is necessary to keep the mean chromaticities constant while varying the luminance-redness correlation.) To this end, I use a type of stimuli that allows to independently vary single statistics of the chromatic distribution of patches in the stimulus (means and standard deviations of l, s and luminance as well as correlations between these three chromatic measures). The procedure I use to create this type of stimuli differs only slightly from that of Mausfeld and Andres who used this type of stimuli first (cf. Andres, 1997; Mausfeld, 1998; Mausfeld & Andres, 2002). The specific stimuli for experiment 1 will be described in the following Section.

While in chapter 4 the chromatic scene statistics have been considered on a logarithmic scale for l, s and luminance (because in this logarithmic color space the regularities of illumination induced shifts on the level of individual surfaces can be suitably analysed), I use for the experiments the more common linear measures l, s and luminance. This is unproblematic since the analysis of the real world on the level of scene statistics yields the same results when carried out with the linear measures instead of the logarithmic ones. I have also additionally carried out several of the following experiments using the logarithmic measures instead of the linear ones and I have in no case found results substantially different from the ones presented below.

5.2.2 Methods

Stimuli

The stimuli consist of overlapping circles of various colours with the test field in the center of the display on top of all circles in that area (cf. Plate IV for examples of these stimuli). The colours of the circles are chosen such that the chromatic distribution of the pixels within the surround takes on for the means, standard deviations and correlations in (l,s,$luminance$) space the values given in Tab. 5.1. As the independent variable of experiment 1, the redness-luminance correlation is varied in five conditions with the values -1.0, -0.7, 0.0, +0.7, and +1.0. (If, in this as well as in the following experiments, the value of -1.0 or +1.0 was not realizable for mathematical reasons, the closest realizable value, which

usually deviated less than 0.01 from the target value, was used. Such cases are very rare and I will denote them anyway by -1.0 or +1.0.) The values of all other statistics are chosen to resemble the values of natural scenes, i.e. for each statistic the value of that statistic averaged over all twelve scenes of Ruderman et al. (1998) under 7000 K daylight illumination is used. Only the mean and the standard deviation of luminance have been determined differently: they have been set to values lower than the natural ones such that all surround colours in the stimuli lie within the gamut of the monitor. Another specification of the chromatic distribution of the surround colours which is equivalent to the above description can be given in terms of eigenvectors: Tab. C.5 in Appendix C.3 lists the eigenvectors and the corresponding standard deviations of the chromatic distribution in $(l, s, luminance)$-space for all five conditions of luminance-redness correlation.

As the examples in the top row of Plate IV show, in the stimuli with luminance-redness correlation of -1.0, the reddish circles are darker than the greenish ones while the opposite holds for the stimuli of the condition with luminance-redness correlation of +1.0. All other statistics besides the luminance-redness correlation are kept constant for all five conditions, not only on average over the entire area of the stimuli but also within each of five concentric regions around the central test field. These disjunct, concentric, and annular-shaped regions have increasing diameters such that they consecutively cover the whole stimulus. The innermost region adjacent to the test field has an outer diameter of 9° in the field of vision. The test field itself has the same diameter as all other circles in the stimuli, viz. 1.5° in the field of vision. The whole stimulus display measures 70° × 55° in the field of vision.

Subjects

Eleven subjects (4 male, 7 female, aged from 20 to 33 years) participated in experiment 1. All subjects had normal or corrected-to-normal vision and showed no colour vision deficiency, as tested by the Ishihara colour plates. All subjects except for one (JG, the author) were naïve about the design and purpose of the experiment.

Procedure

The stimuli were presented on a carefully calibrated CRT monitor (Sony GDM-F500R), driven by a PC with a standard graphics board with 24 bit (8 bit per channel) colour resolution. The spatial resolution was set to 1280 × 960 pixel

Statistic	*Value*
Means	
l	0.6877
s	1.1466
luminance	20.0 cd/m^2
Standard deviations	
l	0.005
s	0.1536
luminance	5.0
Correlations	
l and luminance	5 conditions:
	-1.0; -0.7; 0.0; 0.7; 1.0
s and luminance	-0.1153
l and s	-0.2133

Table 5.1: Chromatic statistics of stimuli for experiment 1. The listed values are the means, standard deviations, and correlations of the chromatic distribution of the pixels within the sourround of the stimuli.

(horizontal × vertical) at a refresh rate of 85 Hz, non-interlaced. Subjects viewed the 21-inch screen (53 cm diagonal) of the monitor from a distance of 55 cm in a dark room through a tunnel covered inside with matte black felt such that the only visible object was the stimulus displayed on the screen.

Subjects had to adjust the test field that was located in the center of the stimulus on top of all other circles in that area to look grey. It was explained to them that the test field should appear perfectly achromatic without any reddish, greenish, yellowish, bluish or otherwise coloured tinge. To this end, subjects could change the chromaticity of the test field in the MacLeod-Boynton (l,s) chromaticity plane by means of the four arrow keys of the keyboard, while the luminance of the test field was fixed at 20 cd/m^2 which is equal to the mean luminance of each respective stimulus. Subjects made themselves familiar with the directions of colours in the adjustment plane as well as with all other details of the experimental procedure in test trials before the actual experiment.

Note that observers were instructed to judge the colour of the test field, but no further interpretation was indicated whether this was related to the light reaching their eyes from the test field or rather to the underlying surface property of the test patch. Instructions that provide different interpretations in this respect can, under some conditions, affect subjects' judgment of colour appearance (Arend & Reeves, 1986), but the range of conditions where this instruction factor makes a noticeable difference is not yet known. My own experience as well as interviews

with subjects after they performed in the experiment suggests that for the task and stimuli used here the results would not differ if instructions were provided that emphasized the properties of the surface material the test field was made of. Subjects were further instructed that they should not adjust the test field to make it look like one of the patches in the surround that they assume to be grey. Instead, they should adjust the test field until that field appears grey to them.

I have found in experiments using stimuli and procedures very similar to the one presented here that the viewing instructions given to subjects can substantially effect the results of grey settings (Golz, 2004). Particularly, if subjects fixate the test field during the entire presentation of the stimuli and do not visually explore the stimuli by looking around in the surround, such a viewing behaviour leads to smaller surround induced shifts of the grey settings (lower degrees colour constancy), probably due to differences in foveal state of adaptation as well as due to reduced spatiotemporal modulation in the retinal periphery (cf. Appendix B).

In order to allow the viewing behaviour of the subjects to be as natural as possible, and to avoid any visual mechanism that is sensitive to spatiotemporal modulation in the periphery and potentially relevant for the perception of surface and illumination colours to be prevented from proper operation by artificial viewing behaviour, I have adopted for the experiments presented here the following viewing instructions that were provided and explained to the subjects in the test trials before the actual experiment:

> "... Adjust the central test field until it looks grey to you. Then explore the scene by looking around in the surround for about ten seconds. To this end, move your gaze to different parts of the scene along a circle around the center halfway to the border of the display. Then look at the test field again to make sure it still looks perfectly grey. If it doesnt, adjust it until it does. Keep alternately looking around and then adjusting the test field until a beep informs you that you can end the trial by pressing the [ENTER] key. ..."

Subjects could only accept their settings and end the trial after 60 seconds, which was indicated by the beep. This rather long minimum presentation time for each stimulus was chosen to allow the visual system to arrive at a steady state for any mechanism affecting the perception of the test field colour before the subject's final adjustment for the respective trial was recorded. In order to balance any remaining carry-over effect from one experimental condition to the next, the succession of the five conditions of correlation in the surround were presented twice, first ascending from correlation condition -1.0 to +1.0 and then backwards in

descending order from correlation condition +1.0 to -1.0. Within this succession of conditions, each condition was presented four times. In the four trials of a particular condition, slightly different versions of stimuli for that condition were used in which merely the spatial layout of the circles in the surround varied, but not the value of the correlation or of any other chromatic statistic. Thus, the experimental procedure resulted for each subject in eight settings per condition. The entire experiment with the total of 40 settings was performed in one session which usually lasted about 50 minutes.

5.2.3 Results

The results for all of the eleven subjects show a dependence of the grey settings on the luminance-redness correlation in the surround that is in accordance with the hypothesis stated above: the higher the luminance-redness correlation, the higher on average the l-value of subjects' grey settings (cf. Fig. 5.1).[1] As Tab. 5.2 shows, this effect is individually significant for all of the eleven subjects.[2] This table also provides a measure for the size of the effect, viz. $\hat{\psi}$. This measure ranges from 0.0011 (subject TP) to 0.0034 (subjects GW, JG, and JJ) with an average value of 0.00255. Here, it may suffice to say that for each subject the l-value of the grey settings increases, on average, by its respective $\hat{\psi}$ if the luminance-redness correlation increases by 1.0. (More precisely as explained in Appendix A, $\hat{\psi}$ is the estimated slope of the linear dependence of the l-value of the grey settings on the luminance-redness correlation.) Thus, for the average subject, the l-value of the grey setting increases by $2\times0.00255 = 0.0051$ when the luminance-redness correlation increases from -1.0 to +1.0. For comparison, the average standard deviation of l in the natural images of Ruderman et al. (1998) is 0.005 (this value has also been chosen for the standard deviation of l in the surround of the experimental stimuli) and the difference in l between the two most extreme unimpeded daylights (4000 K and 20 000 K) is 0.04, so approximately eight times the above-mentioned average subject shift of 0.0051.

[1] Note that throughout this whole chapter, in all figures of experimental results the ordinate will always cover the same range of l-values (or s-values for experiment 4 and 5) in order to make the figures easily comparable. The upper and lower limit of this range might differ, though, due to individual differences in the overall level of the chromaticity of grey settings.

[2] The test that is used for the statistical analysis of the experimental results of this as well as of all following experiments is described in Appendix A. Throughout this chapter, I will regard a subject's results as significant if $p < 0.05$ (i.e. $t_{emp} > t_{35;0.05}$). In Tab. 5.2 and all of the following tables of experimental results, additional information for readers who prefer a different significance level is indicated by: *: $p < 0.05$; **: $p < 0.01$; ***: $p < 0.001$; $^{n.s.}$: not significant.

Subject	$\hat{\psi}$	t_{emp}
AH	0.0028	8.52 ***
GW	0.0034	18.21 ***
HS	0.0031	7.47 ***
JG	0.0034	18.40 ***
JJ	0.0034	10.19 ***
LS	0.0025	12.78 ***
ML	0.0020	4.17 ***
MP	0.0026	8.85 ***
RI	0.0015	8.27 ***
SC	0.0023	8.91 ***
TP	0.0011	5.64 ***

Table 5.2: Results of experiment 1.

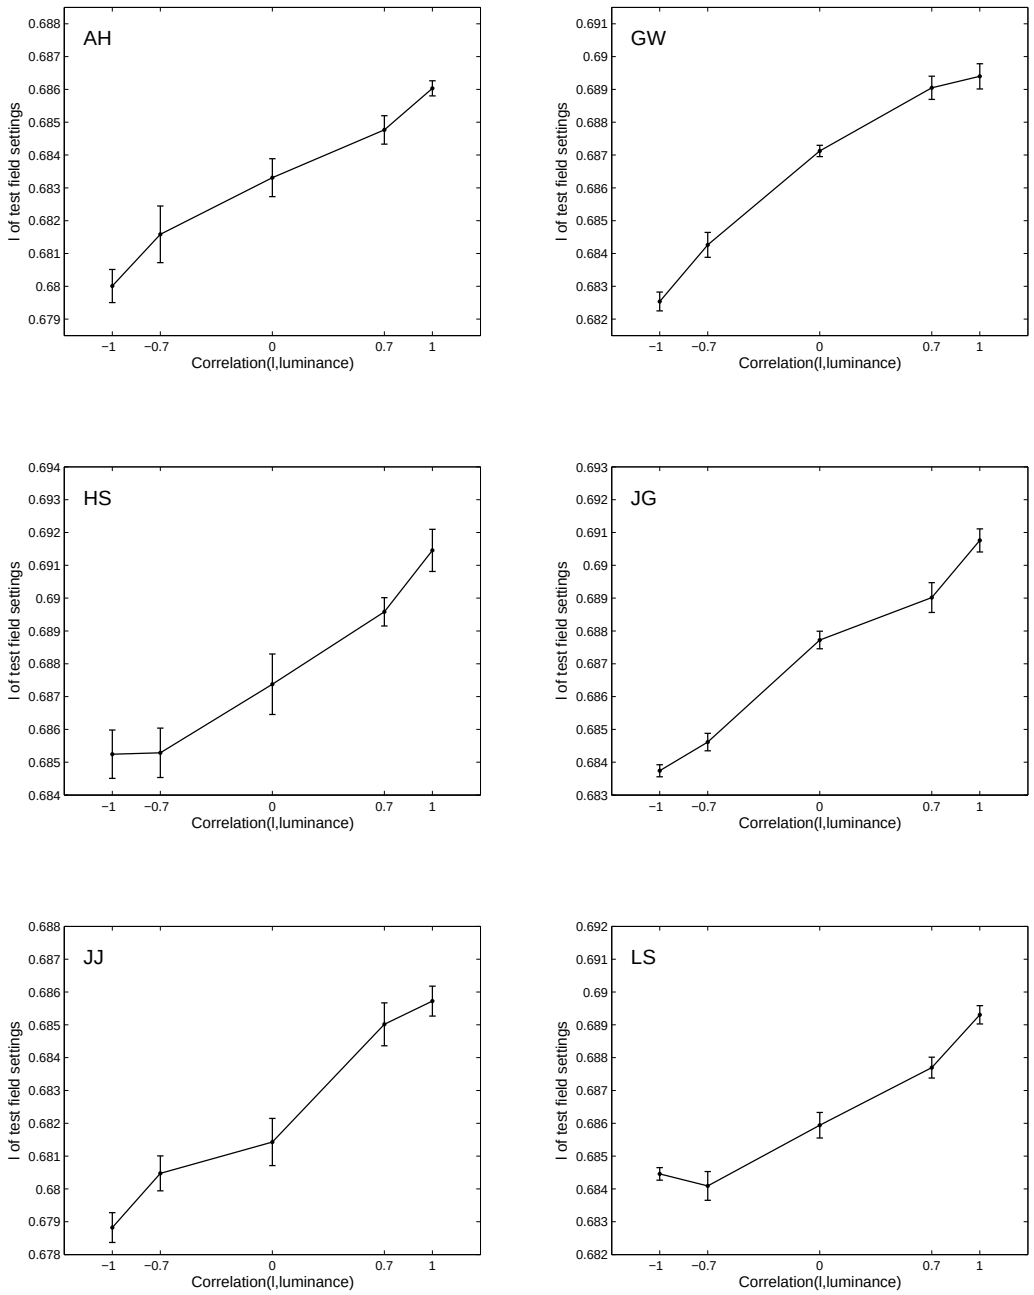

Figure 5.1: Results of experiment 1.

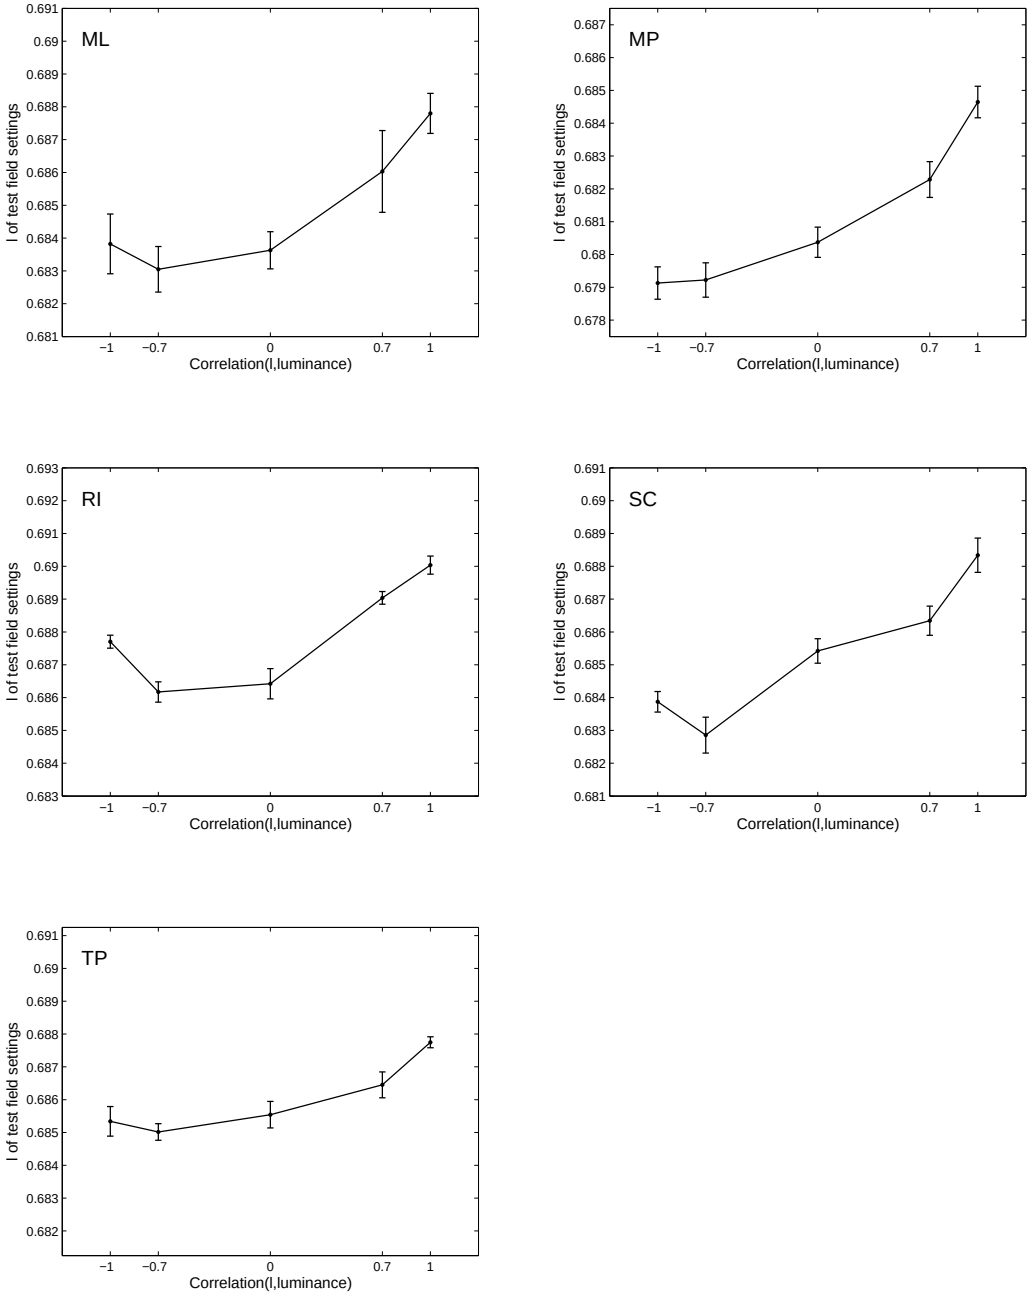

Figure 5.1: (Continuation)

5.2.4 Discussion

In summery, experiment 1 has yielded strong evidence that the visual system indeed takes into account the luminance-redness correlation for the perception of surface and illumination colour. This effect seems to be robust in that the results of all eleven subjects tested are significant. The size of this effect is substantial but not overwhelmingly large. This is not surprising, since the chromatic regularity in the natural world images that has suggested the luminance-redness correlation as a useful cue is itself not very large.

This experimental result nicely confirms the similar finding of earlier work that I carried out in cooperation with Donald I. A. MacLeod (Golz & MacLeod, 2002). Both experiments yield results that allow the same conclusion despite the fact that they have been carried out in different labs with slight differences in experimental setup, calibration procedure, experimental procedure, stimuli and with different subjects, so this is in my opinion convincing evidence that the human visual system exploits the luminance-redness correlation as a cue for the illumination.

But in experiment 1 as well as in Golz and MacLeod (2002), only stimuli with neutral (greyish) mean chromaticity have been used. In the next experiment, I will extend the investigation of the luminance-redness correlation to other backgrounds.

5.3 Experiment 2

5.3.1 Luminance-redness correlation: non-neutral stimuli

The aim of experiment 2 is to survey the bearing of the basic effect of the luminance-redness correlation found in experiment 1 and test whether it also holds for stimuli with non-neutral mean chromaticities. To this end, four different chromaticities are used for the mean chromaticity of the stimuli. I will label these four surround types “blue”, “green”, “red”, and “yellow” since this (very roughly) describes how these chromaticities appear when viewed in isolation. For each of these four surround types, a sub-experiment is carried out (in a separate session) that very much resembles experiment 1, except for the mean chromaticity of the stimuli. That is, for each of the four mean surround chromaticities, the luminance-redness correlation in the surround is varied in five conditions and subjects have to adjust a central test field such that it appears grey to them. The hypothesis is that, in each of the four sub-experiments the following relation

holds: the higher the luminance-redness correlation is in the surround, the higher is the l-value of subjects' grey settings for the test field. The rationale for this relation is exactly analogous to the one in experiment 1 (see page 61). Neither the Gaussian World analysis nor the real world simulation have yielded any evidence that the illumination-dependence of the luminance-redness correlation hinges on the mean chromaticity of the retinal image. Both analyses rather suggest that this correlation can figure as a signed cue to the illumination no matter what the mean chromaticity is. Thus, the above-mentioned relation should hold in each of the four sub-experiments with different mean surround chromaticity.

5.3.2 Methods

Stimuli

The stimuli used in experiment 2 are equal to those used in experiment 1 in every aspect except for the mean chromaticities of the surround, i.e. they share the following features (for details see page 62):

- They consist of overlapping circles of various colours with the test field in the center of the display on top of all circles in that area.
- They have the same size and spatial layout as the stimuli in experiment 1.
- The redness-luminance correlation is varied in five conditions with the values -1.0, -0.7, 0.0, +0.7, and +1.0.
- The values of all other statistics are kept constant and are chosen to resemble the values of natural scenes (cf. Tab. 5.1).

Only the mean chromaticities of the surround differ from the neutral one of experiment 1. The four mean chromaticities used (see Tab. 5.3) are determined as follows:

- The chromaticity of the blue surround equals the chromaticity of a 20 000 K daylight.
- The chromaticity of the yellow surround equals the chromaticity of a 4 000 K daylight.

Statistic	*Value*
blue surround	
mean l	0.670
mean s	1.861
green surround	
mean l	0.661
mean s	0.884
red surround	
mean l	0.722
mean s	1.212
yellow surround	
mean l	0.712
mean s	0.587

Table 5.3: Mean chromaticities of the four different types of stimuli for experiment 2.

- The chromaticities of the red and the green surround are chosen such that the connection line of these two chromaticities is orthogonal to the connection line of the blue and yellow chromaticities when plotted in the CIE (u',v') chromaticity diagram. The intersection of these two connection lines lies very close to the grey chromaticity that was used as mean surround chromaticity in experiment 1 and all four chromaticities are approximately equally distant from the intersection point.

The chromaticities of the red and the green surround are equal to two chromaticities that Delahunt and Brainard (2004) used in their study (their chromaticities 'Red_60' and 'Green_60') while the chromaticities of the blue and the yellow surround are very close to their 'Blue_60' and 'Yellow_60'.

Subjects

Four subjects (all female, aged from 20 to 25 years) participated in experiment 2. All subjects had normal or corrected-to-normal vision and showed no colour vision deficiency, as tested by the Ishihara colour plates. All subjects had previously participated in experiment 1 and were naïve about the design and purpose of the experiment.

Subject	$\hat{\psi}$	t_{emp}
AH	0.0028	6.39 ***
LS	0.0021	6.69 ***
RI	0.0013	6.33 ***
SC	0.0031	10.29 ***

Table 5.4: Results of experiment 2 for blue stimuli.

Subject	$\hat{\psi}$	t_{emp}
AH	0.0024	6.10 ***
LS	0.0020	7.09 ***
RI	0.0021	7.70 ***
SC	0.0033	11.55 ***

Table 5.5: Results of experiment 2 for green stimuli.

Procedure

The procedure equaled in every aspect (experimental setup, task, instructions, sequence of presentations etc.) exactly that of experiment 1 (see page 63), except that instead of one session for the entire experiment 1 four sessions were needed for experiment 2 — one session for each mean surround chromaticity. All subjects performed the four sessions in the order blue, green, red, and finally yellow mean surround chromaticity.

5.3.3 Results

The results for all of the four mean surround chromaticities and for all of the four subjects show the expected dependence of the grey settings on the luminance-redness correlation in the surround (as Fig. 5.2, 5.3, 5.4, and 5.5 show for the blue, green, red, and yellow background respectively): the higher the luminance-redness correlation, the higher on average the l-value of subjects' grey settings. All of these effects are individually significant as can be seen in Tab. 5.4, 5.5, 5.6, and 5.7.[3] The measure for the size of the effect $\hat{\psi}$ is, averaged over the four subjects, 0.002325, 0.00245, 0.0022, and 0.00255 for the blue, green, red, and yellow background respectively. These values are very similar to the average effect size in experiment 1.

[3]Note that due to the way the hypothesis has been stated above, no adjustment of the α-level for the four simultaneous tests per subject is required.

Subject	$\hat{\psi}$	t_{emp}
AH	0.0010	3.28 **
LS	0.0023	8.02 ***
RI	0.0021	8.68 ***
SC	0.0034	10.10 ***

Table 5.6: Results of experiment 2 for red stimuli.

Subject	$\hat{\psi}$	t_{emp}
AH	0.0012	4.23 ***
LS	0.0036	11.40 ***
RI	0.0018	7.09 ***
SC	0.0036	12.41 ***

Table 5.7: Results of experiment 2 for yellow stimuli.

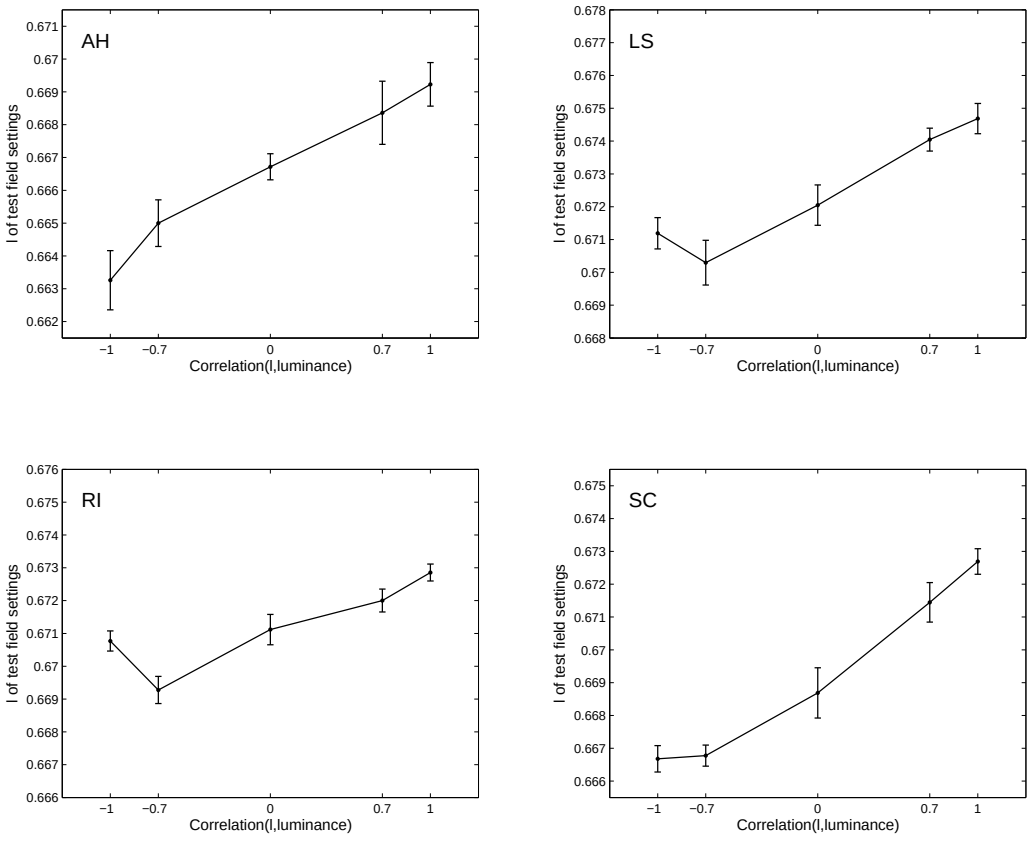

Figure 5.2: Results of experiment 2 for blue stimuli.

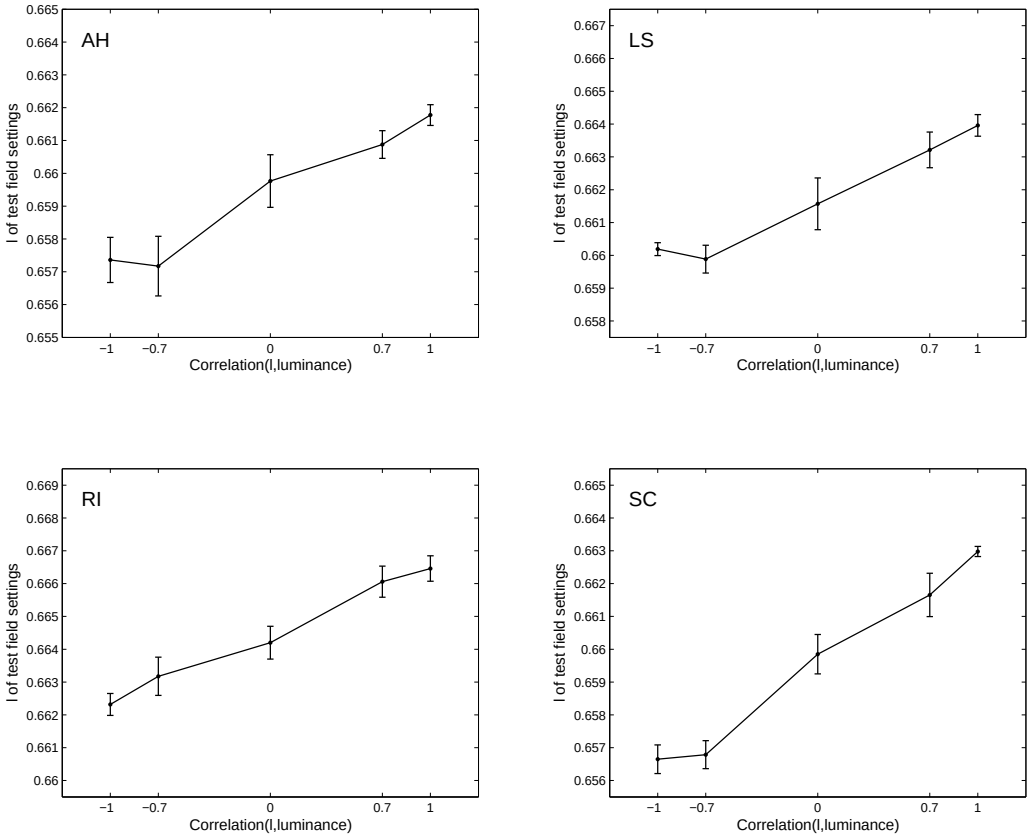

Figure 5.3: Results of experiment 2 for green stimuli.

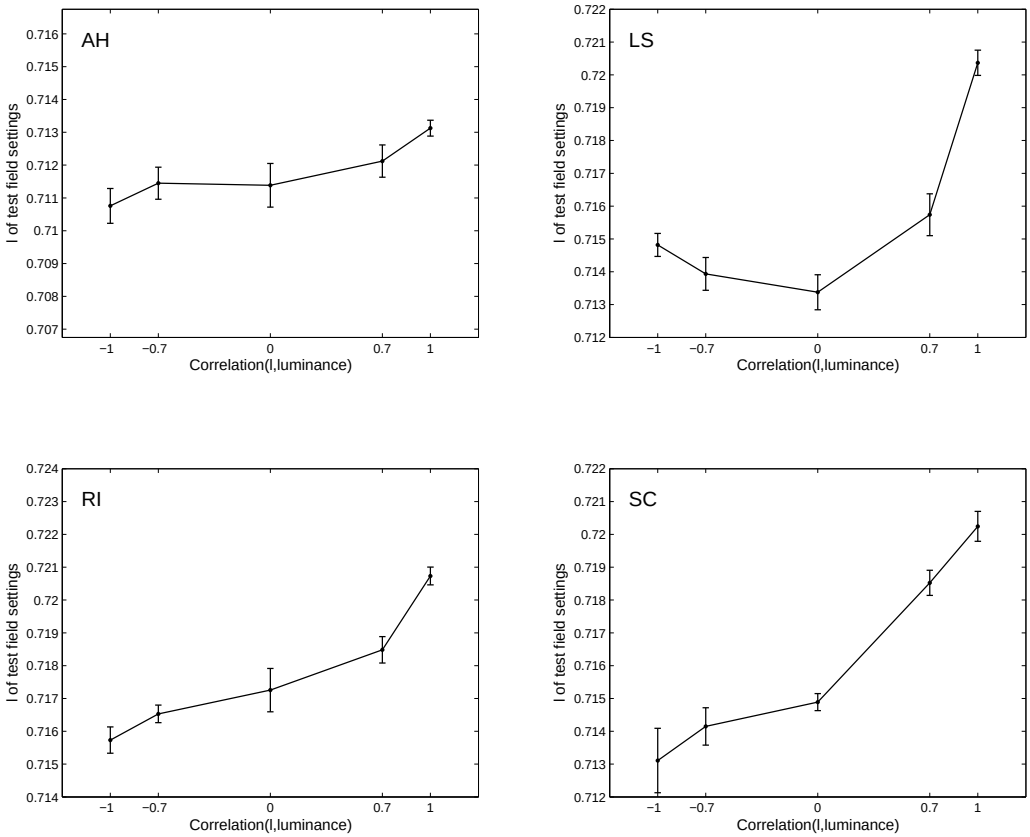

Figure 5.4: Results of experiment 2 for red stimuli.

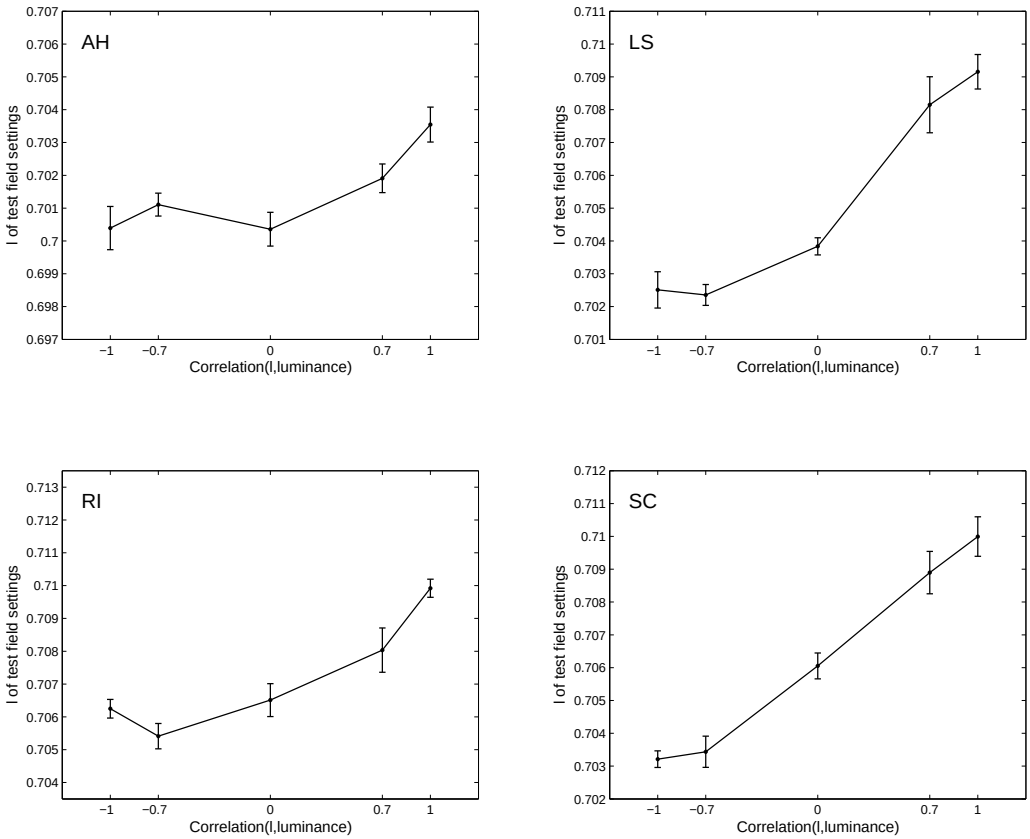

Figure 5.5: Results of experiment 2 for yellow stimuli.

5.3.4 Discussion

The results of experiment 2 extend the basic effect of the luminance-redness correlation found in experiment 1 to surrounds with non-neutral mean chromaticities. For the four surround colours used, the results were all significant and similar to the results of experiment 1. No indication was found that the influence of the luminance-redness correlation on the perceived colour of the test field differed for the different surrounds. Thus, it seems that the visual system uses the luminance-redness correlation as a signed cue independently of the mean image chromaticity and that the way in which these two cues are integrated does not deviate drastically from simple additivity.

5.4 Experiment 3

5.4.1 Luminance-redness correlation: luminance-weighted mean

The basic idea of experiment 1 (as well as experiment 2) is to vary the luminance-redness correlation while at the same time keeping all other statistics of the chromatic distribution of the colours in the surround constant. Especially the mean chromaticity of the surround has to be the same for all conditions of luminance-redness correlation since this statistic is known to strongly influence the test field settings. The mean chromaticitiy is calculated as the space average of the chromaticity values l and s within the surround. One possible problem for the conclusions based on this experiment is that one could calculate the mean chromaticity in a different way (namely by weighting the chromaticity of each pixel in the surround by its luminance, formula see below) and that the so calculated luminance-weighted mean of l (but not of s) covaries with the luminance-redness correlation. The reason for this covariation is obvious: when the (non-luminance-weighted) mean chromaticity remains constant but the luminance-redness correlation increases, the more reddish pixel become lighter. Thus they get more highly weighted when entering into the luminance-weighted mean redness and the weights of the darker greenish pixels decrease. Thus, the luminance-weighted mean redness shifts toward a more reddish chromaticity (higher l value). This also shows that it is not possible to keep constant at the same time the non-luminance-weighted mean redness as well as the luminance-weighted mean redness when the luminance-redness correlation is varied.

Consequently, the question arises whether the results of experiment 1 are not really due to the manipulation of the luminance-redness correlation but due to the

variation in luminance-weighted mean redness. In earlier work (MacLeod & Golz, 2003) Donald I. A. MacLeod and I have shown that this worry is unwarranted: the estimated shifts of grey settings induced by changes in luminance-weighted mean chromaticity are substantially smaller than the experimentally observed shifts that result from changes in luminance-redness correlation. Here I will address and refute this objection more directly: experiment 3 will be equal to experiment 1 in every aspect except that now the luminance-weighted mean of l will be kept constant for all conditions of luminance-redness correlation. Then the non-luminance-weighted mean of l covaries inevitably with the luminance-redness correlation. However, if the visual system takes only the non-luminance-weighted mean into account, the changes in subjects' grey settings would be in the opposite direction of the changes expected by the manipulation of the luminance-redness correlation. Thus, if the results show the expected dependence (the hypothesis is like in experiment 1 that the higher the luminance-redness correlation is in the surround, the higher is the l-value of subjects' grey settings for the test field) this could not be explained by the assumption that the visual system relies only on either the luminance-weighted mean chromaticity or the non-luminance-weighted mean chromaticity. Instead, such results would be strong evidence that the visual system takes into account the luminance-redness correlation for the perception of surface and illumination colour.

5.4.2 Methods

The luminance-weighted mean of l for the n pixels in the surround (or in each of the concentric areas for which it is kept constant as well, see page 62) is calculated as follows:

$$mean_{luminance-weighted}(l) = \sum_{i=1}^{n} \frac{l_i * luminance_i}{n * mean(luminance)}$$

This is a weighted mean of the l values of all pixels where the weights are each pixels' luminance divided by the average luminance. Since $l = L/luminance$ and $luminance = L + M$, one can also write:

$$\begin{aligned} mean_{luminance-weighted}(l) &= \sum_{i=1}^{n} \frac{L_i}{n * mean(luminance)} \\ &= \frac{mean(L)}{mean(luminance)} \\ &= \frac{mean(L)}{mean(L) + mean(M)} \end{aligned}$$

While for the non-luminance weighted mean of l the cone excitations L and M are first transformed to the MacLeod-Boynton chromaticity coordinate $l = L/(L+M)$ and than averaged over all pixels, for the luminance-weighted mean the cone excitations are first averaged, and then the resulting means are subjected to a corresponding transformation.

Stimuli

The stimuli used in experiment 3 are equal to those used in experiment 1 in every aspect except that instead of the non-luminance-weighted mean of l within the surround the luminance-weighted mean of l is kept constant with a value of 0.6877 for all five conditions of luminance-redness correlation. In summary, the stimuli used in experiment 3 share the following features with those used in experiment 1 (for details see page 62):

- They consist of overlapping circles of various colours with the test field in the center of the display on top of all circles in that area.
- They have the same size and spatial layout as the stimuli in experiment 1.
- The redness-luminance correlation is varied in five conditions with the values -1.0, -0.7, 0.0, +0.7, and +1.0.
- The values of all other statistics are kept constant and are chosen to resemble the values of natural scenes (cf. Tab. 5.1).

Subjects

Eight subjects (3 male, 5 female, aged from 20 to 33 years) participated in experiment 3. All subjects had previously participated in experiment 1, had normal or corrected-to-normal vision and showed no colour vision deficiency, as tested by the Ishihara colour plates. All subjects except for one (JG, the author) were naïve about the design and purpose of the experiment.

Procedure

The procedure equaled in every aspect (experimental setup, task, instructions, sequence of presentations etc.) exactly that of experiment 1 (see page 63).

Subject	$\hat{\psi}$	t_{emp}
GW	0.0022	12.04 ***
HS	0.0016	5.54 ***
JG	0.0026	13.11 ***
JJ	0.0024	8.95 ***
ML	0.0013	4.01 ***
MP	0.0029	9.56 ***
RI	0.0013	10.05 ***
SC	0.0015	6.13 ***

Table 5.8: Results of experiment 3.

5.4.3 Results

The results for all of the eight subjects show the expected dependence of the grey settings on the luminance-redness correlation in the surround (cf. Fig. 5.6): the higher the luminance-redness correlation, the higher on average the l-value of subjects' grey settings. These effects are individually significant for all of the eight subjects, as Tab. 5.8 shows. The measure for the size of the effect $\hat{\psi}$ ranges from 0.0013 (subjects ML and RI) to 0.0029 (subject MP) with an average value of 0.001975. This average of all eight subjects is approximately one fifth lower than the average effect size of 0.00255 in experiment 1.

5.4.4 Discussion

The results of experiment 3 clearly show that the basic effect of the luminance-redness correlation found in experiment 1 is not caused by variation in luminance-weighted mean redness within the surrounds. An only slightly weaker and still substantial effect is found for all of the eight subjects in experiment 3 when the luminance-weighted mean redness is kept constant for all conditions of luminance-redness correlation. This is strong evidence that the visual system indeed takes into account the luminance-redness correlation for the perception of surface and illumination colour. Thus, it seems that the visual system has internalized the corresponding statistical regularity of the chromatic world found in the Gaussian World analysis as well as in the real world simulation, namely that a high luminance-redness correlation within the retinal image indicates a reddish illumination.

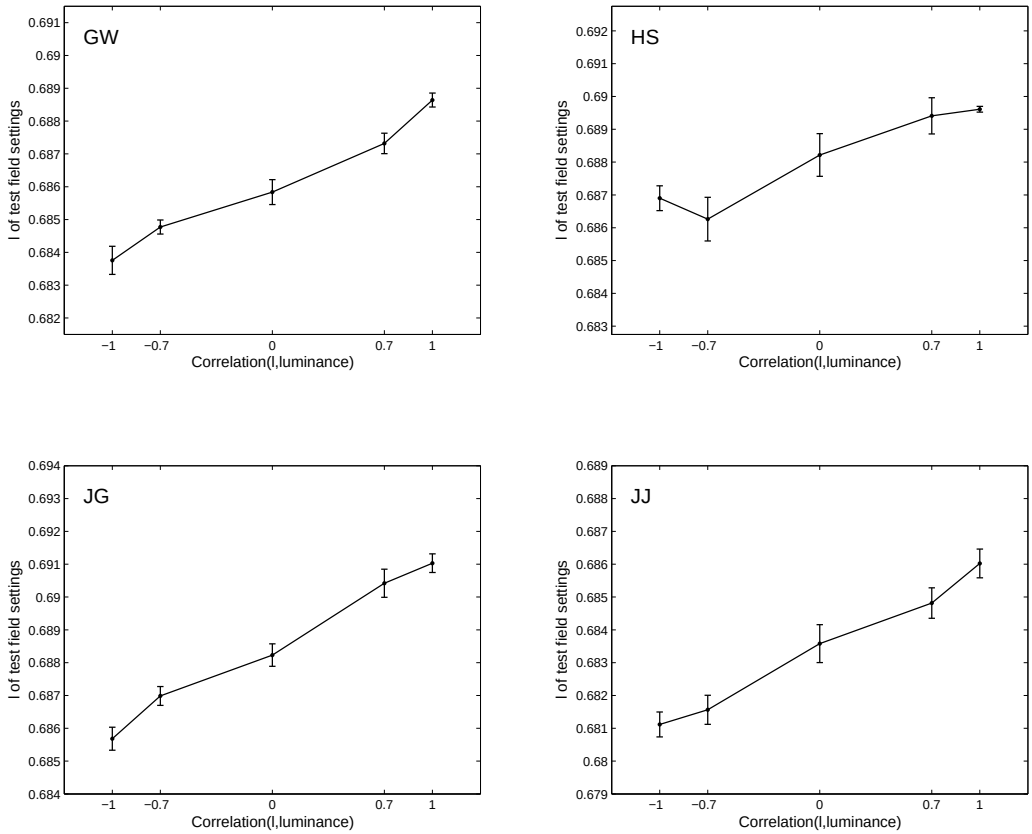

Figure 5.6: Results of experiment 3.

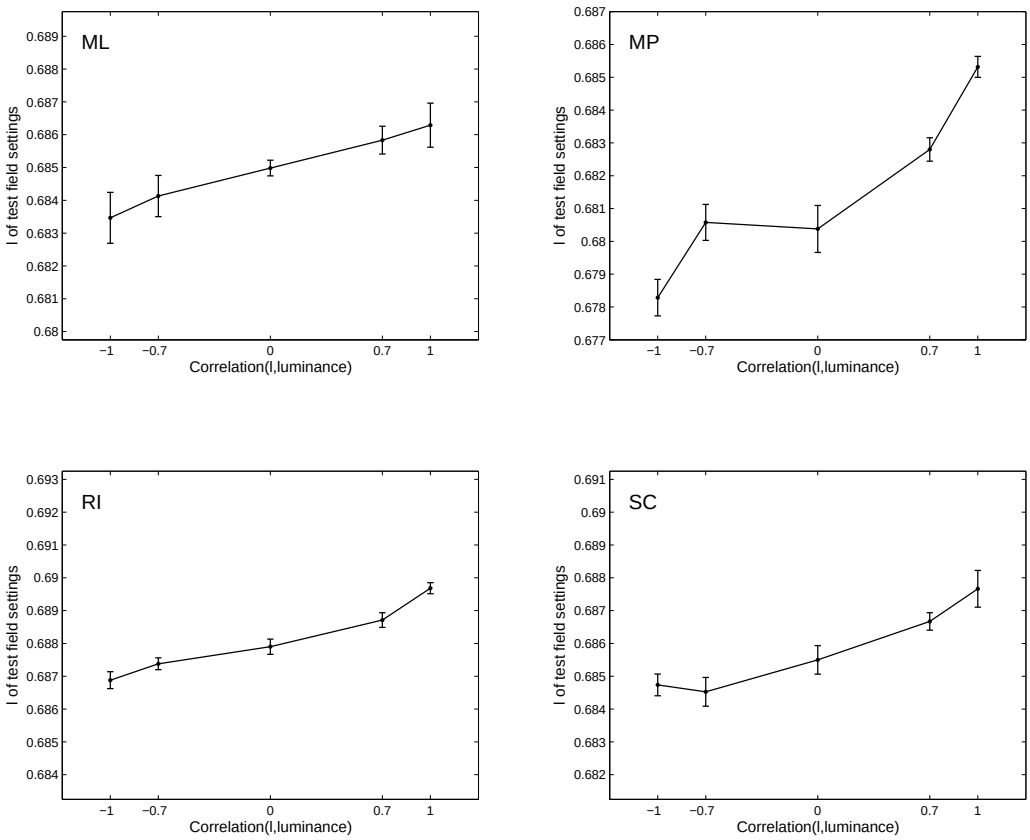

Figure 5.6: (Continuation)

5.5 Experiment 4

5.5.1 Luminance-blueness correlation: basic experiment

The aim of this experiment is to investigate whether the human visual system takes into account for the perception of surface and illumination colour the correlation between luminance and blueness (s value) within the retinal image (very much like experiment 1 did for the luminance-redness correlation). More precisely, the hypothesis is — based on the results of the Gaussian World analysis and the real world simulation — that the higher the luminance-blueness correlation in the stimulus is the more bluish becomes the illumination inferred by the visual system. Consequently, the visual system would have to take this estimated illumination into account when determining perceived surface colours: if the estimated illumination becomes more bluish, a retinal area of constant chromaticity would look like a less bluish surface. For the task of experiment 4, viz. grey settings in variegated surrounds with five different conditions of luminance-blueness correlation, the following prediction results: for higher correlation values, the visual system infers a more bluish illumination and in order to keep the appearance of the test field constantly grey subjects adjust the test field to more bluish chromaticities. (A test field with a chromaticity that looks grey in an a condition with low luminance-blueness correlation would look yellowish in a condition with high luminance-blueness correlation because, loosely speaking, more of the constant retinal input for the test field is attributed to the more bluish illumination and thus less blueness can be attributed to the surface. In order to make the test field appear grey they thus would have to pick a more bluish chromaticity.) Thus, the precise hypothesis is that the higher the luminance-blueness correlation is in the surround, the higher is the s-value of subjects' grey settings for the test field.

5.5.2 Methods

Stimuli

For experiment 4, the same type of stimuli is used as for the first three experiments, the only difference being that now it is the luminance-blueness correlation that is varied as the independent variable in five conditions with the values -1.0, -0.7, 0.0, +0.7, and +1.0. The luminance-redness correlation is set to 0.078 (which is the average value of that statistic in the natural images of Ruderman et al. (1998) under 7000 K daylight illumination) and kept constant for all conditions of luminance-blueness correlation. Tab. 5.9 lists the values of all statistics of the chromatic distribution within the surround of the stimuli for experiment 4. Tab.

C.6 in Appendix C.3 provides another specification of the chromatic distribution of the surround colours which is equivalent to the above description: it lists the eigenvectors and the corresponding standard deviations of the chromatic distribution in $(l, s, luminance)$-space for all five conditions of luminance-blueness correlation. The middle row of Plate IV shows examples of the stimuli for experiment 4. One can see that in the stimuli with luminance-blueness correlation of -1.0, the bluish circles are darker than the yellowish/brownish ones while the opposite holds for the stimuli for condition with luminance-blueness correlation of +1.0.

Besides the above-mentioned differences due to the different independent variable, the stimuli used in experiment 4 are equal to those used in experiment 1 in every other aspect. In summary, the stimuli used in experiment 4 share the following features with those used in experiment 1 (for details see page 62):

- They consist of overlapping circles of various colours with the test field in the center of the display on top of all circles in that area.
- They have the same size and spatial layout as the stimuli in experiment 1.
- The values of all other statistics except for the independent variable (luminance-blueness correlation) are kept constant and are chosen to resemble the values of natural scenes (cf. Tab. 5.9).

Subjects

Six subjects (3 male, 3 female, aged from 20 to 33 years) participated in experiment 4. All subjects had normal or corrected-to-normal vision and showed no colour vision deficiency, as tested by the Ishihara colour plates. All subjects had previously participated in experiment 1 and were naïve about the design and purpose of the experiment.

Procedure

The procedure equaled in every aspect (experimental setup, task, instructions, sequence of presentations etc.) exactly that of experiment 1 (see page 63).

Statistic	*Value*
Means	
l	0.6877
s	1.1466
luminance	20.0 cd/m^2
Standard deviations	
l	0.005
s	0.1536
luminance	5.0
Correlations	
l and luminance	0.078
s and luminance	5 conditions:
	-1.0; -0.7; 0.0; 0.7; 1.0
l and s	-0.2133

Table 5.9: Chromatic statistics of stimuli for experiment 4. The listed values are the means, standard deviations, and correlations of the chromatic distribution of the pixels within the sourround of the stimuli.

5.5.3 Results

The results for all of the six subjects show a dependence of the grey settings on the luminance-blueness correlation in the surround that is in accordance with the hypothesis stated above: the higher the luminance-blueness correlation, the higher on average the s-value of subjects' grey settings (cf. Fig. 5.7).[4] But as Tab. 5.10 shows, this effect is individually significant only for five of the six subjects. (The non-significant results of subject TP will not be taken into consideration in the following discussion of the size of the effect.) The size of the effect $\hat{\psi}$ ranges from 0.0381 (subject GW) to 0.0824 (subject JJ) with an average value of 0.055. Thus, for the average subject, the s-value of the grey setting increases by $2\times0.055 = 0.11$ when the luminance-blueness correlation increases from -1.0 to +1.0. For comparison, the average standard deviation of s in the natural images of Ruderman et al. (1998) is 0.1536 (this value has also been chosen for the standard deviation of s in the surround of the experimental stimuli) and the difference in s between the two most extreme unimpeded daylights (4000 K and

[4]Note that in all subplots in Fig. 5.7 (as well as in Fig. 5.8 for experiment 5) the ordinate always cover the same range of s-values in order to make the figures easily comparable. The upper and lower limit of this range might differ, though, due to individual differences in the overall level of the chromaticity of grey settings. Because subject JJ shows an extraordinarily large effect in experiment 5 which requires a large scale for the ordinate, the effects in the subplots of many other subjects in experiment 4 and 5 appear comparatively less salient due to this large scale common to all subplots.

Subject	$\hat{\psi}$	t_{emp}
GW	0.0381	3.92 ***
JJ	0.0824	5.20 ***
MP	0.0661	8.17 ***
RI	0.0453	7.93 ***
SC	0.0428	9.53 ***
TP	0.0075	1.47 $^{n.s.}$

Table 5.10: Results of experiment 4.

20 000 K) is 1.27, so a little bit more than eleven times the above-mentioned average subject shift of 0.11.

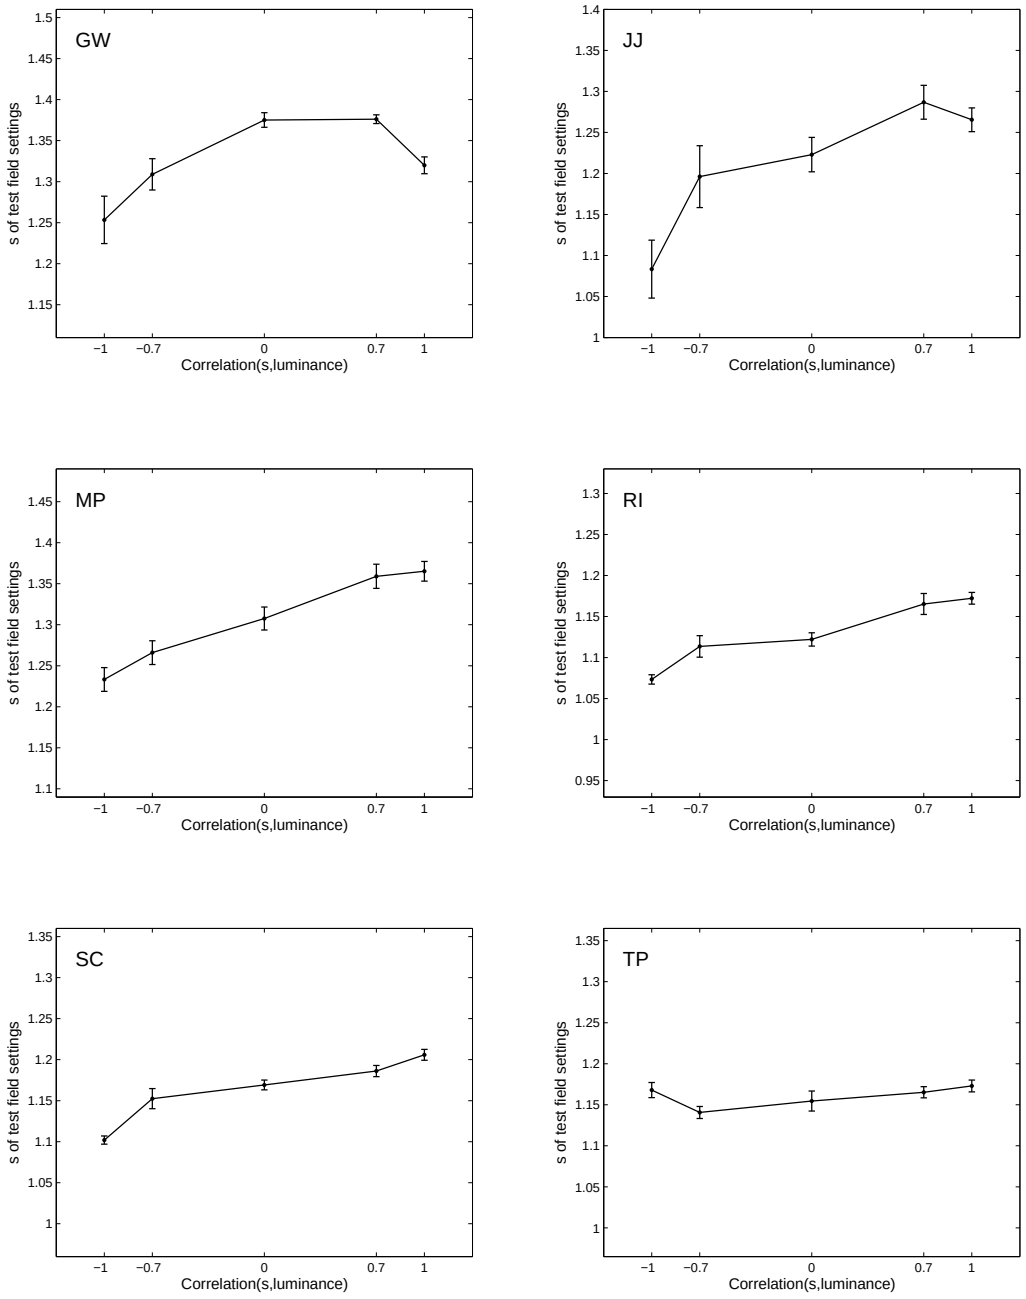

Figure 5.7: Results of experiment 4.

5.5.4 Discussion

In summery, experiment 4 has yielded strong evidence that the visual system indeed takes into account the luminance-blueness correlation for the perception of surface and illumination colour. This effect seems to be robust in that the results of five of six subjects tested are significant and the size of this effect is substantial but not overwhelmingly large.

The theoretical evidence that the correlation between luminance and blueness in the retinal image is helpful for estimating the illuminant colour has been found to be ambiguous in the following sense: the Gaussian World analysis predicts that the bluish surfaces (with high s values) increase in luminance more than e.g. greyish or yellowish surfaces when the illumination becomes bluish (cf. page 29). This regularity has indeed been found to hold for all scenes in the analysis of the real world simulation on the level of individual pixels (cf. page 41). But surprisingly, in the analysis of the real world simulation on the level of scene statistics the correlation between luminance and blueness within a scene is almost independent of illumination (cf. page 51). With respect to this mixed theoretical evidence it is very interesting that the human visual system seems to have internalized the luminance-blueness correlation as a cue for the illumination as the results of experiment 4 suggest.

5.6 Experiment 5

5.6.1 Luminance-blueness correlation: luminance-weighted mean

In experiment 4 the luminance-redness correlation has been varied while at the same time keeping all other statistics of the chromatic distribution of the colours in the surround constant, especially the mean chromaticity of the surround since this statistic is known to strongly influence the test field settings. But, like in the case of mean redness in experiment 3 (as explained on page 78), the mean blueness can also be calculated in two ways: a non-luminance-weighted mean of s (which has been kept constant for all conditions of luminance-blueness correlation in experiment 4) and a luminance-weighted mean of s (which has covaried with the luminance-blueness correlation in experiment 4). The reason for this covariation is obvious: when the (non-luminance-weighted) mean chromaticity remains constant but the luminance-blueness correlation increases, the more bluish pixel become lighter. Thus they get more highly weighted when entering

into the luminance-weighted mean blueness and the weights of the darker yellowish/brownish pixels decrease. Thus, the luminance-weighted mean blueness shifts toward a more bluish chromaticity (higher s value). This also shows that it is not possible to keep constant at the same time the non-luminance-weighted mean blueness as well as the luminance-weighted mean blueness when the luminance-blueness correlation is varied.

Consequently, the question arises whether the results of experiment 4 are not really due to the manipulation of the luminance-blueness correlation but due to the variation in luminance-weighted mean blueness. The aim of experiment 5 is therefore to show that the luminance-blueness correlation also systematically influences subjects' grey settings when the luminance-weighted mean chromaticity is kept constant (very much like what experiment 3 showed for the luminance-redness correlation). Thus, experiment 5 will be equal to experiment 4 in every aspect except that now the luminance-weighted mean of s will be kept constant for all conditions of luminance-blueness correlation. Then the non-luminance-weighted mean of s covaries inevitably with the luminance-blueness correlation. However, if the visual system takes only the non-luminance-weighted mean into account, the changes in subjects' grey settings would be in the opposite direction of the changes expected by the manipulation of the luminance-blueness correlation. Thus, if the results show the expected dependence (the hypothesis is like in experiment 4 that the higher the luminance-blueness correlation is in the surround, the higher is the s-value of subjects' grey settings for the test field) this could not be explained by the assumption that the visual system relies only on either the luminance-weighted mean chromaticity or the non-luminance-weighted mean chromaticity. Instead, such results would be strong evidence that the visual system takes into account the luminance-blueness correlation for the perception of surface and illumination colour.

5.6.2 Methods

The luminance-weighted mean of s for the n pixels in the surround (or in each of the concentric areas for which it is kept constant as well, see page 62) is calculated as follows:

$$mean_{luminance-weighted}(s) = \sum_{i=1}^{n} \frac{s_i * luminance_i}{n * mean(luminance)}$$

This is a weighted mean of the s values of all pixels where the weights are each pixels' luminance divided by the average luminance. Since $s = S/luminance$ and $luminance = L + M$, one can also write:

$$
\begin{aligned}
mean_{luminance-weighted}(s) &= \sum_{i=1}^{n} \frac{S_i}{n * mean(luminance)} \\
&= \frac{mean(S)}{mean(luminance)} \\
&= \frac{mean(S)}{mean(L) + mean(M)}
\end{aligned}
$$

While for the non-luminance weighted mean of s the cone excitations L, M, and S are first transformed to the MacLeod-Boynton chromaticity coordinate $s = S/(L + M)$ and than averaged over all pixels, for the luminance-weighted mean the cone excitations are first averaged and then the resulting means are subjected to a corresponding transformation.

Stimuli

Subjects

The same six subjects as in experiment 4 (3 male, 3 female, aged from 20 to 33 years) participated in experiment 5. All subjects had normal or corrected-to-normal vision, showed no colour vision deficiency, as tested by the Ishihara colour plates, and were naïve about the design and purpose of the experiment.

Procedure

The procedure equaled in every aspect (experimental setup, task, instructions, sequence of presentations etc.) exactly that of experiment 1 (see page 63).

5.6.3 Results

The results for all of the six subjects show the expected dependence of the grey settings on the luminance-blueness correlation in the surround (cf. Fig. 5.8): the higher the luminance-blueness correlation, the higher on average the s-value of subjects' grey settings. As Tab. 5.11 shows, these effects are individually significant for all of the six subjects (even for subject TP whose results were not significant in experiment 4). The measure for the size of the effect $\hat{\psi}$ ranges from 0.0134 (subjects TP) to 0.1187 (subject JJ) with an average value of 0.052. This average of all eight subjects is very similar to the average effect size of 0.055 in experiment 4.

Subject	$\hat{\psi}$	t_{emp}
GW	0.0514	5.05 ***
JJ	0.1187	9.37 ***
MP	0.0329	3.57 ***
RI	0.0403	8.59 ***
SC	0.0557	10.60 ***
TP	0.0134	2.32 *

Table 5.11: Results of experiment 5.

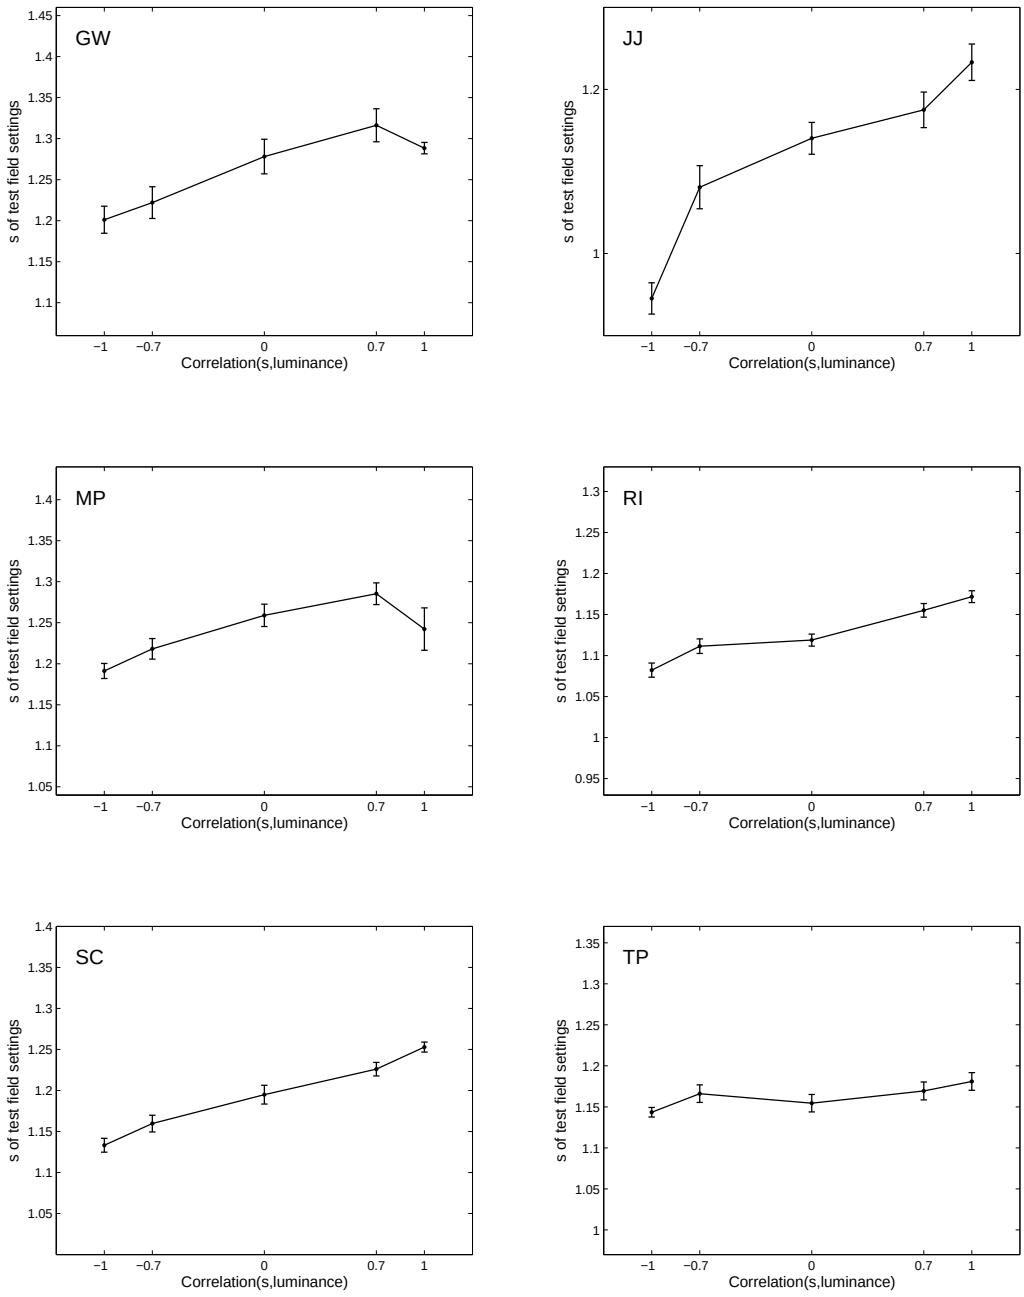

Figure 5.8: Results of experiment 5.

5.6.4 Discussion

The results of experiment 5 clearly show that the basic effect of the luminance-blueness correlation found in experiment 4 is not caused by variation in luminance-weighted mean blueness within the surrounds. A very robust and substantial effect, the size of which is similar to that found in experiment 4, is found for all of the six subjects in experiment 5 when the luminance-weighted mean blueness is kept constant for all conditions of luminance-blueness correlation. This is strong evidence that the visual system indeed takes into account the luminance-blueness correlation for the perception of surface and illumination colour. Thus, it seems that the visual system has internalized the corresponding statistical regularity found in the Gaussian World analysis and — with somewhat mixed results — in the real world simulation, namely that a high luminance-blueness correlation within the retinal image indicates a bluish illumination.

5.7 Experiment 6

5.7.1 Redness-blueness correlation: basic experiment

The aim of this experiment is to investigate whether the human visual system takes into account for the perception of surface and illumination colour the correlation between redness (l value) and blueness (s value) within the retinal image (very much like experiment 1 did for the luminance-redness correlation and experiment 4 did for the luminance-blueness correlation). Based on the results of the Gaussian World analysis and the real world simulation suggesting that the redness-blueness correlation increases when the illumination changes toward a more reddish 4000 K illuminant, the hypothesis is that the higher the redness-blueness correlation in the stimulus is the more reddish becomes the illumination inferred by the visual system. (The more reddish 4000 K illuminant is also less bluish than daylight illuminants with higher colour temperature, but, for the sake of simplicity, I will restrict the analysis of the experimental results to the chromaticity axis l). Consequently, the visual system would have to take this estimated illumination into account when determining perceived surface colours: if the estimated illumination becomes more reddish, a retinal area of constant chromaticity would look like a less reddish surface. For the task of experiment 6, viz. grey settings in variegated surrounds with five different conditions of redness-blueness correlation, the following prediction results: for higher correlation values, the visual system infers a more reddish illumination and in order to keep the appearance of the test field constantly grey subjects adjust the test field to more reddish chromaticities. (A test field with a chromaticity that looks grey

in an a condition with low redness-blueness correlation would look greenish in a condition with high redness-blueness correlation because, loosely speaking, more of the constant retinal input for the test field is attributed to the more reddish illumination and thus less redness can be attributed to the surface. In order to make the test field appear grey they thus would have to pick a more reddish chromaticity.) Thus, the precise hypothesis is that the higher the redness-blueness correlation is in the surround, the higher is the l-value of subjects' grey settings for the test field.

5.7.2 Methods

Stimuli

For experiment 6, the same type of stimuli is used as for the previous experiments, the only difference being that now it is the redness-blueness correlation that is varied as the independent variable in five conditions with the values -1.0, -0.7, 0.0, +0.7, and +1.0. The luminance-redness correlation is set to 0.078 and the luminance-blueness correlation to -0.1153 (which are the average values of these statistics in the natural images of Ruderman et al. (1998) under 7000 K daylight illumination) and they are kept constant for all conditions of luminance-blueness correlation. Tab. 5.12 lists the values of all statistics of the chromatic distribution within the surround of the stimuli for experiment 6. Tab. C.7 in Appendix C.3 provides another specification of the chromatic distribution of the surround colours which is equivalent to the above description: it lists the eigenvectors and the corresponding standard deviations of the chromatic distribution in $(l, s, luminance)$-space for all five conditions of redness-blueness correlation. The bottom row of Plate IV shows examples of the stimuli for experiment 6.

Besides the above-mentioned differences due to the different independent variable, the stimuli used in experiment 6 are equal to those used in experiment 1 in every other aspect. In summary, the stimuli used in experiment 6 share the following features with those used in experiment 1 (for details see page 62):

- They consist of overlapping circles of various colours with the test field in the center of the display on top of all circles in that area.
- They have the same size and spatial layout as the stimuli in experiment 1.
- The values of all other statistics except for the independent variable (redness-blueness correlation) are kept constant and are chosen to resemble the values of natural scenes (cf. Tab. 5.12).

Statistic	*Value*
Means	
l	0.6877
s	1.1466
luminance	20.0 cd/m^2
Standard deviations	
l	0.005
s	0.1536
luminance	5.0
Correlations	
l and luminance	0.078
s and luminance	-0.1153
l and s	5 conditions:
	-1.0; -0.7; 0.0; 0.7; 1.0

Table 5.12: Chromatic statistics of stimuli for experiment 6. The listed values are the means, standard deviations, and correlations of the chromatic distribution of the pixels within the sourround of the stimuli.

Subjects

Seven subjects (2 male, 5 female, aged from 20 to 25 years) participated in experiment 6. All subjects had normal or corrected-to-normal vision, showed no colour vision deficiency, as tested by the Ishihara colour plates, and were naïve about the design and purpose of the experiment. All subjects had previously participated in experiment 1.

Procedure

5.7.3 Results

The results of only four of the seven subjects show a dependence of the grey settings on the redness-blueness correlation in the surround that is in accordance with the hypothesis stated above: the higher the redness-blueness correlation, the higher on average the l-value of subjects' grey settings (cf. Fig. 5.9). Though the effects for subjects JJ, ML, RI, and TP are individually significant (cf. Tab. 5.13), at best the effect for subject ML can be considered substantial. The size of the effect $\hat{\psi}$ of these four subjects ranges from 0.00039 (subject TP) to 0.00175 (subject ML) with an average value of 0.00082. This average size of the four

Subject	$\hat{\psi}$	t_{emp}
HS	0.00022	0.72 $^{n.s.}$
JJ	0.00067	2.43 *
ML	0.00175	5.95 ***
MP	0.00044	1.23 $^{n.s.}$
RI	0.00045	2.92 *
SC	0.00011	0.53 $^{n.s.}$
TP	0.00039	2.01 *

Table 5.13: Results of experiment 6.

significant effects is less than one third of the average effect size of 0.00255 in experiment 1. Another comparison shows that the average effect of the significant subjects in experiment 6 is rather weak. For those subjects, the l-value of the grey setting increases on average by $2\times0.00082 = 0.00164$ when the redness-blueness correlation increases from -1.0 to +1.0. For comparison, the average standard deviation of l in the natural images of Ruderman et al. (1998) is 0.005 (this value has also been chosen for the standard deviation of l in the surround of the experimental stimuli) and the difference in l between the two most extreme unimpeded daylights (4000 K and 20 000 K) is 0.04, so approximately 24 times the above-mentioned average subject shift of 0.00164.

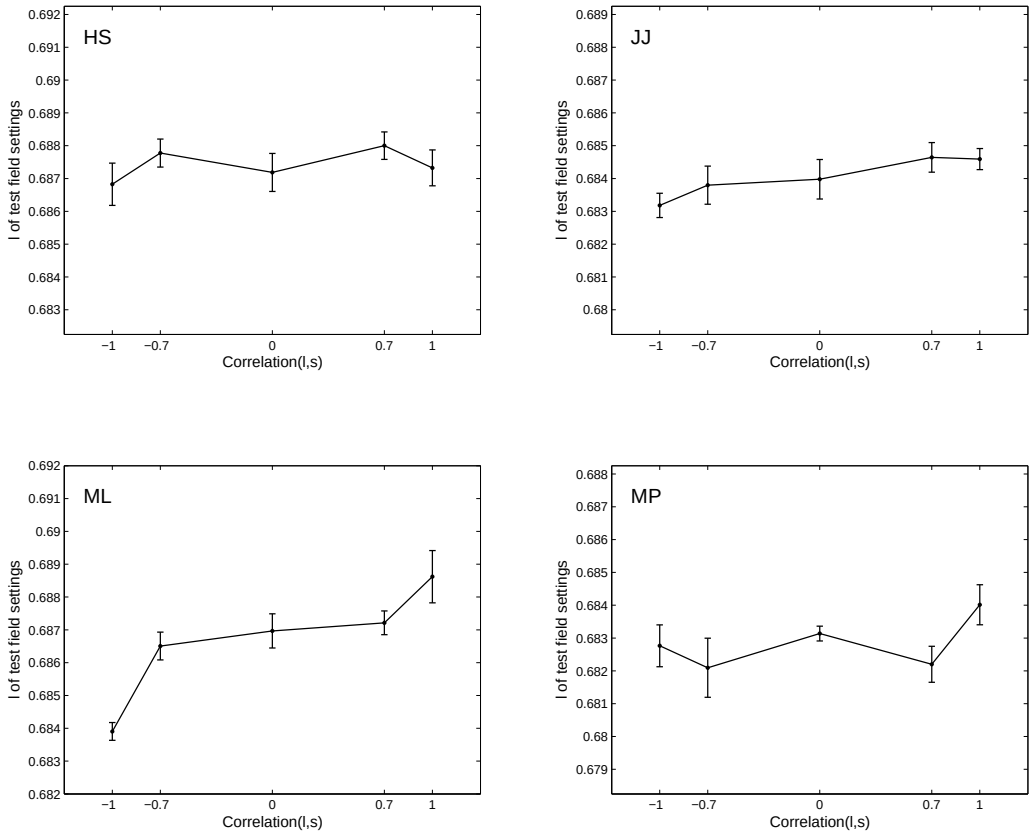

Figure 5.9: Results of experiment 6.

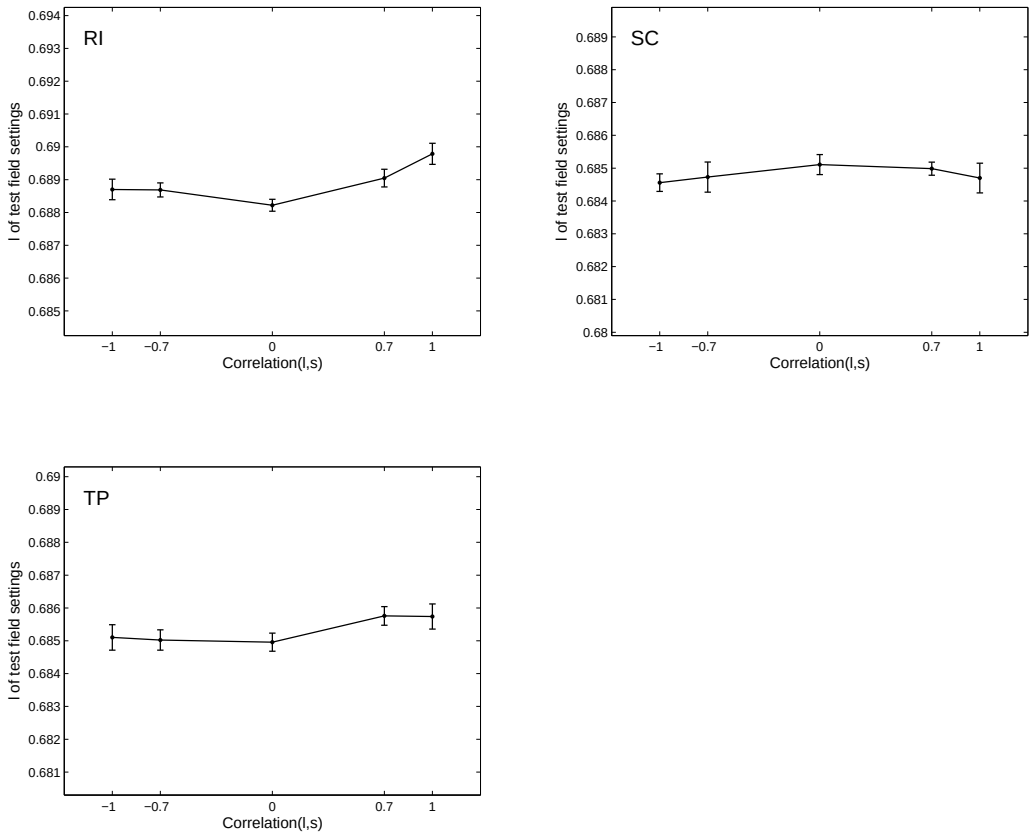

Figure 5.9: (Continuation)

5.7.4 Discussion

Unlike the results of the previous five experiments, all of which have shown robust and substantial effects of the respective correlations tested, the results of experiment 6 yield more mixed and less clear-cut evidence that the redness-blueness correlation within the retinal image is used as a cue for the illumination by the human visual system. The expected influence of this statistic on the grey settings was individually significant only for four of seven subjects tested in experiment 6. But for three of these four subjects the effect was rather weak.

One reason for these ambiguous results could be that the visual processes of the subjects mirror the fact that the usefulness of the redness-blueness correlation as a cue to the illuminant has been found in the theoretical analyses of chapter 3 and 4 to be ambiguous, too: The Gaussian World idealization predicts (cf. page 32) that the light reflected from surfaces with sufficiently narrow-band reflectance will undergo a smaller chromaticity shift under changes of illumination than the light from broad-band surfaces. Accordingly, the illumination-induced shift in the

real world simulation is greater for bluish and greyish surfaces than for yellowish surfaces which are more shift resistant due to their smaller bandwidth (cf. page 42). This leads, in the analysis of the real world simulation on the level of scene statistics, to an increase in the redness-blueness correlation when the illumination changes toward the more reddish/yellowish 4000 K illuminant. Unfortunately, this statistic depends also systematically on the (illumination-independent) scene-averaged surface redness, in a way that weakens the usefulness of this statistic as a cue to the illumination: the resulting constellation of this statistic to the mean image redness (Fig. 4.13) resembles to some degree the schematic constellation (d) in Plate I, where two cues are correlated both within and across illumination clusters such that the combination of the two cues is only as useful as each cue is alone (cf. section 2.3.2). On the basis of the real world simulation it thus remains an open question whether the redness-blueness correlation is helpful for estimating the illuminant in the natural environment.

The analysis of the real world simulation on the level of scene statistics indicates that it could be the case that the redness-blueness correlation is more useful for some illuminants than for others (see Fig. 4.13). Also, it might well be the case that the bandwidth effect underlying the illumination-dependence of the redness-blueness correlation is particularly strong for illuminants that themselves have a restricted bandwidth. Such illuminants do not occur in the set of unimpeded daylight illuminants from which the sample of illuminants used in the real world simulation were drawn, but they can occur in the natural environment, e.g. when light is filtered and reflected by leaves in forests (Endler, 1993). From these considerations it follows that the redness-blueness correlation could exert a stronger influence on subjects' grey settings than the one found in experiment 6 if the mean chromaticity of the stimuli is biased (instead of neutrally grey as in experiment 6). This possibility will be tested in the next experiment.

5.8 Experiment 7

5.8.1 Redness-blueness correlation: non-neutral stimuli

Following the considerations in the discussion of experiment 6, the aim of experiment 7 is to test whether larger effects of the redness-blueness correlation than the ones found in experiment 6 can be obtained by use of stimuli with non-neutral mean chromaticities. To this end, the same four chromaticities as in experiment 2 (labelled "blue", "green", "red", and "yellow") are used for the mean chromaticity of the stimuli in experiment 7. For each of these four surround types, a sub-experiment is carried out (in a separate session) that very much resembles

experiment 6, except for the mean chromaticity of the stimuli. That is, for each of the four mean surround chromaticities, the redness-blueness correlation in the surround is varied in five conditions and subjects have to adjust a central test field such that it appears grey to them. For each of the four sub-experiments the hypothesis is that the higher the redness-blueness correlation is in the surround, the higher is the l-value of subjects' grey settings for the test field. The rationale for this is exactly analogous to the one in experiment 6 (see page 94). Neither the Gaussian World analysis of chapter 3 nor the real world simulation of chapter 4 provide any specific predictions for which mean image chromaticities this dependence on the redness-blueness correlation should be stronger and for which it should be weaker. Thus, the treatment of this question by use of the four different mean surround chromaticities is purely exploratory.

5.8.2 Methods

Stimuli

The stimuli used in experiment 7 are equal to those used in experiment 6 in every aspect except for the mean chromaticities of the surround, i.e. they share the following features (for details see also page 62):

- They consist of overlapping circles of various colours with the test field in the center of the display on top of all circles in that area.
- They have the same size and spatial layout as the stimuli in experiment 1.
- The redness-blueness correlation is varied in five conditions with the values -1.0, -0.7, 0.0, +0.7, and +1.0.
- The values of all other statistics are kept constant and are chosen to resemble the values of natural scenes (cf. Tab. 5.12).

Only the mean chromaticities of the surround differ from the neutral one of experiment 6. The four mean chromaticities used are the same as those used in experiment 2 (see Tab. 5.3). For a description of how they are determined see page 71.

Subjects

The same seven subjects as in experiment 6 (2 male, 5 female, aged from 20 to 25 years) participated in experiment 7. All subjects had normal or corrected-to-normal vision, showed no colour vision deficiency, as tested by the Ishihara colour plates, and were naïve about the design and purpose of the experiment.

Procedure

The procedure equaled in every aspect (experimental setup, task, instructions, sequence of presentations etc.) exactly that of experiment 1 (see page 63), except that instead of one session for the entire experiment 1 four sessions were needed for experiment 7 — one session for each mean surround chromaticity. All subjects performed the four sessions in the order blue, green, red, and finally yellow mean surround chromaticity.

5.8.3 Results

The results of experiment 7 differ substantially for the four different mean surround chromaticities. In the case of the green surround (cf. Fig. 5.11) as well as in the case of the yellow surround (cf. Fig. 5.13) the results show for all of the seven subjects the expected dependence of the grey settings on the redness-blueness correlation within the surround: the higher the redness-blue correlation, the higher on average the l-value of subjects' grey settings. These effects are individually significant for all seven subjects for the green surround (cf. Tab. 5.15) as well as for the yellow surround (cf. Tab. 5.17).

But no such effects are obvious for the blue surround (cf. Fig. 5.10) as well as for the red surround (cf. Fig. 5.12). Though two of the seven subjects (subjects ML and SC) have a significant result in the case of the blue surround and one subject (subject MP) has a significant result in the case of the red surround, these effect are rather small (the measure $\hat{\psi}$ is less than 0.001 in all of these three cases).

Very different from these negative results for the blue and the red surround, it emerges clearly an effect of the redness-blueness correlation for the green as well as for the yellow surround. Not only are for the latter surrounds, as mentioned above, the results of all of the seven subjects significant, but also are these effects substantial with an effect size averaged over all subjects of 0.0019 for the green surround and 0.00153 for the yellow surround. For comparison, the average effect size of the luminance-redness correlation in experiment 1 was 0.00255.

Subject	$\hat{\psi}$	t_{emp}
HS	-0.00059	-1.63 $^{n.s.}$
JJ	-0.00004	-0.09 $^{n.s.}$
ML	0.00091	3.00 **
MP	-0.00031	-1.21 $^{n.s.}$
RI	-0.00001	-0.07 $^{n.s.}$
SC	0.00050	2.22 *
TP	-0.00081	-3.15 $^{n.s.}$

Table 5.14: Results of experiment 7 for blue stimuli.

Subject	$\hat{\psi}$	t_{emp}
HS	0.00244	9.07 ***
JJ	0.00228	6.59 ***
ML	0.00296	8.72 ***
MP	0.00141	5.72 ***
RI	0.00211	8.30 ***
SC	0.00079	2.69 **
TP	0.00128	5.83 ***

Table 5.15: Results of experiment 7 for green stimuli.

Subject	$\hat{\psi}$	t_{emp}
HS	-0.00054	-1.43 $^{n.s.}$
JJ	0.00106	1.59 $^{n.s.}$
ML	-0.00001	-0.02 $^{n.s.}$
MP	0.00093	2.55 **
RI	-0.00018	-0.48 $^{n.s.}$
SC	-0.00009	-0.23 $^{n.s.}$
TP	-0.00051	-1.96 $^{n.s.}$

Table 5.16: Results of experiment 7 for red stimuli.

Subject	$\hat{\psi}$	t_{emp}
HS	0.00217	6.07 ***
JJ	0.00112	3.20 **
ML	0.00197	4.00 ***
MP	0.00142	4.20 ***
RI	0.00102	4.59 ***
SC	0.00160	5.77 ***
TP	0.00139	3.72 ***

Table 5.17: Results of experiment 7 for yellow stimuli.

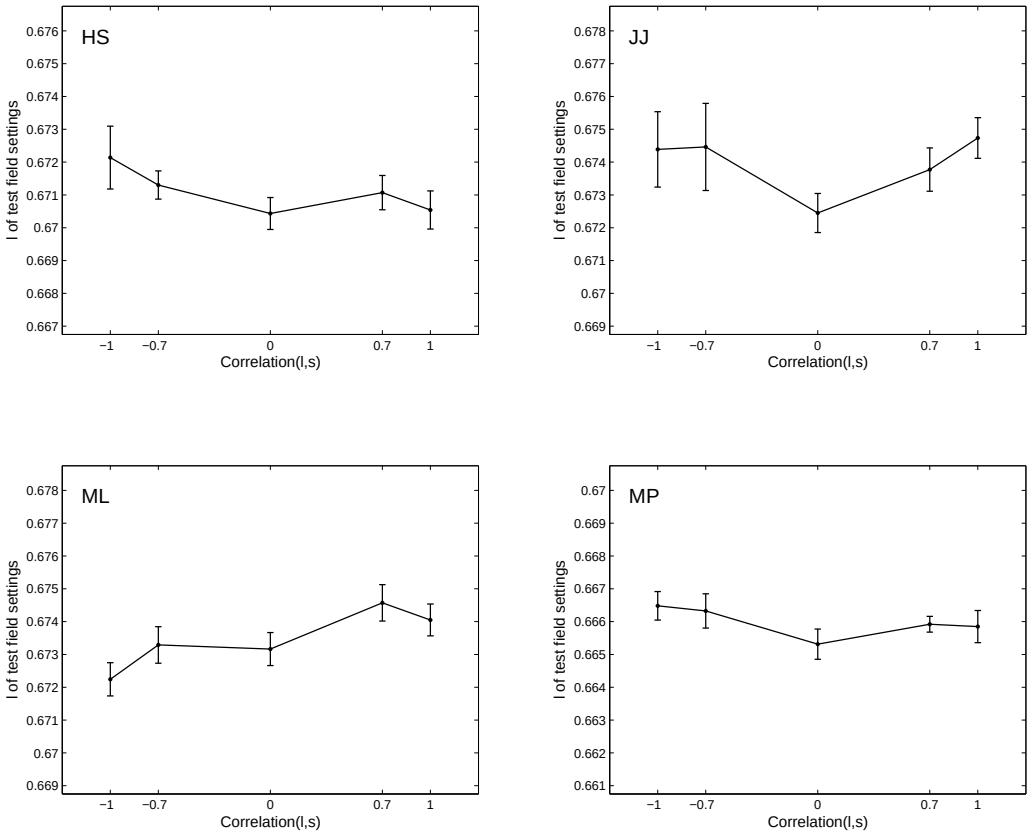

Figure 5.10: Results of experiment 7 for blue stimuli.

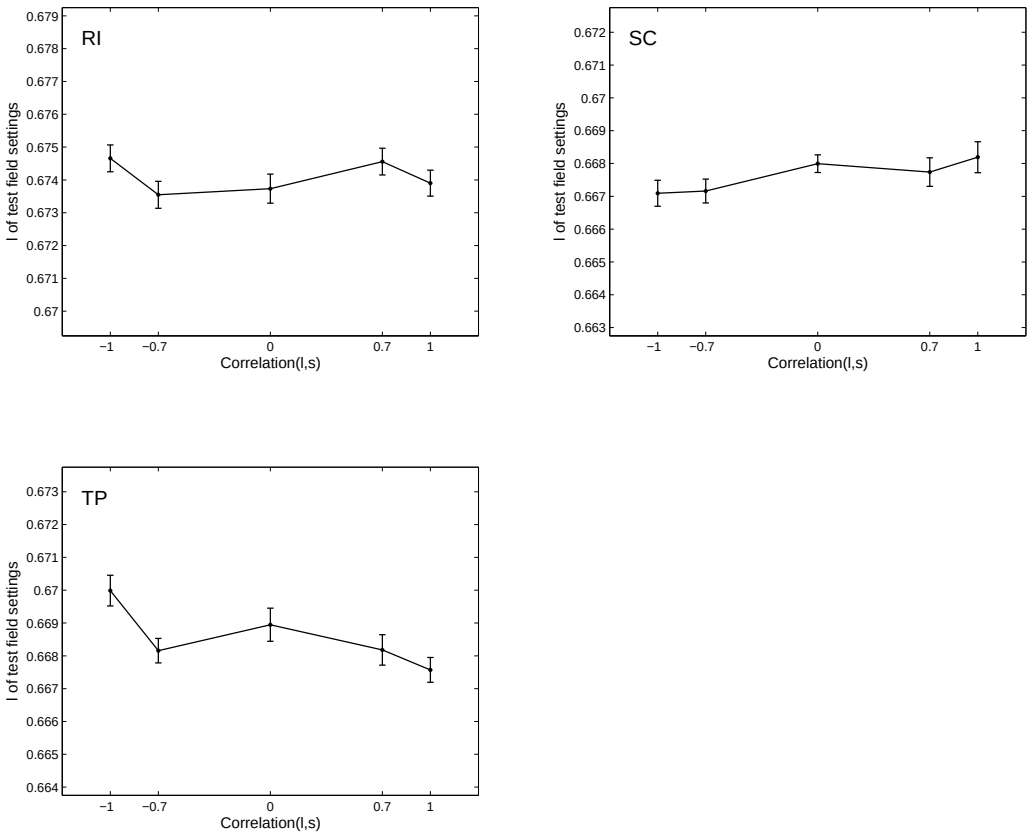

Figure 5.10: (Continuation)

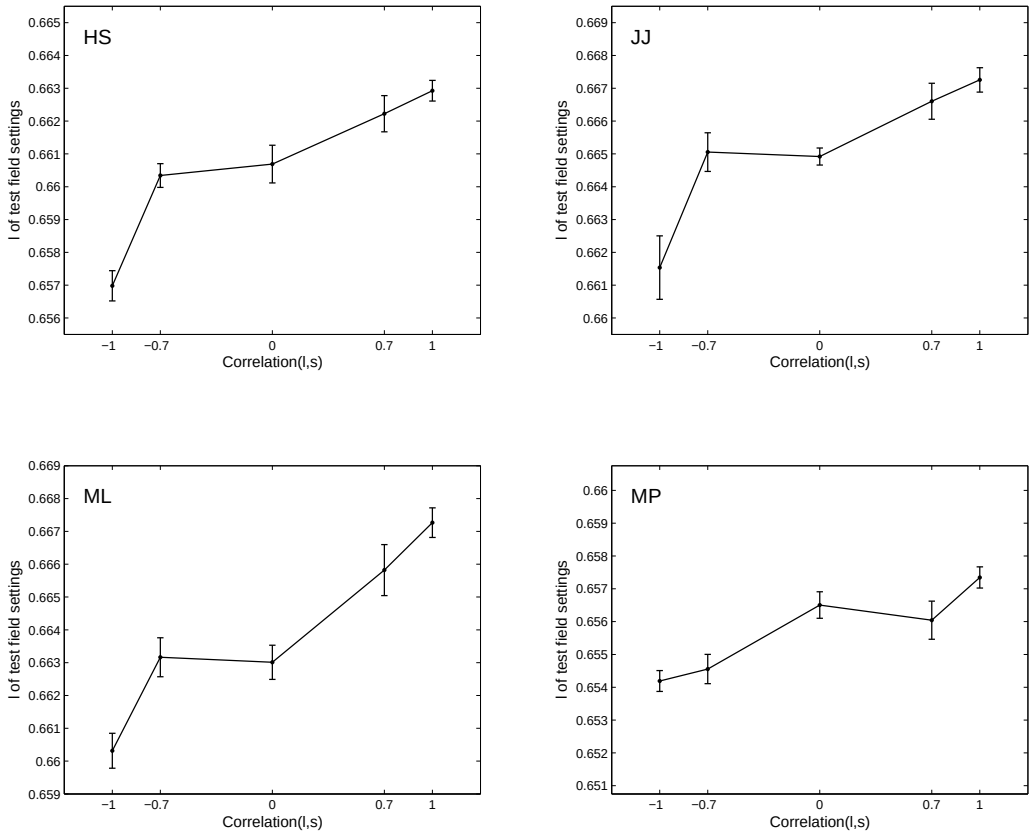

Figure 5.11: Results of experiment 7 for green stimuli.

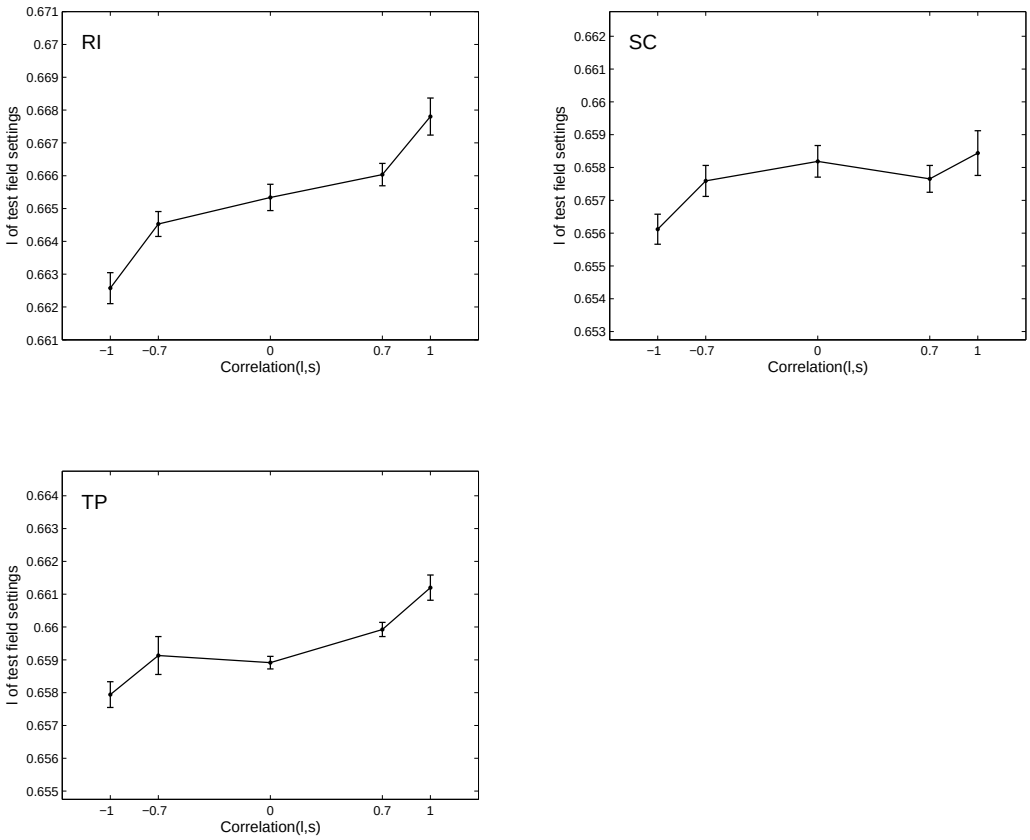

Figure 5.11: (Continuation)

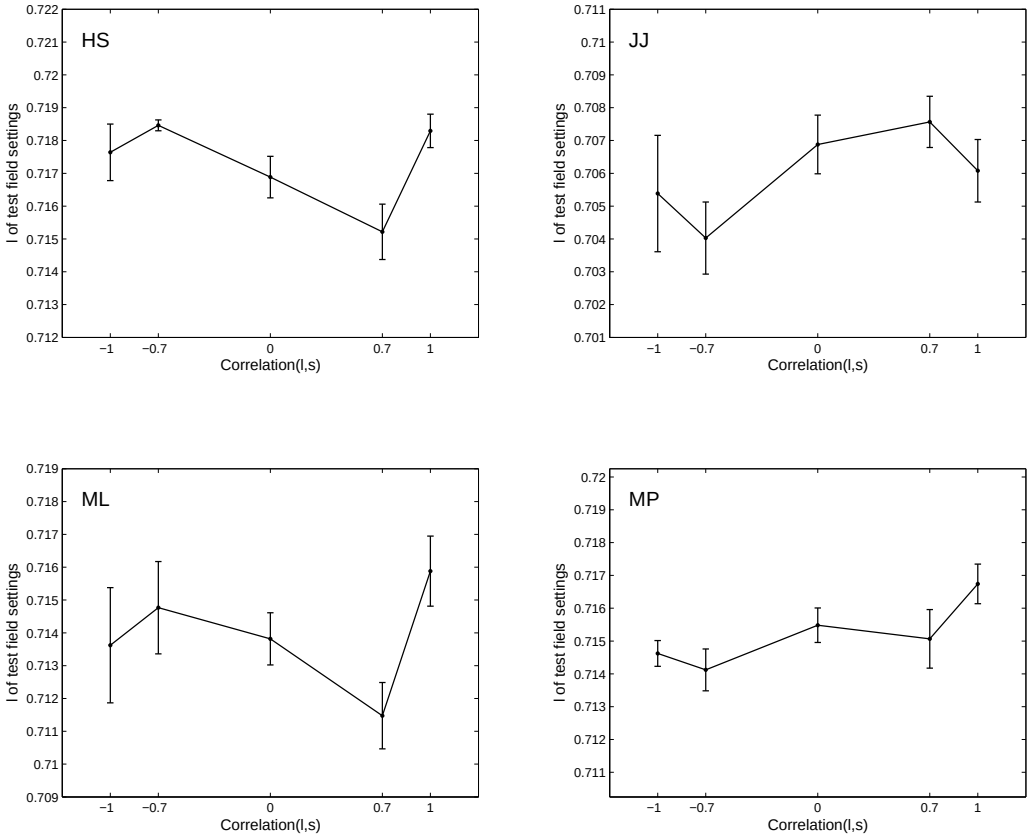

Figure 5.12: Results of experiment 7 for red stimuli.

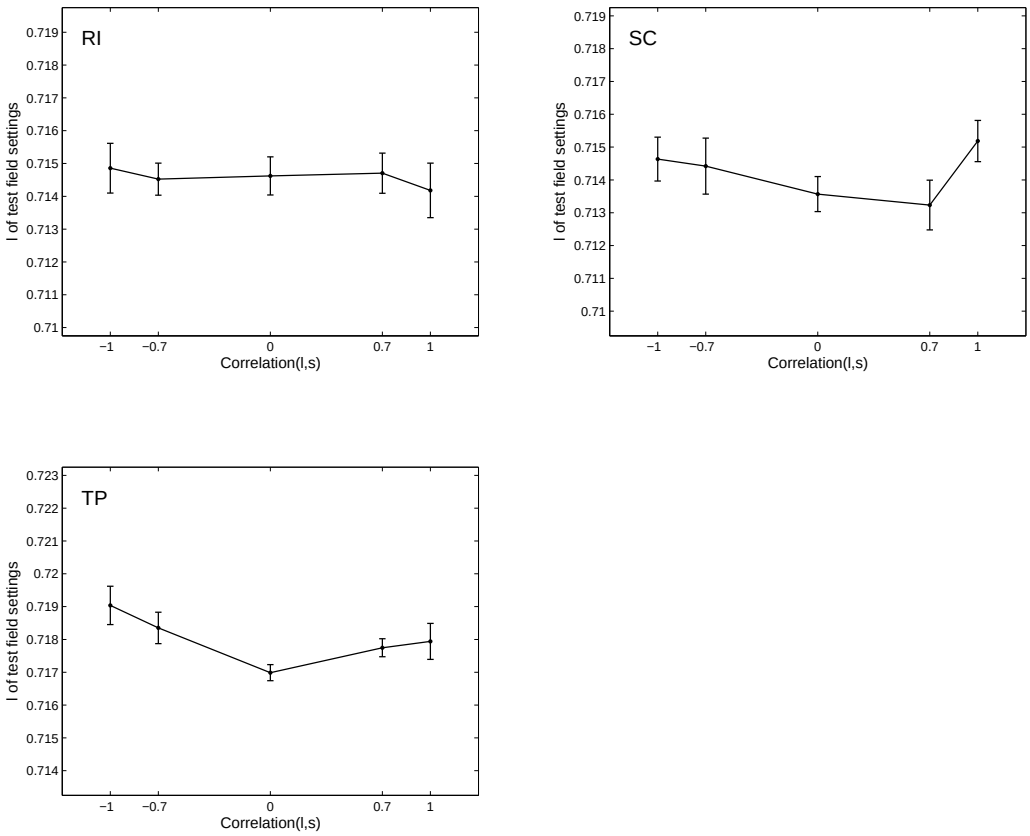

Figure 5.12: (Continuation)

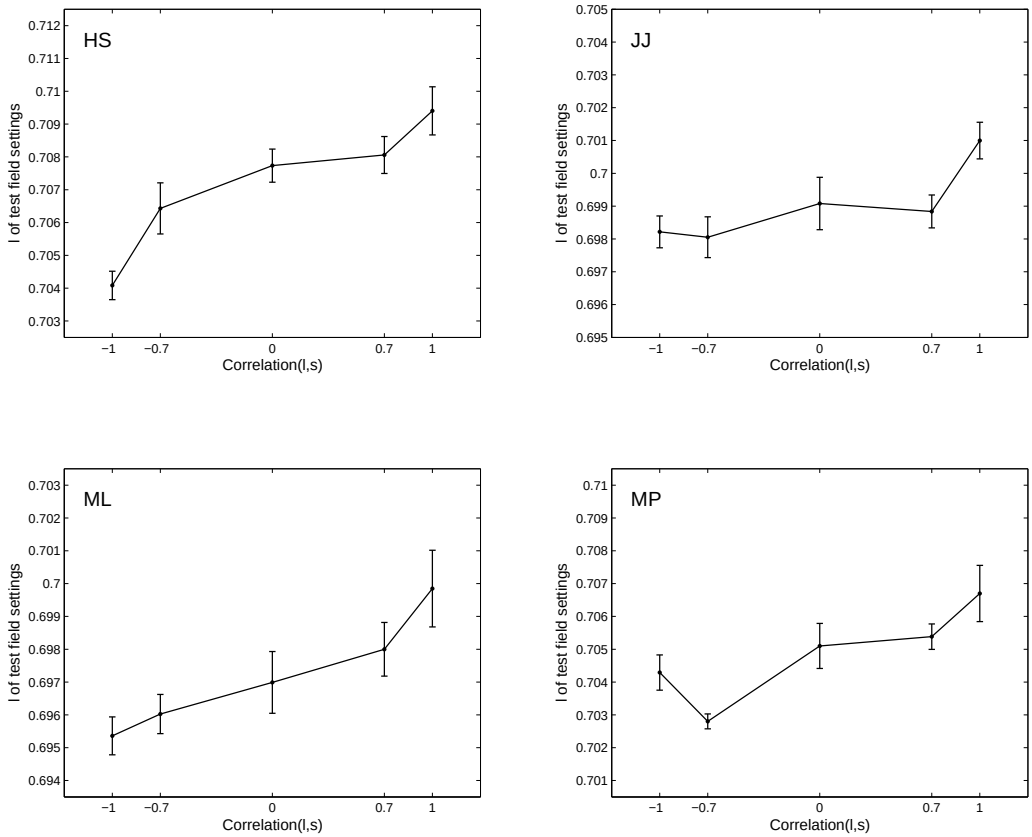

Figure 5.13: Results of experiment 7 for yellow stimuli.

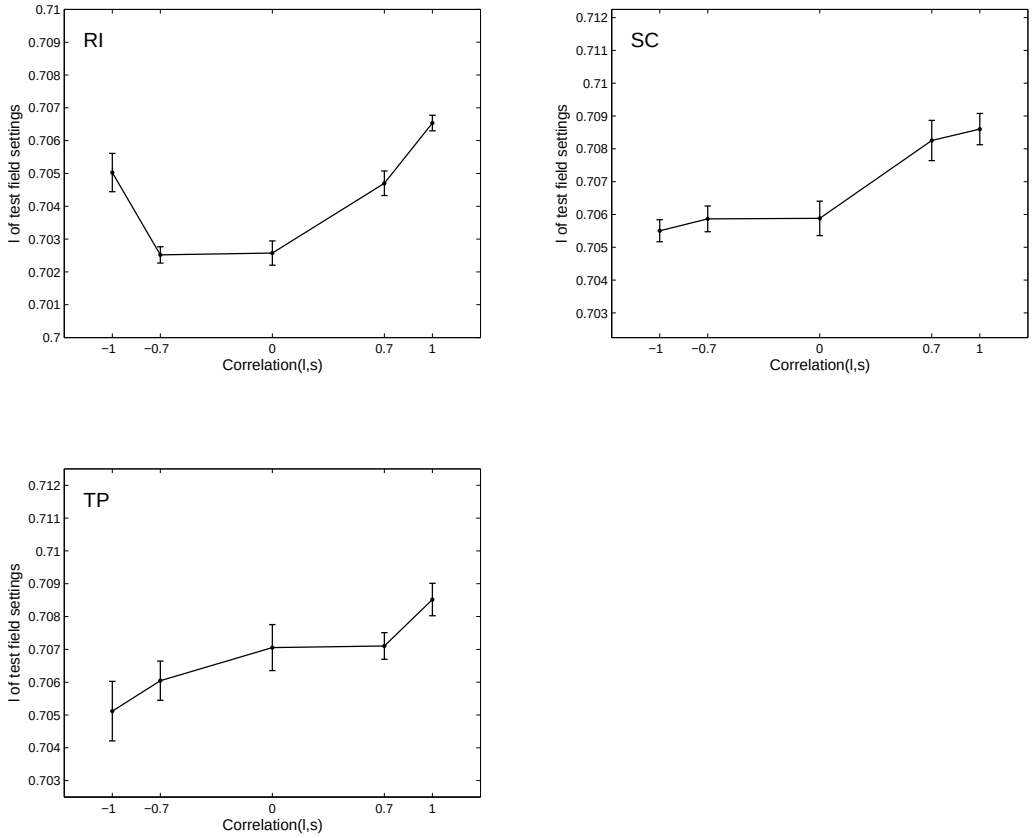

Figure 5.13: (Continuation)

5.8.4 Discussion

The results of experiment 7 together with the results of experiment 6 suggest that the degree to which the visual system takes into account the redness-blueness correlation depends on the mean chromaticity of the retinal image. In experiment 7, a substantial and robust effect of this statistic was found for green as well as for yellow surrounds, while no substantial effect was found for blue as well as for red surrounds. Experiment 6, in which a grey surround was used, yielded mixed results that can be regarded as intermediate between the two patterns of results found in experiment 7. The reason for this dependence of the effect of the redness-blueness correlation on the mean image chromaticity is not obvious, neither from the analysis of the Gaussian World in chapter 3 nor from the real world simulation in chapter 4, and thus remains on open question. Further research in order to address this question is desirable. A first step in this direction could be to run experiments with a larger sample of mean surround chromaticities in order to obtain a more thorough picture of the relation between the mean image chromaticity and the redness-blueness correlation effect. This should then be

compared with an analysis of the natural chromatic environment that addresses the relation of the redness-blueness correlation and the mean retinal image chromaticity with regard to the resulting diagnostic value of this higher order scene statistic. To this end, one could e.g. extend the real world simulation of chapter 4 to include illuminants with chromaticities outside of the loci of typical unimpeded daylights. Of special interest would be lights that occur in the natural environment but have more narrow bandwidths than the daylights used in chapter 4, e.g. lights in forests filtered by leaves.

Chapter 6

Discussion

In this final chapter, I will summarize in a brief review the results of the theoretical and empirical work presented in the previous three chapters and discuss how these results relate to each other. Some remarks about remaining open questions and further research that, in my opinion, is desirable will end this chapter.

6.1 The visual system and the world(s)

In this work, the role of chromatic scene statistics for the perception of surface and illumination colour is addressed by asking the following two questions:

1. Which statistics of the chromatic distribution in retinal images of natural scenes depend systematically on the illumination and are thus potential cues for the illumination that the visual system could exploit?
2. Which of these potential cues are taken into account for the perception of surface and illumination colour by the human visual system?

The first question refers to the chromatic relation between distal surfaces and illuminants on the one hand and the proximal stimulus one the other hand. Thus, it is part of a computational analysis in the sense of Marr (1982) and could be regarded as belonging to the realm of ecological optics. I have tried to tackle this question by investigating the regularities of the chromatic shifts of surfaces under changes of illumination in an idealized world that is mathematically tractable, the Gaussian World, in chapter 3 as well as in a simulation of the real world

in chapter 4. Higher order statistics of the chromatic distribution in the retinal image that could catch these regularities have then been tested with respect to their usefulness for estimating the illuminant by analysing the simulated real world — additionally to the previous analysis on the level of individual pixels of single scenes — on the level of scene statistics for the complete set of scenes. Three chromatic scene statistics have been identified as potential cues for the illuminant:

- the luminance-redness correlation,
- the luminance-blueness correlation,
- and the redness-blueness correlation.

Finally, in chapter 5, I have addressed for these three statistics the second above-mentioned question concerning the utilization of potential cues by the visual system by means of psychophysical experiments.

These two questions are, though distinct, tightly related in the following sense: if a particular scene statistic that has been shown to be diagnostic for the illumination in the theoretical analysis is experimentally found to be taken into account by the human visual system in the expected way, then it is likely that this experimental result is due to an adaption of the human visual system to the way in which in the natural environment this statistic depends on the illumination.

But such a case with clear-cut and consistent results for all parts of the theoretical and experimental analysis is given only for one of the three potential cues: the luminance-redness correlation. For the other two statistics the results have been less consistent and the overall conclusions are less certain decidable, as the following concluding summaries for each of the three potential cues will show.

6.1.1 Summary for the luminance-redness correlation

All three parts of the theoretical analysis propose that the correlation between luminance and redness in the retinal image is helpful for estimating the illuminant colour:

1. The algebraic solution for cone excitations in the Gaussian World predicts that surfaces similar to the illuminant in chromaticity are rendered as relatively lighter in the retinal image than surfaces dissimilar from the illuminant in chromaticity (cf. page 29). When the lighting gets e.g. more reddish, the reddish surfaces therefore should get lighter, relative to e.g. greyish or greenish surfaces.

2. This illumination-dependent regularity has been found to hold for the pixels of natural scenes in the real world simulation (cf. page 41).

3. The analysis of the natural scenes on the level of chromatic scene statistics has shown that this illumination-dependent regularity makes the luminance-redness correlation a useful cue for the illumination, even though this illumination-dependent regularity interacts with a second regularity, such that the correlation value of a given scene is almost unaffected by changes in illumination. Nevertheless, in the real world simulation this statistic is diagnostic for the illuminant, in that, for instance, a high luminance-redness correlation indicates a reddish illumination (cf. page 49), much like the consideration of the schematic example for the use of higher order scene statistics in Plate I(c).

In accordance with these theoretical findings, the experimental work presented in chapter 5 has yielded strong evidence that the human visual system uses the luminance-redness correlation for estimating the illuminant:

1. The results for all of the eleven subjects of experiment 1 show the expected dependence of the grey settings on the luminance-redness correlation within the surround of variegated stimuli. These effects are individually significant and substantial for all of the eleven subjects tested. These results confirm the findings of an earlier experiment that I have carried out in cooperation with Donald I. A. MacLeod (Golz & MacLeod, 2002).

2. Experiment 2 extends the basic effect of the luminance-redness correlation found for stimuli with neutral mean surround chromaticities in experiment 1 to surrounds with non-neutral mean chromaticities. The obtained effects are individually significant and substantial for all of the four surround chromaticities and for all of the four subjects tested in experiment 2.

3. Experiment 3 shows that this effect holds even if (instead of the non-luminance-weighted mean chromaticity within the surround) the luminance-weighted mean chromaticity is kept constant for all conditions of luminance-redness correlation. The effects obtained in this experiment were individually significant and substantial for all of the eight subjects tested.

In summary, the robust and substantial effects of the luminance-redness correlation found in these experiments are in best agreement with the predictions of the theoretical analyses and indicate that the human visual system exploits this chromatic scene statistic as a cue for the perception of surface and illumination colour.

6.1.2 Summary for the luminance-blueness correlation

The theoretical evidence that the correlation between luminance and blueness in the retinal image is helpful for estimating the illuminant colour is more mixed than in the case of luminance-redness correlation in that only two of the three parts of the theoretical analysis yield positive results:

1. Similar to the illumination-dependent regularity underlying the luminance-redness correlation, the Gaussian World analysis predicts that the bluish surfaces increase in luminance compared to e.g. greyish or yellowish surfaces when the illumination becomes bluish (cf. page 29).
2. This regularity has been found to hold for all scenes in the analysis of the real world simulation on the level of individual pixels (cf. page 41).
3. But surprisingly, in the analysis of the real world simulation on the level of scene statistics the correlation between luminance and blueness within a scene is almost independent of illumination (cf. page 51).

With respect to this mixed theoretical findings it is very interesting that experiment 4 and 5 yield strong evidence that the human visual system takes into account the luminance-blueness correlation for estimating the illuminant:

1. The results for all of the six subjects tested in experiment 4 show the expected dependence of the grey settings on the luminance-blueness correlation in the surround. These effects are individually significant and substantial for all of the six subjects.
2. The same six subjects show also significant and substantial effects in experiment 5, where (instead of the non-luminance-weighted mean chromaticity within the surround) the luminance-weighted mean chromaticity is kept constant for all conditions of luminance-blueness correlation.

Because of the agreement of the results of two of the three parts of the theoretical analysis on the one hand and the experiments on the other hand, one might be tempted to conclude that the human visual system takes into account the luminance-blueness correlation because it thereby can exploit a diagnostic value of this statistic for estimating the illuminant. But the question remains, why the third part of the theoretical analysis, i.e. the analysis of the real world simulation on the level of scene statistics, has not predicted such a diagnostic value for this statistic. It could of course be the case that the scenes used in the real

world simulation are just not representative in this respect for the chromatic environment to which the human visual system has adapted. But this point clearly desires further research and until the reason for the discrepancy concerning the findings for the luminance-blueness correlation has not been found, no final conclusion can be drawn.

6.1.3 Summary for the redness-blueness correlation

While the potential usefulness for estimating the illuminant of the two correlation statistics above result from a chromaticity-dependent increase in intensity for surfaces chromatically similar to the illuminant, rests the diagnostic value of the correlation between redness and blueness on a illumination-dependent regularity that is mediated by differences in surface reflectance bandwidth. This regularity on the level of individual surfaces is clearly observed in the first two parts of the theoretical analysis, but it is less clear whether the results for the redness-blueness correlation in the third part of the theoretical analysis indicate that this statistic is useful for estimating the illuminant:

1. The Gaussian World proposes (cf. page 32) that the light reflected from surfaces with narrow reflectance bandwidths will undergo a smaller chromaticity shift under changes of illumination than the light from surfaces with broader reflectance spectra.

2. Accordingly is the illumination-induced shift in the real world simulation greater for bluish and greyish surfaces than for yellowish surfaces which are more shift resistant due to their smaller bandwidth (cf. page 42).

3. This leads, in the analysis of the real world simulation on the level of scene statistics, to an increase in the redness-blueness correlation when the illumination changes toward the more reddish 4000 K daylight. Unfortunately, this statistic depends also systematically on the (illumination-independent) scene-averaged surface redness in a way that weakens the usefulness of this statistic as a cue to the illumination: the resulting constellation of this statistic to the mean image redness (Fig. 4.13) resembles to some degree the schematic constellation (d) in Plate I, where two cues are correlated both within and across illumination clusters such that the combination of the two cues is only as useful as each cue is alone (cf. section 2.3.2). On the basis of the real world simulation it thus remains an open question whether the redness-blueness correlation is helpful in the natural environment for estimating the illuminant.

The experimental evidence that the redness-luminance correlation within the retinal image is taken into account by the human visual system for the perception of surface and illumination colour is more mixed than the very clear-cut experimental results for both of the previous two scene statistics:

1. The expected influence of the redness-blueness correlation on the grey settings was only found for four of the seven subjects tested in experiment 6. And for three of these four subjects with significant results, the effects are rather weak.

2. The effect of the redness-blueness correlation differs substantially for the surrounds with different mean chromaticities used in experiment 7. For green as well as for yellow surrounds an individually significant and substantial effect was found for all of the seven subjects while no substantial effect was found for blue as well as for red surrounds.

The results for the grey surrounds used in experiment 6 can be regarded as intermediate between the two patterns of results found in experiment 7 (especially since the same seven subjects participated in both experiments). It thus seems as if the effect of the redness-blueness correlation depends on the mean image chromaticity. In order to clarify the exact nature of and the reason for this relation, further research is desirable. On the experimental side a larger sample of mean surround chromaticities should be tested in order to obtain a more thorough picture of the relation between the mean image chromaticity and the redness-blueness correlation effect. On the theoretical side the question of the diagnostic value of the redness-blueness correlation and its dependence on the mean retinal image chromaticity should be addressed, e.g. by extending the real world simulation of chapter 4 to include illuminants with chromaticities outside of the loci of typical unimpeded daylights. Of special interest would be lights that occur in the natural environment but have more narrow bandwidths than the daylights used in chapter 4, e.g. lights in forests filtered by leaves.

6.2 Open questions

Two questions that remain open with respect to the luminance-blueness correlation and the redness-blueness correlation have been already mentioned in the above concluding summaries for these statistics. Here, I will go beyond the evaluation of the results of chapter 3, 4, and 5 and will discuss the relation of the presented research to other issues that have not been addressed directly in this work.

In chapter 3 and 4, I have presented several regularities of the chromatic shifts of surfaces under changes of illumination in order to derive potential cues that the visual system could exploit for estimating the illuminant. One of the first questions that follow is whether the visual system — additionally to using a particular cue for solving the estimation problem (cf. section 2.3) — does also take into account the underlying regularity for determining the perceived colour of individual surfaces, i.e. to solve the compensation problem (cf. section 2.4). For instance, the analysis of the Gaussian World as well as the analysis of the real world simulation have shown that surfaces similar to the illuminant in chromaticity are rendered as lighter than surfaces dissimilar from the illuminant in chromaticity: when the lighting gets reddish, the reddish surfaces should get higher luminances, relative to e.g. greyish or greenish surfaces. Does the visual system compensate for these different luminance shifts for surfaces of different chromaticities? To this end, the visual system could determine the perceived lightness of each surface by taking into account the retinal information corresponding to that surface and the estimated illuminant colour in a non-normalization-compatible way by assigning, for instance, a lower perceived lightness to a reddish surface than to an isoluminant greenish surface, if the estimated illuminant is reddish.

Whether the visual system indeed improves the constancy of perceived lightness of surfaces by taking into account this regularity of chromaticity-dependent luminance shifts in the retinal image is not evident from the extensive literature on colour and lightness matching. The same question has also not yet been addressed for the other regularities of chromatic shifts of surfaces under changes of illumination presented in chapter 3 and 4. Further research addressing these questions is thus desirable.

In order to make a phenomenon like colour perception — that depends in a complex way on numerous factors — accessible to scientific investigation, one has to restrict oneself to consider only a few of these factors in a particular study. The restrictions that one is willing or forced to make will depend among other things on the aim and the assumptions of the perspective in which the investigation is embedded, the hitherto accumulated knowledge about the phenomenon, and the methods that are available in this field of research. The restrictions that I have made in this work is to consider only the information available in the chromatic distribution of the retinal images that result from scenes with co-planar diffusely reflecting surfaces that are lit by a single uniform illuminant. This leaves aside many other factors that could play a role for the perception of surface and illumination colour.

For instance, by moving from a two-dimensional world of co-planar surfaces to a three-dimensional world with surfaces of different orientations, one can investigate cues to the illuminant that depend on the three-dimensional geometry of the

scene like e.g. mutual reflections (Funt et al., 1991; Funt & Drew, 1993; Bloj et al., 1999). One further step would be to go beyond diffusely reflecting surfaces and consider the bi-directional reflectance density function (see e.g. Maloney, 2003), which describes the way in which the spectrum of a light reflected from a surface in a certain direction depends on the direction from which light is arriving at the surface. The consideration of this type of information has lead to the investigation of another cue to the illuminant, namely specular reflections (D'Zmura & Lennie, 1986; Lee, 1986; Tominaga & Wandell, 1989; Maloney & Yang, 2003), as well as of the perception of material properties (see e.g. Adelson, 2001). Abandoning finally the restriction to a single homogenous illumination leads, e.g., to the question how the visual system determines for different regions of the scene the illuminants that it assumes to be effective in these regions (see e.g. Adelson, 2000).

The question how the results presented in this work fit into and interact with these further factors of the more complex three-dimensional world is an open question and research to address this question is currently facing some major challenges. One the one hand we still lack a theoretical model of our natural environment that integrates all the geometrical and spectral idealizations in order to allow predictions of cues for the perception of surface and illumination colour in the complex three-dimensional world (similar to what the Gaussian World allowed for the much simpler two-dimensional world). On the other hand, a theoretical framework that allows to describes in a unified way the internal architecture and processes underlying the perception of surfaces and illumination in all of these worlds of different complexity has only recently started to emerge (Mausfeld, 2003). This framework assumes for the arsenal of theoretical building blocks for instance internal representations for illuminants as well as for surfaces. In such a framework, the interdependencies of colour with geometrical aspects of the perceived scene as well as other factors can be formulated by an explanatory vocabulary that is more natural and unified than ad-hoc mechanisms that have been tailored for specific types of stimuli, as e.g. the context effects often invoked by colour scientists. The way in which for instance the internal parameters for colour in these representations depend on (and exert influence on) other parameters of these as well as other internal representations has yet to be fully explored.

Bibliography

Adelson, E. (2001). On seeing stuff: the perception of materials by humans and machines. In B. E. Rogowitz & T. N. Pappas (Eds.), *Proceedings of the SPIE, Vol. 4299: Human Vision and Electronic Imaging VI* (p. 1-12). The International Society for Optical Engineering SPIE.

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

Andres, J. (1996). Das Allgemeine Lineare Modell. In E. Erdfelder, R. Mausfeld, I. Meiser, & G. Rudinger (Eds.), *Handbuch Quantitative Methoden* (p. 185-200). Weinheim: PVU-Beltz.

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

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

Bloj, M., Kersten, D., & Hurlbert, A. C. (1999). Perception of three-dimensional shape influences colour perception through mutual illumination. *Nature*, *402*, 877-879.

Brill, M. H. (1978). A device performing illuminant-invariant assessment of chromatic relations. *Journal of Theoretical Biology*, *71*, 473-478.

Brown, R. O. (2003). Backgrounds and illuminants: the yin and yang of color constancy. In R. Mausfeld & D. Heyer (Eds.), *Colour perception: Mind and the physical world* (p. 247-272). Oxford: Oxford University Press.

Buchsbaum, G. (1980). A spatial processor model for object colour perception. *Journal of The Franklin Institute*, *310*, 1-26.

Cornelissen, F. W., & Brenner, E. (1995). Simultaneous colour constancy revisited: an analysis of viewing strategies. *Vision Research*, *35*(17), 2431-2448.

Dannemiller, J. L. (1993). Rank orderings of photoreceptor photon catches from natural objects are nearly illuminant-invariant. *Vision Research*, *33*, 131-140.

Delahunt, P. B., & Brainard, D. H. (2004). Does human color constancy in-

corporate statistical regularities of natural daylight? *Journal of Vision*, *4*, 57-81.

D'Zmura, M., & Lennie, P. (1986). Mechanisms of color constancy. *Journal of the Optical Society of America A*, *3*, 1662-1672.

Eisner, A., & MacLeod, D. I. A. (1980). Blue-sensitive cones do not contribute to luminance. *Journal of the Optical Society of America*, *70*, 121-122.

Endler, J. A. (1993). The color of light in forests and its implications. *Ecological Monographs*, *63*, 1-27.

Foster, D. H., & Nascimento, S. M. (1994). Relational colour constancy from invariant cone-excitation ratios. *Proceedings of the Royal Society of London. Series B: Biological Sciences*, *257*(1349), 115-21.

Funt, B. V., & Drew, M. S. (1993). Color space analysis of mutual illumination. *IEEE Trans. Pattern Anal. Mach. Intell.*, *15*, 1319-1326.

Funt, B. V., Drew, M. S., & Ho, J. (1991). Color constancy from mutual reflection. *International Journal of Computer Vision*, *6*, 5-24.

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. L., & Ramachandran, V. (1992). Red rooms in white light appear different from white rooms in red light. *Proceeding of the National Academy of Sciences USA*, *96*, 307-312.

Golz, J. (2004). Influence of subjects viewing behavior on grey settings. *Perception*, *33 (Supplement)*, 59.

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

Helmholtz, H. von. (1896). *Handbuch der Physiologischen Optik* (2nd ed.). Hamburg: Voss.

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

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

Koenderink, J. J., & van Doorn, A. J. (2003). Perspectives on color space. In R. Mausfeld & D. Heyer (Eds.), *Colour perception: Mind and the physical world* (p. 1-56). Oxford: Oxford University Press.

Koffka, K. (1932). Some remarks on the theory of colour constancy. *Psychologische Forschung*, *16*, 329-354.

Kraft, J. M., & Brainard, D. H. (1999). Mechanisms of color constancy under nearly natural viewing. *Invest. Ophthalmol. Vis. Sci.*, *33*, 756.

Lee, H.-C. (1986). Method for computing the scene-illuminant chromaticity from specular highlights. *Journal of the Optical Society of America A*, *3*, 1694-1699.

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

MacLeod, D. I. A. (1985). Receptoral constraints on color appearance. In D. Otterson & S. Zeki (Eds.), *Central and peripheral mechanisms of color vision.* London: MacMillan.

MacLeod, D. I. A., & Boynton, R. M. (1979). Chromaticity diagram showing cone excitation by stimuli of equal luminance. *Journal of the Optical Society of America*, *69*(8), 1183-6.

MacLeod, D. I. A., & Golz, J. (2003). A computational analysis of colour constancy. In R. Mausfeld & D. Heyer (Eds.), *Colour perception: Mind and the physical world* (p. 204-242). Oxford: Oxford University Press.

MacLeod, D. I. A., & von der Twer, T. (2003). The pleistochrome: optimal opponent codes for natural colours. In R. Mausfeld & D. Heyer (Eds.), *Colour perception: Mind and the physical world* (p. 155-184). Oxford: Oxford University Press.

Maloney, L. T. (1986). Evaluation of linear models of surface spectral reflectance with small numbers of parameters. *Journal of the Optical Society of America A*, *3*, 1673-1683.

Maloney, L. T. (2003). Surface colour perception and environmental constraints. In R. Mausfeld & D. Heyer (Eds.), *Colour perception: Mind and the physical world* (p. 279-300). Oxford: Oxford University Press.

Maloney, T. M., & Yang, J. N. (2003). The illumination estimation hypothesis and surface color perception. In R. Mausfeld & D. Heyer (Eds.), *Colour perception: Mind and the physical world* (p. 335-358). Oxford: Oxford University Press.

Marr, D. (1982). *Vision.* New York: W. H. Freeman and Company.

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. (2002). The physicalistic trap in perception theory. In D. Heyer & R. Mausfeld (Eds.), *Perception and the physical world* (p. 75-112). Chichester: Wiley.

Mausfeld, R. (2003). '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. 382-430). Oxford: Oxford University 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), 209-24.

Moon, P., & Spencer, D. E. (1945). Polynomial representation of spectral curves. *Journal of the Optical Society of America*, *35*, 597-600.

Munsell Color Company. (1976). *Munsell book of color.* Baltimore: Munsell Color.

Páraga, C. A., Brelstaff, G., Troscianko, T., & Moorehead, I. R. (1998). Color and luminance information in natural scenes. *Journal of the Optical Society of America A*, *15*, 563-569.

Ruderman, D. L., Cronin, T. W., & Chiao, C. C. (1998). Statistics of cone responses to natural images: implications for visual coding. *Journal of the Optical Society of America A*, *15*, 2036-2045.

Sällström, P. (1973). *Colour and physics: Some remarks concerning the physical aspects of human colour vision.* Report No. 73-09, Institute of Physics, University of Stockholm, Stockholm.

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.

Tominaga, S., & Wandell, B. A. (1989). Standard surface-reflectance model and illuminant estimation. *Journal of the Optical Society of America A*, *6*, 576-584.

Ward, J. H., & Jennings, E. (1973). *Introduction to linear models.* Englewood Cliffs, N.J.: Prentice-Hall.

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. 65-110). Oxford: Oxford University Press.

Webster, M. A., & Mollon, J. D. (1997). Adaptation an the color statistics of natural images. *Vision Research*, *37*, 3283-3298.

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

Worthey, J. A., & Brill, M. H. (1986). Heuristic analysis of Von Kries color constancy. *Journal of the Optical Society of America A*, *3*, 1708-12.

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

Appendix A

Statistical Test

In this appendix I will describe the statistical test that is used in chapter 5 to evaluate the results of all seven experiments. This test analyses the linear dependence of the test field settings (dependent variable) on the experimentally manipulated chromatic correlation within the surround (independent variable). For this, a particular model within the framework of the General Linear Model will be specified and used. As we will see, the parametric function that will be adopted can also be regarded as a contrast within the framework of the one-factorial ANOVA. A full description of the General Linear Model is beyond the scope of this work. For an introduction to this topic see e.g. Ward and Jennings (1973) or Andres (1996).

In order to introduce the specific way in which the relation between independent and dependent variable will be analysed, let us consider the following theoretical situation: in an experiment with J conditions, each condition j is presented n_j times and the total number of trials is $N = \sum_{j=1}^{J} n_j$. The value of the independent variable X for condition j is denoted by x_j and the expected value of the dependent variable Y for condition j is denoted by μ_j. Imagine that the expected values for all conditions were known, than the following procedure is suitable to address the question whether the dependent variable increases when the independent variable increases: fit a Gaussian regression line to the expected values and test whether the slope of that regression line is greater than zero; in doing so, weight each expected value μ_j by the corresponding number of trials n_j, such that as data for each condition μ_j enters n_j times into the linear regression.

The slope of a linear regression line for the usual situation where Y is to be predicted by X is:

$$b = \frac{Cov(X,Y)}{S_X^2} = \frac{\frac{1}{N}\sum_{i=1}^{N}(x_i - M_X)(y_i - M_Y)}{S_X^2}$$

Correspondingly, for the present situation where μ is to be predicted the slope is:

$$b = \frac{\frac{1}{N}\sum_{i=1}^{N}(x_i - M_X)(\mu_i - M_\mu)}{S_X^2}$$

Instead of summing up all N data points in a single cycle, we will sum up condition by condition:

$$b = \frac{\frac{1}{N}\sum_{j=1}^{J}\sum_{k=1}^{n_j}(x_{jk} - M_X)(\mu_{jk} - M_\mu)}{S_X^2}$$

For each condition j the n_j data points are all identical, namely (x_j, μ_j). Thus we can write and then rearrange the numerator as follows:

$$\begin{aligned} b &= \frac{\frac{1}{N}\sum_{j=1}^{J} n_j(x_j - M_X)(\mu_j - M_\mu)}{S_X^2} \\ &= \frac{\frac{1}{N}\sum_{j=1}^{J} n_j(x_j - M_X)\mu_j + \frac{1}{N}\sum_{j=1}^{J} n_j(x_j - M_X)(-M_\mu)}{S_X^2} \end{aligned}$$

The second sum in the numerator yields zero, so we obtain for the slope b:

$$b = \frac{\frac{1}{N}\sum_{j=1}^{J} n_j(x_j - M_X)\mu_j}{S_X^2}$$

A slight rearrangement makes obvious that this equation can be regarded as a parametric function within the framework of the General Linear Model. It can also be viewed as a contrast within the framework of the one-factorial ANOVA. We will denote this parametric function by ψ:

$$\psi := \sum_{j=1}^{J} \frac{n_j(x_j - M_X)}{NS_X^2}\mu_j$$

By means of this parametric function, we can now formulate the above-mentioned question whether the slope of the regression line is greater than zero by the following pair of hypotheses:

$$\begin{aligned} H_0 &: \psi = 0 \\ H_1 &: \psi > 0 \end{aligned}$$

Or as will be done in below, one can treat the more general question whether the slope is greater than a real number ψ_0:

$$\begin{aligned} H_0 &: \psi = \psi_0 \\ H_1 &: \psi > \psi_0 \end{aligned}$$

Of course, we are not in the lucky situation to know the μ_j and thus ψ, so that we are in need for a statistical test that allows to decide about the above hypotheses. Such a test that is provided by the General Linear Model will be described next. First, we need to introduce some ingredients for the test statistic that will be used.

Let $\boldsymbol{\beta}$ be the $(J \times 1)$ vector with μ_j as components and let $\mathbf{c}$ be the $(J \times 1)$ vector with the coefficients for μ_j in the parametric function ψ as components:

$$\boldsymbol{\beta} = \begin{pmatrix} \mu_1 \\ \vdots \\ \mu_J \end{pmatrix}, \ \mathbf{c} = \begin{pmatrix} \frac{n_1(x_1 - M_X)}{NS_X^2} \\ \vdots \\ \frac{n_J(x_J - M_X)}{NS_X^2} \end{pmatrix}$$

Then, we can write ψ as:

$$\psi = \mathbf{c}'\boldsymbol{\beta} = \sum_{j=1}^{J} c_j \mu_j$$

Let $\mathbf{y}$ be the $(N \times 1)$ vector containing the measurements of the dependent variable for one subject in the following order: first the n_1 measurements for condition 1, then the n_2 measurements for condition 2 and so forth until condition J.

$$\mathbf{y} = \begin{pmatrix} y_1 \\ \vdots \\ y_N \end{pmatrix}$$

The design matrix $\mathbf{X}$ that specifies the particular model that is considered within the framework of the General Linear Model, is in our present case a $(N \times J)$ matrix with zeros everywhere except ones in the first column for the first n_1 rows, then in the second column for the next n_2 rows and so forth:

$$\mathbf{X} = \begin{pmatrix} 1 & 0 & \dots & 0 \\ \vdots & \vdots & \vdots & \vdots \\ 1 & 0 & \dots & 0 \\ 0 & 1 & \dots & 0 \\ \vdots & \vdots & \vdots & \vdots \\ 0 & 1 & \dots & 0 \\ \vdots & \vdots & \vdots & \vdots \\ \vdots & \vdots & \vdots & \vdots \\ 0 & 0 & \dots & 1 \\ \vdots & \vdots & \vdots & \vdots \\ 0 & 0 & \dots & 1 \end{pmatrix} \begin{matrix} \text{row } 1 \\ \vdots \\ \\ \\ \vdots \\ \\ \vdots \\ \vdots \\ \\ \vdots \\ \text{row } N \end{matrix} \begin{matrix} \left.\vphantom{\begin{matrix} a \\ b \\ c \end{matrix}}\right\} n_1 \text{ times} \\ \\ \left.\vphantom{\begin{matrix} a \\ b \end{matrix}}\right\} n_2 \text{ times} \\ \\ \\ \left.\vphantom{\begin{matrix} a \\ b \end{matrix}}\right\} n_J \text{ times} \end{matrix}$$

Since this design matrix has full rank, we can obtain the following Gauß-Markoff-estimators for $\boldsymbol{\beta}$ and ψ:

$$\hat{\boldsymbol{\beta}} = (\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'\mathbf{y}$$

$$\hat{\psi} = \mathbf{c}'\hat{\boldsymbol{\beta}} = \mathbf{c}'(\mathbf{X}'\mathbf{X})^{-1}\mathbf{X}'\mathbf{y} = (\mathbf{X}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{c})'\mathbf{y}$$

We will set $\mathbf{a} := \mathbf{X}(\mathbf{X}'\mathbf{X})^{-1}\mathbf{c}$. This vector allows to directly specify how to estimate ψ by means of $\mathbf{y}$:

$$\hat{\psi} = \mathbf{a}'\mathbf{y}$$

The norm of this vector, which will be needed below for the test statistic, can be obtained easily by:

$$\|\mathbf{a}\| = \sqrt{\mathbf{c}'(\mathbf{X}'\mathbf{X})^{-1}\mathbf{c}}$$

The last element needed for the test statistic is an unbiased estimator for the variance σ^2, which is obtained as follows (where $\hat{\mathbf{y}}$ is given by $\hat{\mathbf{y}} = \mathbf{X}\hat{\boldsymbol{\beta}}$):

$$s^2 = \frac{\|\mathbf{y}\|^2 - \|\hat{\mathbf{y}}\|^2}{N - J}$$

Then, we can use the following statistic t_{emp} to test whether the slope of the regression line is greater than ψ_0:

$$t_{emp} = \frac{\hat{\psi} - \psi_0}{\|\mathbf{a}\|\, s}$$

This test statistic is t-distributed with $N - J$ degrees of freedom and the non-centrality parameter $\delta = (\psi - \psi_0)/(\|\mathbf{a}\|\,\sigma)$. Thus, we can test the hypotheses

$$\begin{aligned} H_0 &: \psi = \psi_0 \\ H_1 &: \psi > \psi_0 \end{aligned}$$

by means of the test statistic t_{emp} on the significance level α as follows: the null hypothesis H_0 has to be rejected, if $t_{emp} > t_{N-J;\alpha}$ (where $t_{N-J;\alpha}$ denotes the upper α-fractile of the Student's t-distribution with $N - J$ degrees of freedom).

Appendix B

Viewing Instructions

Some of the details of the procedure of the seven experiments presented in chapter 5 (namely the viewing instructions and the total presentation time per trial) have been chosen by taking into account experiments that I have carried out to investigate the influence of viewing instructions on subjects grey seetings. Since these experiments have not been published in a journal yet, but only presented at the European Conference of Visual Perception 2004, I will give a brief overview of the main results in this appendix.

In the seven experiments of chapter 5 as well as in numerous other studies on colour constancy and colour induction the experimental task is to adjust a test field until it appears achromatic. Subject's viewing behaviour during these grey settings are often not controlled or reported. If it is left to the subjects to decide how much time they spend looking at the test field, increased variance or even artifacts might be the consequence. In some studies, subjects are explicitly instructed to visually explore the stimulus by looking around in the surround, but the influence of this difference in subjects viewing behaviour on the resulting grey settings has previously not been investigated. (For a related work with respect to results of matching experiments, see Cornelissen and Brenner (1995).)

B.1 Experiment I: basic effect of viewing instructions

Subjects had to adjust a central test field in a variegated surround to make it appear achromatic. The surround stimuli used equaled the stimuli of experiment 1 with zero luminance-redness correlation (cf. Plate IV), only that various different

mean surround chromaticities were used in order to measure the degree of colour constancy in these coloured surrounds (i.e. the shift of grey setting induced by the biased mean surround chromaticity). The question that was addressed is whether the degree of colour constancy in a given surround depends on the viewing instruction given to the subjects. In alternating trials, two viewing instructions ("Focus" vs. "Explore") were displayed shortly before each trial. These instructions were explained to the subjects before the experiment as follows:

> "Focus: Adjust the central test field until it looks grey to you. Keep looking at the test field, do not look around in the surround. If the test field changes its colour appearance while you are looking at it, adjust it such that it looks perfectly grey again. Keep looking at the test field and make it look grey until a beep informs you that you can end the trial by pressing the [ENTER] key."

> "Explore: Adjust the central test field until it looks grey to you. Then explore the scene by looking around in the surround for about ten seconds. To this end, move your gaze to different parts of the scene along a circle around the center halfway to the border of the display. Then look at the test field again to make sure it still looks perfectly grey. If it doesnt, adjust it until it does. Keep alternately looking around and then adjusting the test field until a beep informs you that you can end the trial by pressing the [ENTER] key."

Subjects could only accept their settings and end the trial after 60 seconds, indicated by the beep. This made the total time of adaptation to the stimulus equal for the two viewing conditions.

Results

In the "Explore" condition, the grey settings shifted subtiantially more towards the mean of the surround (higher degree of colour constancy, see Plate V(a)-(d)). The effect is most obvious for surrounds that differ in s-value from chromatically neutral ones, but is also present for surrounds with an offset on the l-axis (see Plate V(e)).

Note that, for surround *blue25*, subjects could adjust the test field only by modifying the s-value along a line through the surround chromaticity and a neutral chromaticity (7000 K). This task allowed satisfactory grey settings for all subjects and was used in all of the following experiments.

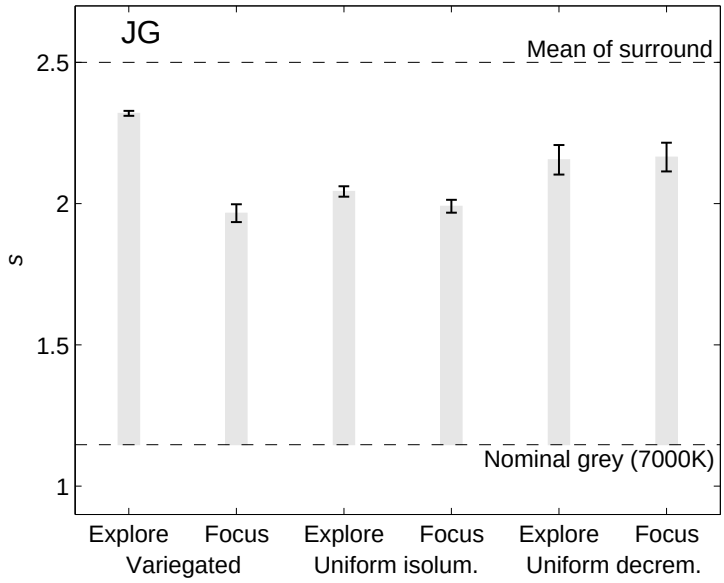

Figure B.1: Mean grey settings (±1 SEM) of experiment II for subject JG. An effect of viewing instructions is only obvious for the variegated surround.

B.2 Experiment II: uniform stimuli

The same instructions and methods as in experiment I were used, but now a uniform surround, whose chromaticity equalled the mean of the variegated surround, was also included. Since subjects could possibly produce lower shifts of the now isoluminant test field in the uniform surround because they avoid making the test field undistinguishable from the surround, another condition with uniform surrround but decremental test field was included.

Results

The effect of viewing instruction is strongly reduced or even absent when a uniform surround (with isoluminant as well as decremental test field) is used instead of a variegated one (cf. Fig. B.1).

B.3 Experiment III: local or global cause?

Among the possible mechanisms that could be causing the effect of viewing instructions are:

- In the "Focus" condition, the fovea only adapts to the test field (which at the beginning of each trial starts at a neutral chromaticity), whereas in the "Explore" condition, it also adapts to the coloured surround. The resulting difference in the foveal state of adaptation could explain the dependence of the grey settings on the viewing instructions.
- The output of mechanisms sensitive to spatiotemporal modulation in the periphery could decrease when the test field is fixated due to reduced modulation of the signals from the surround.

Experiments IIIa and IIIb aimed to explore these possible contributions to the effect of viewing instructions.

Experiment IIIa: varying the test field start chromaticity

The same instructions and methods as in experiment I were used. In order to investigate whether the effect of viewing instructions is caused by the initial chromaticity of the test field, two starting chromaticities were used:

- a neutral chromaticity of a 7000 K daylight ((l,s) = (0.6887, 1.1466)) and
- a saturated blue one ((l,s) = (0.6887, 3.5)).

Results

If differences in foveal state of adaptation were the only cause for the dependence of the grey settings on the viewing instructions, the settings in the "Focus" condition with blue test field start chromaticity should have s-values substantially higher than in the corresponding "Explore" condition (and even higher than s = 2.5, the mean of the surround). Fig. B.2 shows that this is clearly not the case. Thus, this local factor seems to contribute only partially to the observed effect of viewing instruction.

Experiment IIIb: varying the spatiotemporal modulation in the surround

The same instructions and methods as in experiment I were used, but the following two conditions were also included:

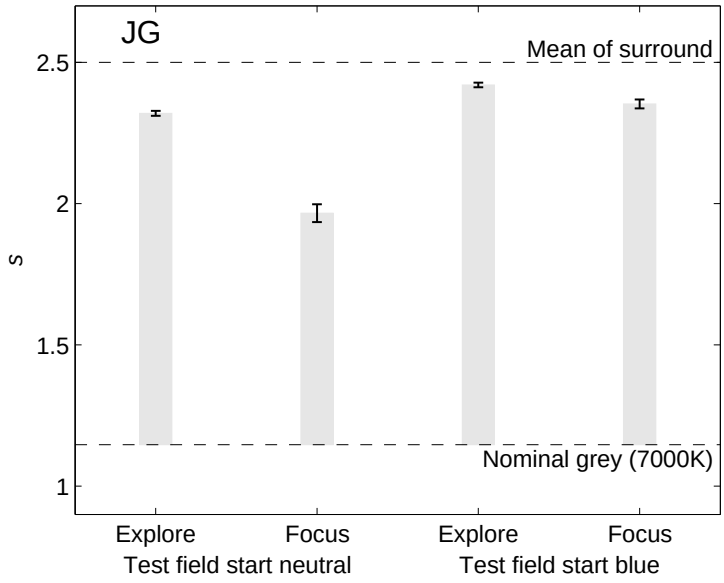

Figure B.2: Mean grey settings (±1 SEM) of experiment IIIa for subject JG. Letting the test field start in blue does not reverse the effect of viewing instruction, which suggests that differences in foveal state of adaptation contributes only partially to the effect.

- "Moving test field": the test field and the immediate context (a circular region cut out from the variegated surround adjacent to the test field) moved on the stationary background (see Fig. B.3(a)).
- "Moving surround": the test field and the immediate context remained stationary, but the surround moved (see Fig. B.3(b)).

In both conditions, subjects fixated a cross at the center of the test field and continuously adjusted the colour to look grey without looking around in the surround. Thus, concerning the foveal state of adaptation these two conditions are equal to the "Focus" condition and should yield comparable grey settings, if this local factor is the only cause of the effect of viewing instructions.

Results

As Fig. B.4 shows, the conditions "Moving test field" and "Moving surround" give rise to grey settings with s-values higher than in the "Focus" condition. This finding, too, suggests that differences in foveal state of adaptation contribute only partially to the observed effect of viewing instruction and indicates a more global mechanism.

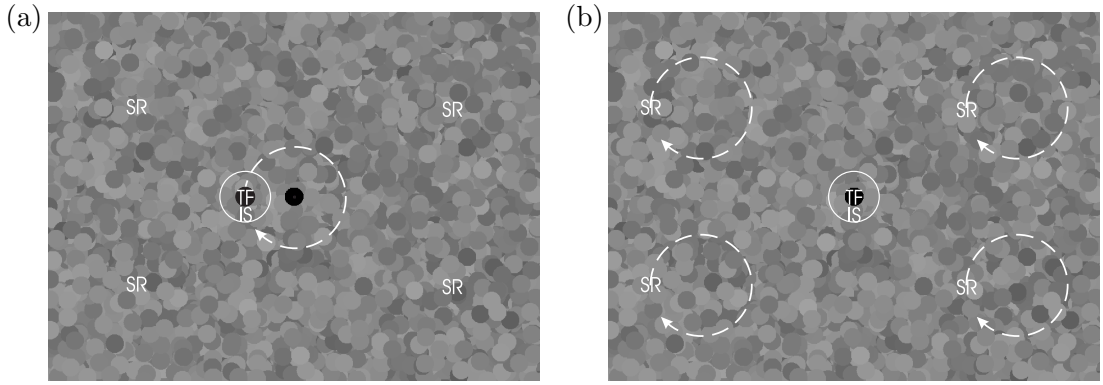

Figure B.3: Depiction of the stimulus regions and their movements in experiment IIIb. TF: test field, IS: immediate surround, SR: surround. (a) "Moving test field" condition: only the test field and its immediate surround moves. (b) "Moving surround" condition: the surround moves rigidly along a circular trajectory.

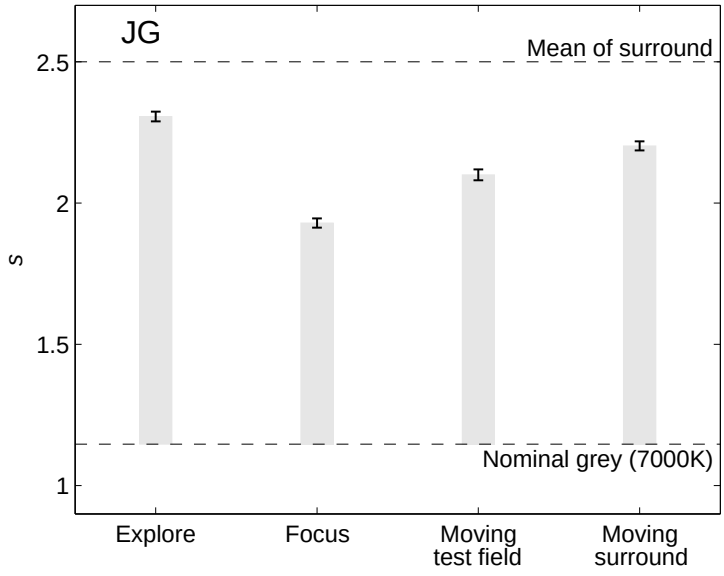

Figure B.4: Mean grey settings (±1 SEM) of experiment IIIb for subject JG. Introducing spatiotemporal modulation in the surround while keeping local factors constant yields higher s-values than the "Focus" condition. This suggests that differences in foveal state of adaptation contributes only partially to the effect of viewing instruction.

Appendix C

Supplementary data

C.1 Chromaticity distributions of individual pixels

The following figures show for all of the twelve natural scenes of Ruderman et al. (1998) the distribution of pixel chromaticities in the ($log(l)$,$log(s)$) chromaticity plane under two of the illuminants used in the simulation (cf. section 4.2.2): in each plot, the upper left cluster of data points represents the chromaticities of the respective scene under 20 000 K daylight and the lower right cluster the chromaticities under 4000 K daylight. The result for scene Park4 is discussed as an example in section 4.3.1.

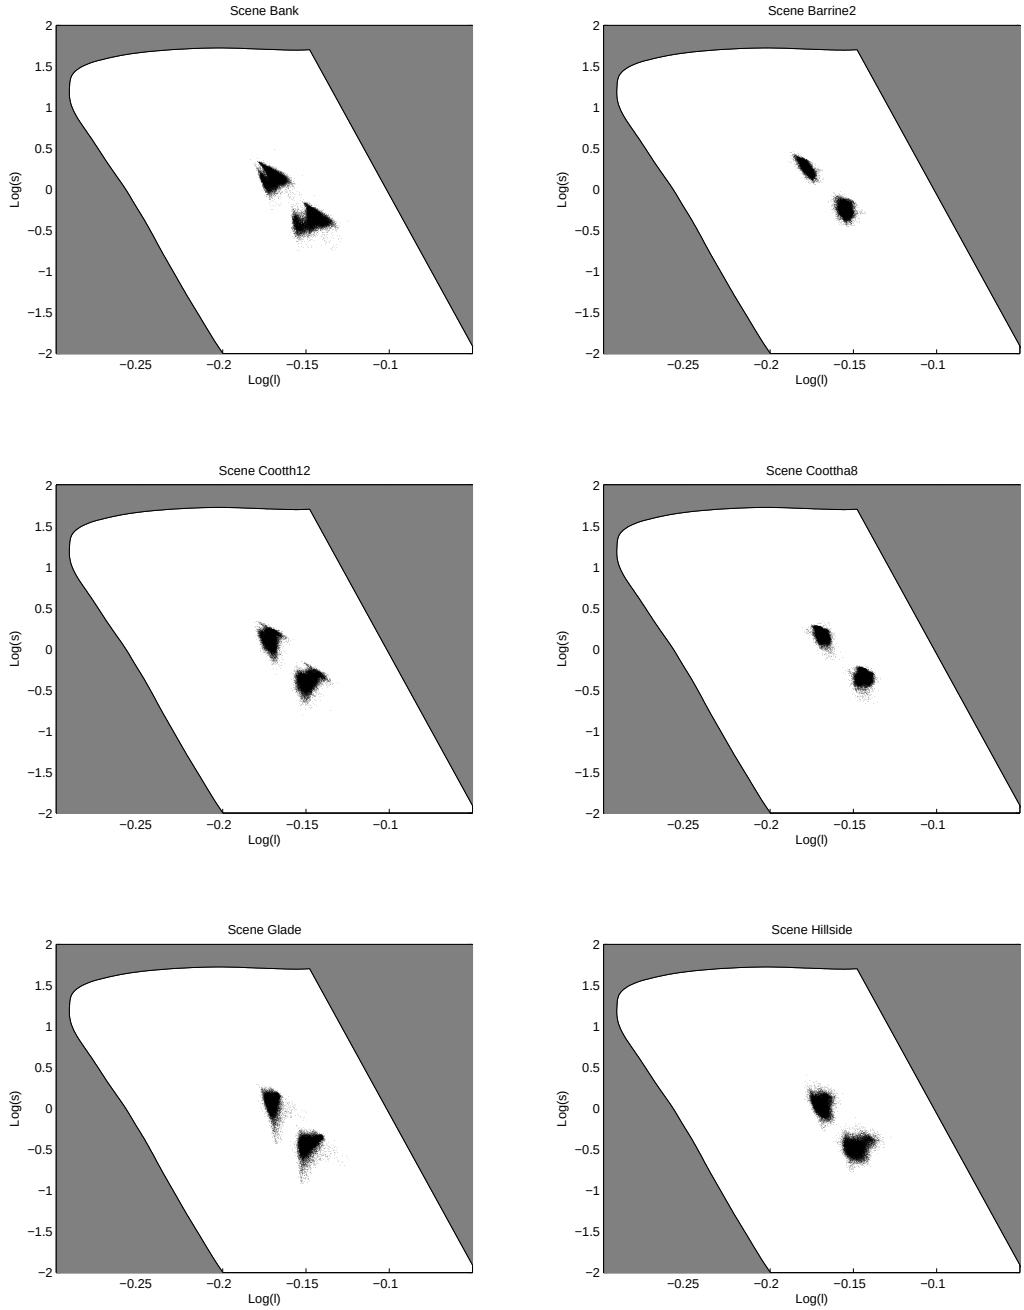

Figure C.1: Distribution of pixel chromaticities for the twelve natural scenes of Ruderman et al. (1998) under two illuminants: 20 000 K daylight (upper left cluster of data points) and 4000 K (lower right cluster).

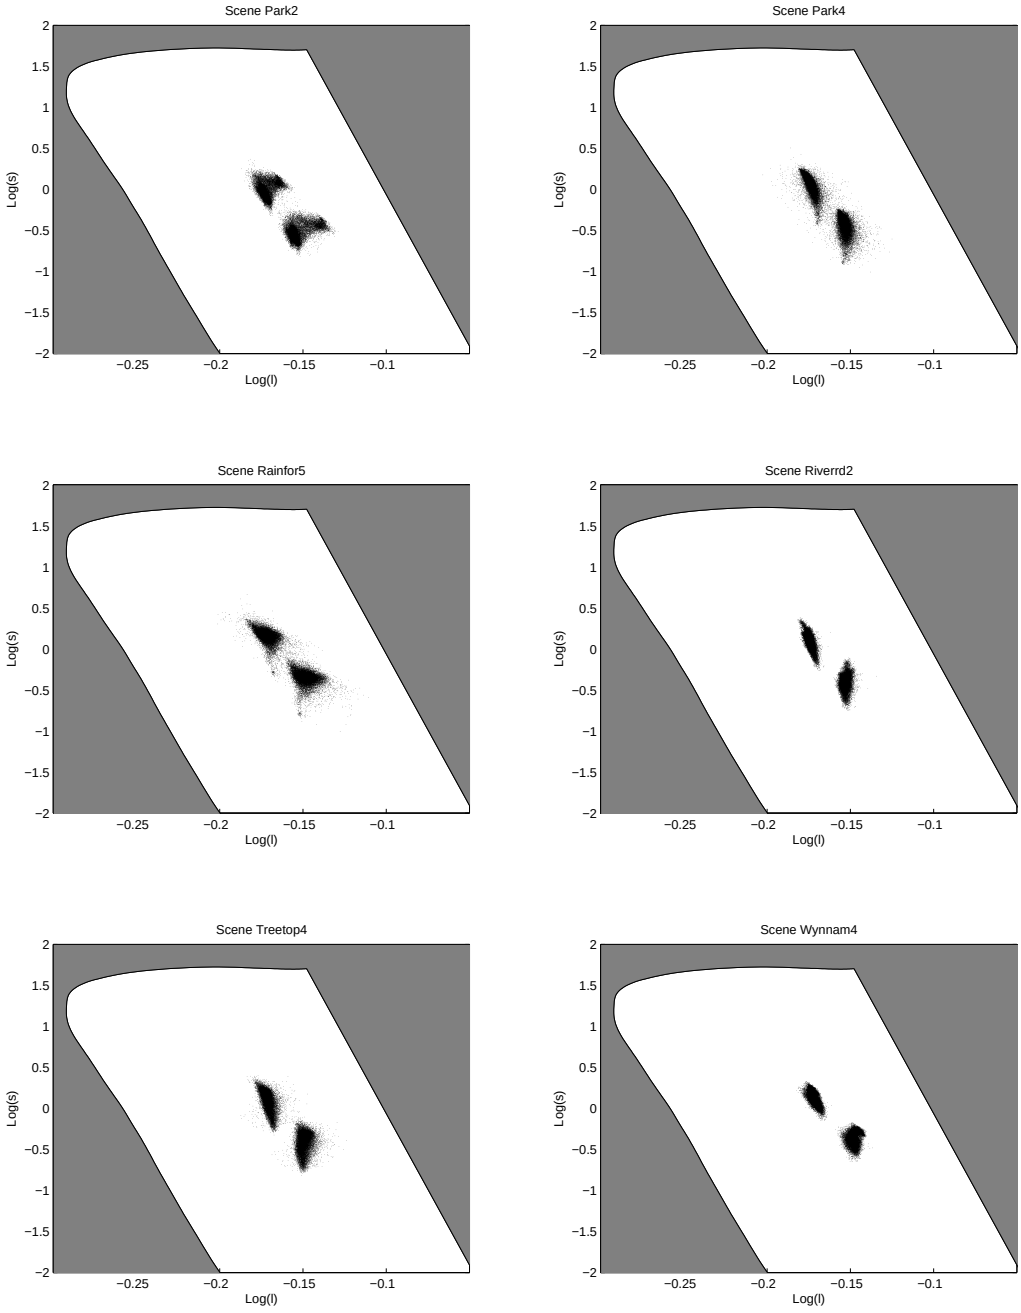

Figure C.1: (Continuation)

C.2 Chromatic scene statistics

The following tables provide values of all chromatic scene statistics of natural scenes that were obtained by the real world simulation described in section 4.2.2. Each of the four tables gives the values of the statistics of all twelve scenes under one of the four illuminants used. The scene names correspond to the original denotation of the hyperspectral reflectance image set on which the simulation is based (cf. Ruderman et al., 1998).

	Means			*Standard deviations*			*Correlations*			*Skewnesses*		
Scene	*log(l)*	*log(s)*	*log(lum)*	*log(l)*	*log(s)*	*log(lum)*	*log(l), log(lum)*	*log(s), log(lum)*	*log(l), log(s)*	*log(l)*	*log(s)*	*log(lum)*
Bank	-0,146647	-0,371533	1,050475	0,005498	0,082807	0,275640	-0,208595	0,249939	0,102505	0,018059	-0,090583	-0,216808
Barrine2	-0,154507	-0,247253	1,354421	0,002267	0,060185	0,187687	0,566657	-0,436722	-0,239483	-0,288916	0,034358	-0,415385
Cootth12	-0,148430	-0,387237	0,900122	0,003773	0,077154	0,187230	0,369019	0,144166	0,242807	0,350656	-0,601714	0,264354
Coottha8	-0,143828	-0,343757	1,105089	0,002601	0,058159	0,108352	-0,057828	-0,112425	0,008038	-0,076008	-0,556022	-0,504925
Glade	-0,148774	-0,446595	0,935791	0,003670	0,086351	0,172874	0,007491	-0,130620	0,286049	0,798059	-1,179404	0,047093
Hillside	-0,148309	-0,476853	1,009019	0,004340	0,077791	0,217747	0,018448	-0,428497	0,102595	0,486866	-0,002419	-0,618615
Park2	-0,152065	-0,519684	1,025018	0,006098	0,093680	0,248600	0,013139	-0,616883	0,351200	1,139647	0,224064	-0,593343
Park4	-0,153148	-0,463416	1,285306	0,003241	0,119991	0,238483	0,057845	0,209119	-0,335244	0,531908	-0,755670	-0,696057
Rainfor5	-0,148234	-0,357449	1,076280	0,005197	0,091031	0,233104	-0,372690	0,074316	-0,330646	0,649915	-0,991263	-0,062223
Riverrd2	-0,152965	-0,415493	1,177823	0,001966	0,093358	0,158978	0,164837	-0,583597	0,107484	0,435495	-0,104631	-0,025681
Treetop4	-0,149134	-0,418081	1,113075	0,003000	0,121077	0,243301	-0,101218	-0,591785	0,105695	0,882365	-0,579110	0,345305
Wynnam4	-0,147979	-0,378337	1,029906	0,002478	0,080354	0,150772	0,024191	0,059675	0,047806	-0,147913	-0,109284	0,483643

Table C.1: Chromatic scene statistics of twelve natural scenes under simulated daylight illumination of 4000 K colour temperature (cf. section 4.2.2).

	Means			*Standard deviations*			*Correlations*			*Skewnesses*		
Scene	*log(l)*	*log(s)*	*log(lum)*	*log(l)*	*log(s)*	*log(lum)*	*log(l), log(lum)*	*log(s), log(lum)*	*log(l), log(s)*	*log(l)*	*log(s)*	*log(lum)*
Bank	-0,155226	-0,182716	1,045179	0,004793	0,083065	0,276680	-0,251521	0,241234	0,011111	0,200862	-0,129795	-0,217209
Barrine2	-0,162884	-0,059376	1,355577	0,002154	0,059841	0,186557	0,603667	-0,445581	-0,450445	-0,362399	0,033099	-0,407125
Cootth12	-0,156845	-0,201974	0,895981	0,003277	0,078571	0,186450	0,362009	0,159751	0,131030	0,312248	-0,601618	0,261981
Coottha8	-0,152923	-0,154048	1,098355	0,002357	0,058802	0,108431	-0,073436	-0,124801	-0,124057	0,036552	-0,556417	-0,515977
Glade	-0,156654	-0,260445	0,931054	0,003107	0,088581	0,172748	0,026891	-0,123215	0,171021	0,827242	-1,185680	0,053386
Hillside	-0,155949	-0,288382	1,003539	0,003780	0,079407	0,217476	0,046151	-0,428847	0,020640	0,382532	-0,006436	-0,621854
Park2	-0,159272	-0,335042	1,021628	0,005209	0,095187	0,248051	0,069197	-0,608156	0,288448	1,089120	0,208499	-0,587572
Park4	-0,160611	-0,279844	1,283346	0,003228	0,121665	0,238690	0,018305	0,212458	-0,449408	0,574729	-0,785848	-0,688406
Rainfor5	-0,156741	-0,169249	1,072163	0,004870	0,090862	0,234532	-0,377862	0,061308	-0,388291	0,570459	-1,067225	-0,059738
Riverrd2	-0,160685	-0,233373	1,176210	0,001782	0,095105	0,158149	0,353897	-0,579447	-0,183994	0,334926	-0,105679	-0,024470
Treetop4	-0,157136	-0,230749	1,108936	0,002761	0,123209	0,242859	0,042093	-0,591734	-0,116913	0,927493	-0,601427	0,344111
Wynnam4	-0,156564	-0,194449	1,025562	0,002221	0,081844	0,150796	0,029919	0,072963	-0,212627	-0,351262	-0,093925	0,493898

Table C.2: Chromatic scene statistics of twelve natural scenes under simulated daylight illumination of 5500 K colour temperature (cf. section 4.2.2).

	Means			*Standard deviations*			*Correlations*			*Skewnesses*		
Scene	*log(l)*	*log(s)*	*log(lum)*	*log(l)*	*log(s)*	*log(lum)*	*log(l), log(lum)*	*log(s), log(lum)*	*log(l), log(s)*	*log(l)*	*log(s)*	*log(lum)*
Bank	-0,162962	-0,011711	1,052380	0,004249	0,082187	0,277746	-0,298235	0,238498	-0,127216	0,374587	-0,162890	-0,217135
Barrine2	-0,170518	0,110032	1,368789	0,002182	0,058682	0,185424	0,602953	-0,451369	-0,636818	-0,372750	0,028836	-0,396711
Cootth12	-0,164424	-0,033914	0,903991	0,002906	0,078775	0,186004	0,342016	0,176974	-0,032730	0,217895	-0,602939	0,265640
Coottha8	-0,161151	0,016921	1,104925	0,002222	0,058687	0,108503	-0,083305	-0,131138	-0,289565	0,122593	-0,551396	-0,523319
Glade	-0,163705	-0,090664	0,937713	0,002676	0,089535	0,172512	0,057672	-0,111615	-0,003351	0,817021	-1,197129	0,055691
Hillside	-0,162756	-0,115821	1,009138	0,003326	0,079961	0,217028	0,088376	-0,427945	-0,102050	0,236923	-0,014606	-0,622232
Park2	-0,165666	-0,165359	1,028602	0,004468	0,095133	0,247046	0,150236	-0,597890	0,184573	0,948052	0,191680	-0,581133
Park4	-0,167267	-0,111574	1,292089	0,003316	0,122023	0,239113	-0,031209	0,218367	-0,559049	0,579131	-0,812555	-0,676017
Rainfor5	-0,164406	0,001165	1,080367	0,004667	0,089946	0,235735	-0,377879	0,054095	-0,460531	0,450582	-1,125819	-0,055047
Riverrd2	-0,167616	-0,067108	1,185641	0,001822	0,095505	0,156887	0,514507	-0,571887	-0,480541	0,203235	-0,114304	-0,011589
Treetop4	-0,164320	-0,060084	1,116396	0,002737	0,123655	0,241805	0,191958	-0,588671	-0,358137	0,785102	-0,623275	0,344317
Wynnam4	-0,164287	-0,027997	1,033576	0,002204	0,082107	0,150833	0,024330	0,092406	-0,479461	-0,305149	-0,085188	0,514037

Table C.3: Chromatic scene statistics of twelve natural scenes under simulated daylight illumination of 8500 K colour temperature (cf. section 4.2.2).

	Means			*Standard deviations*			*Correlations*			*Skewnesses*		
Scene	*log(l)*	*log(s)*	*log(lum)*	*log(l)*	*log(s)*	*log(lum)*	*log(l), log(lum)*	*log(s), log(lum)*	*log(l), log(s)*	*log(l)*	*log(s)*	*log(lum)*
Bank	-0,170273	0,140156	1,071640	0,003917	0,080204	0,278879	-0,343728	0,240620	-0,296187	0,460780	-0,192969	-0,216688
Barrine2	-0,177839	0,259217	1,393934	0,002323	0,056755	0,184253	0,568981	-0,453148	-0,767735	-0,332798	0,022308	-0,383840
Cootth12	-0,171600	0,115784	0,923789	0,002705	0,077820	0,185836	0,299493	0,196304	-0,230758	0,104818	-0,607214	0,274972
Coottha8	-0,168920	0,168164	1,124244	0,002217	0,057779	0,108576	-0,090443	-0,130471	-0,454519	0,168491	-0,545880	-0,527239
Glade	-0,170369	0,061440	0,955419	0,002429	0,089196	0,172163	0,091697	-0,096051	-0,224568	0,773892	-1,214623	0,054683
Hillside	-0,169174	0,039096	1,025443	0,003015	0,079442	0,216387	0,140420	-0,424743	-0,256506	0,077755	-0,026592	-0,620053
Park2	-0,171689	-0,012396	1,045730	0,003915	0,093793	0,245554	0,252731	-0,586273	0,039669	0,674753	0,172651	-0,573900
Park4	-0,173566	0,039516	1,311370	0,003500	0,121041	0,239761	-0,086809	0,227041	-0,653281	0,555485	-0,837765	-0,658701
Rainfor5	-0,171659	0,152427	1,100518	0,004605	0,088174	0,236794	-0,374428	0,052066	-0,534302	0,330278	-1,174173	-0,048065
Riverrd2	-0,174208	0,081850	1,205945	0,002058	0,094565	0,155199	0,597256	-0,560365	-0,688809	0,049761	-0,129467	0,012392
Treetop4	-0,171124	0,092325	1,135114	0,002929	0,122470	0,240161	0,310941	-0,582315	-0,559064	0,559267	-0,645308	0,345881
Wynnam4	-0,171581	0,120226	1,053551	0,002395	0,081043	0,150907	0,002922	0,117002	-0,673051	-0,123826	-0,082513	0,543778

Table C.4: Chromatic scene statistics of twelve natural scenes under simulated daylight illumination of 20 000 K colour temperature (cf. section 4.2.2).

C.3 Experimental stimuli

The following tables specify the chromatic distribution in the stimuli of the experiments presented in chapter 5 in terms of eigenvectors. For further details, see section 5.2.2.

Condition	*Eigenvector*		*Standard deviation*		
	1.	2.	1.	2.	3.
Correlation(l,luminance)=-1.0	(-0.001 , -0.003545, 0.999993)	(-0.010842, 0.999935, 0.003534)	5.000034	0.152584	0.0000001
Correlation(l,luminance)=-0.7	(-0.0007 , -0.003545, 0.999993)	(-0.009706, 0.999947, 0.003538)	5.000033	0.152582	0.003249
Correlation(l,luminance)=0.0	(0.00000002, -0.003545, 0.999994)	(-0.007044, 0.999969, 0.003545)	5.000031	0.152578	0.004883
Correlation(l,luminance)=+0.7	(0.0007 , -0.003545, 0.999993)	(-0.004374, 0.999984, 0.003548)	5.000033	0.152576	0.003508
Correlation(l,luminance)=+1.0	(0.001 , -0.003545, 0.999993)	(-0.003230, 0.999988, 0.003549)	5.000034	0.152575	0.0000001

Table C.5: Eigenvectors and their standard deviations of the chromatic distribution of surround colours in stimuli for experiment 1, 2, and 3. The third eigenvector (not tabelled) is orthogonal to the plane build by the first two eigenvectors.

Condition	*Eigenvector*		*Standard deviation*		
	1.	2.	1.	2.	3.
Correlation(s,luminance)=-1.0	(0.000078, -0.030399 , 0.999538)	(-0.218041, 0.975488, 0.029685)	5.002312	0.022190	0.001229
Correlation(s,luminance)=-0.7	(0.000078, -0.021509 , 0.999769)	(-0.010150, 0.999717, 0.021509)	5.001156	0.109672	0.004859
Correlation(s,luminance)=0.0	(0.000078, -0.0000000005, 0.999999996)	(-0.006950, 0.999976, 0.0000005)	5.0	0.153604	0.004869
Correlation(s,luminance)=+0.7	(0.000078, 0.021509 , 0.999769)	(-0.017133, 0.999622, -0.021505)	5.001156	0.109683	0.004618
Correlation(s,luminance)=+1.0	(0.000078, 0.029327 , 0.999570)	(-0.106001, 0.993939, -0.029154)	5.002151	0.045797	0.001139

Table C.6: Eigenvectors and their standard deviations of the chromatic distribution of surround colours in stimuli for experiment 4 and 5. The third eigenvector (not tabelled) is orthogonal to the plane build by the first two eigenvectors.

Condition	*Eigenvector*		*Standard deviation*		
	1.	2.	1.	2.	3.
Correlation(l,s)=-1.0	(0.999477, 0.032348, 0.000037)	(-0.032348, 0.999470, 0.003546)	0.000681	0.152655	5.000031
Correlation(l,s)=-0.7	(0.999740, 0.022803, 0.000003)	(-0.022803, 0.999734, 0.003546)	0.003570	0.152614	5.000031
Correlation(l,s)=0.0	(0.999999, -0.000297, -0.000079)	(0.000297, 0.999994, 0.003545)	0.004985	0.152575	5.000031
Correlation(l,s)=+0.7	(0.999726, -0.023396, -0.000160)	(0.023396, 0.999720, 0.003542)	0.003479	0.152616	5.000031
Correlation(l,s)=+1.0	(0.999468, -0.032610, -0.000193)	(0.032611, 0.999462, 0.003541)	0.000256	0.152656	5.000031

Table C.7: Eigenvectors and their standard deviations of the chromatic distribution of surround colours in stimuli for experiment 6 and 7. The third eigenvector (not tabelled) is orthogonal to the plane build by the first two eigenvectors.

Appendix D

Colour Plates

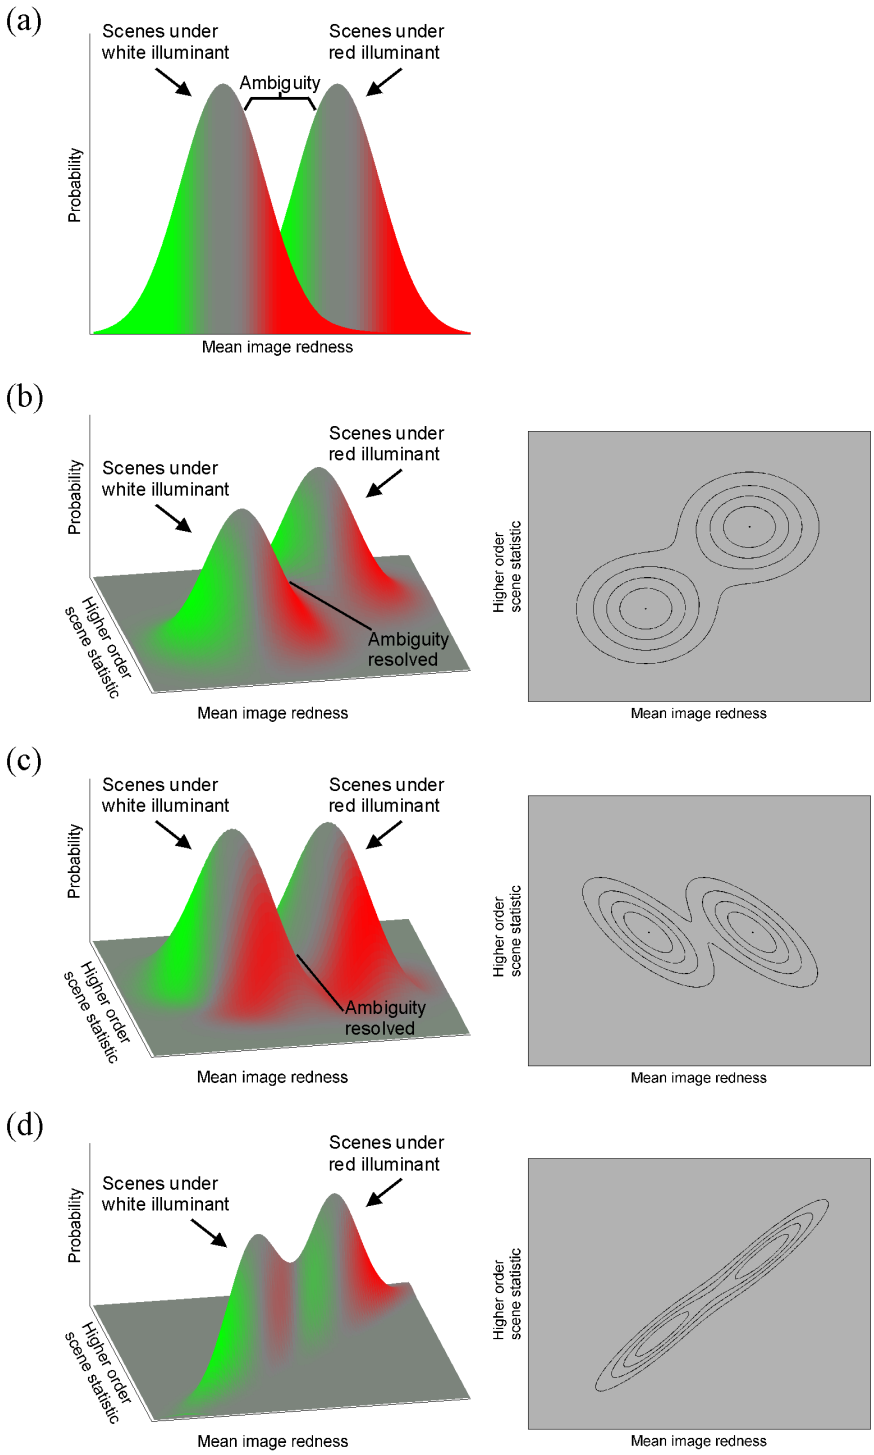

Plate I: (a) The mean image chromaticity is an ambiguous measure for estimating the illuminant. When additionally a higher order scene statistic is taken into account, this ambiguity could be resolved e.g. in constellation (b) and (c), but not in (d). Each plot in the right column shows the corresponding view from above for the respective plot in the left column, depicting probability values by contour lines.

Plate II: The set of hyperspectral images of natural environments collected by Ruderman et al. (1998). Different from this depiction, the original data consist of an estimated surface reflectance function for each pixel. This data set is used for the analysis of chromatic regularities in natural scenes (cf. Section 4.2.1, p. 35).

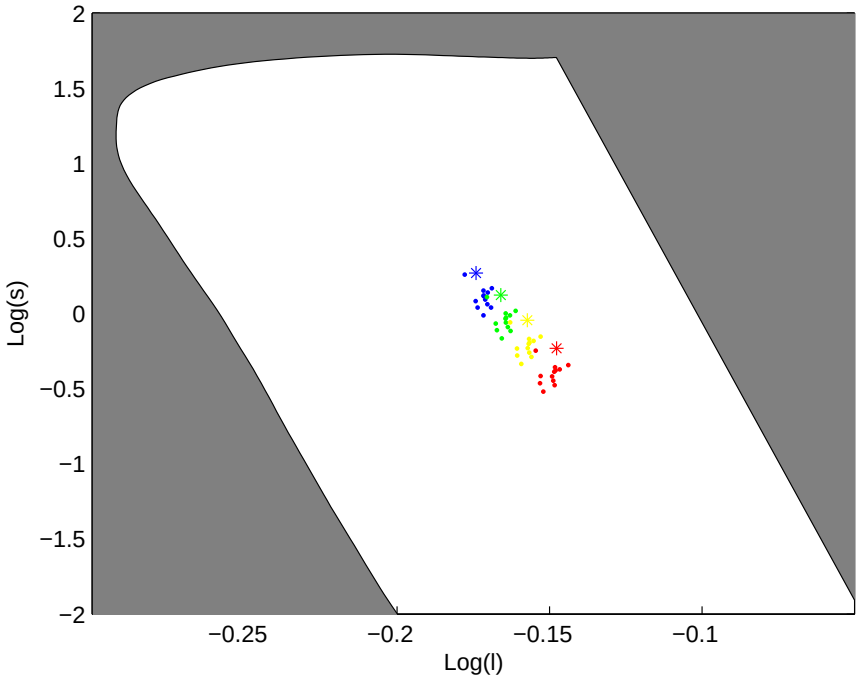

Plate III: Mean chromaticities of the twelve scenes (•) of Ruderman et al. (1998) under four different daylights (∗): colour temperature 4000 K (red), 5500 K (yellow), 8500 K (green), and 20 000 K (blue).

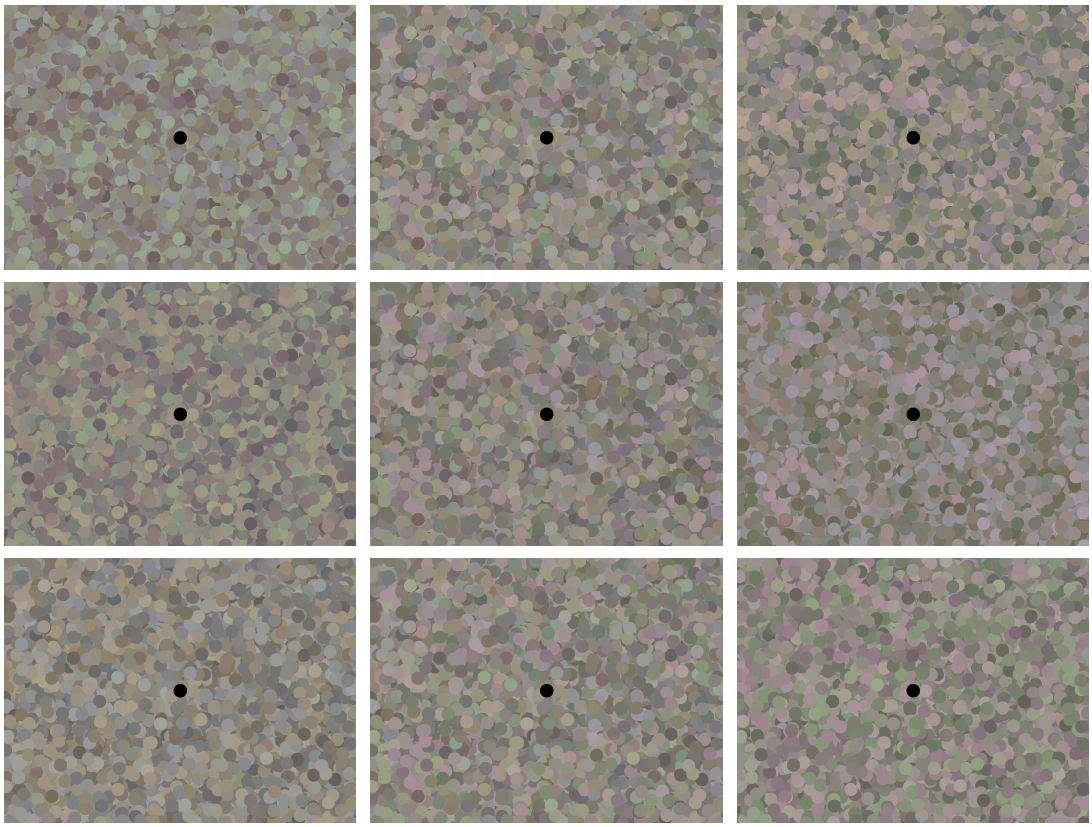

Plate IV: Examples of experimental stimuli of experiment 1 (varying degrees of luminance-redness correlation, top row), experiment 4 (varying degrees of luminance-blueness correlation, middle row), and experiment 6 (varying degrees of redness-blueness correlation, bottom row). The values of the respective correlations are -1.0 in the left column, 0.0 in the middle column, and +1.0 in the right column. The central test field which the subject had to adjust to look grey is indicated in black. Note that this print is not an exact colourimetric reproduction of the stimuli as they were displayed on the calibrated monitor during the experiments.

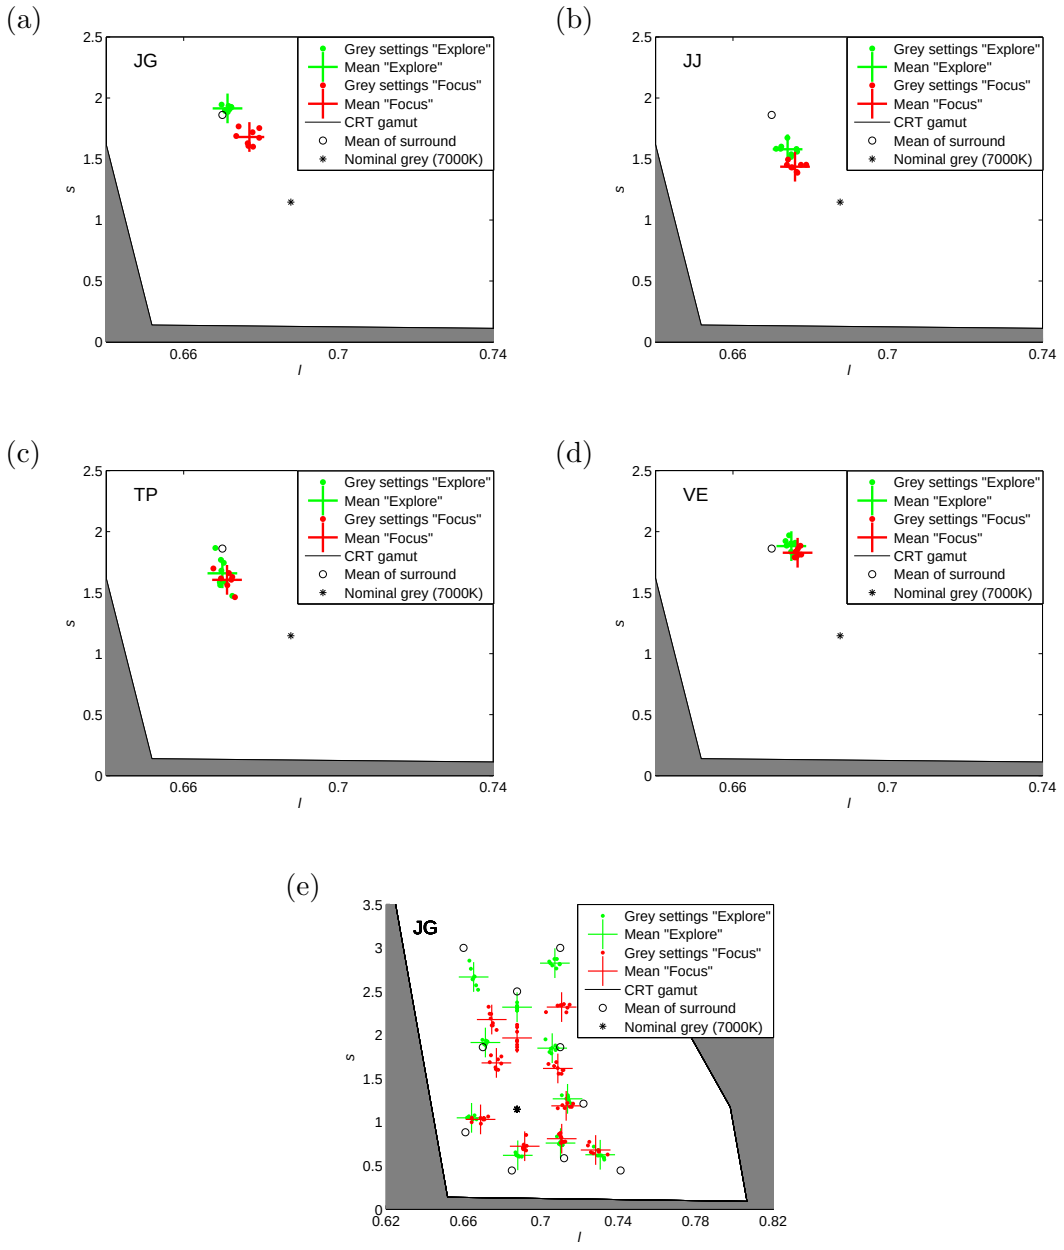

Plate V: Results of experiment I from Appendix B. (a)-(d) Individual grey settings (•) and corresponding mean (+) for the two instructions "Explore" (green) and "Focus" (red) of four subjects for the surround *blue1*. (e) Results of subject JG for all surrounds tested.

Appendix E

Lebenslauf

Persönliche Daten:

- Name: Jürgen Golz
- Geburtsdatum: 5.1.1972
- Geburtsort: Flensburg
- Familienstand: Ledig

Schulausbildung:

- 1978 – 1982: Grundschule Mohrkirch
- 1982 – 1991: Gymnasium Satrup
- Mai 1991: Erlangung der allgemeinen Hochschulreife (Abiturnote: 1,7)

Hochschulstudium:

- 1993 – 1994: Christian-Albrechts-Universität zu Kiel, Diplomstudiengang Physik
- 1995 – 2000: Christian-Albrechts-Universität zu Kiel, Diplomstudiengang Psychologie
- Oktober 1997: Vordiplom (Note „Sehr gut“)
- 1998 – 1999: Einjähriges Auslandsstudium an der UC San Diego (USA) als Stipendiat der Deutsch-Amerikanischen Fulbright Kommission
- September 2000: Diplom (Note „Sehr gut“)
- Seit Oktober 2000: Promotionsstudium an der Christian-Albrechts-Universität zu Kiel im Fach Psychologie

Berufliche Erfahrungen:

- 1991 – 1992: Zivildienst am Holländerhof in Flensburg, Wohnheim für geistig und körperlich behinderte Menschen
- 1994 – 1999: Softwareentwickler für die AHP GmbH, Glücksburg
- 1996 – 2000: Wissenschaftliche Hilfskraft für Prof. Dr. Mausfeld, Institut für Psychologie, Christian-Albrechts-Universität zu Kiel
- Seit Oktober 2000: Wissenschaftlicher Mitarbeiter von Prof. Dr. Mausfeld, Institut für Psychologie, Christian-Albrechts-Universität zu Kiel

Appendix F

Zusammenfassung

Einführung

Die vorliegende Arbeit behandelt theoretisch und experimentell die Frage, wie Menschen in der Lage sind, die Farbe von Oberflächen und Beleuchtungen wahrzunehmen. Diese Fähigkeit mag aufgrund der Mühelosigkeit und Zuverlässigkeit, mit der sie vom visuellen System erbracht wird, zwar einfach erscheinen, jedoch werden, sobald man anfängt über die Beziehung zwischen den chromatischen Eigenschaften von physikalischen Objekten auf der einen Seite und der wahrgenommenen Farbe der beiden entsprechenden psychologischen Repräsentationen auf der anderen Seiten nachzudenken, folgende Probleme offensichtlich: Das Licht, das von einer Oberfäche reflektiert wird, hängt sowohl von der spektralen Reflektanzeigenschaft der Oberfläche als auch von der spektralen Energieverteilung der Beleuchtung ab, so dass diese beiden chromatischen Komponenten in der retinalen Erregung, die den Input des visuellen Systems bildet, verschmolzen sind. Wie ist das visuelle System in der Lage, für die Wahrnehmung von Oberflächen- und Beleuchtungsfarbe diese zwei Komponenten zu entflechten? Mit dieser Leistung erfüllt der Mensch (sowie viele andere Arten) die biologisch wichtige Aufgabe, Objekte trotz sich verändernder Beleuchtungen anhand ihrer Farbe wiederzuerkennen. Unsere Fähigkeit zu beurteilen, ob eine Frucht reif ist, ob eine Schlange oder ein Pilz giftig ist und welches Auto unseres ist, hängt für gewöhnlich nicht von den chromatischen Aspekten der Beleuchtungssituation ab. Diese (annähernde) Konstanz der wahrgenommenen Farbe von Oberflächen unter verschiedenen Beleuchtungen wird Farbkonstanz genannt und ist nur ein Beispiel von vielen Konstanzleistungen, die das menschliche Wahrnehmungssystem erbringt.

Obwohl diese Frage seit mehr als einem Jahrhundert Gegenstand wissenschaftlicher Untersuchungen ist (die erste systematische Studie zur Farbkonstanz wurde

durch Helmholtz vorgenommen), steht eine befriedigende Antwort noch immer aus. Der Ansatz, der in dieser Arbeit verfolgt wird, ist, in einem ersten Schritt theoretisch zu analysieren, welche Statistiken der chromatischen Verteilung der reflektierten Lichter im retinalen Bild von natürlichen Szenen es dem visuellen System ermöglichen könnte, die Beleuchtung zu schätzen und damit Oberflächenfarben in konstanter Weise festzulegen. In einem zweiten Schritt wird die Frage, ob diese Statistiken in der Tat vom menschlichen visuellen System als Hinweisreize für die Wahrnehmung von Oberflächen- und Beleuchtungsfarbe genutzt werden, experimentell untersucht.

Im Einzelnen gliedert sich die Arbeit wie folgt: Kapitel 2 liefert eine detailliertere Darstellung des oben erwähnten Problems der Wahrnehmung von Oberfächen- und Beleuchtungsfarbe sowie des spezifischeren Problems, auf das ich mich in dieser Arbeit konzentriere, nämlich wie das visuelle System die chromatischen Eigenschaften der Beleuchtung schätzen kann. Die Rolle, die chromatische Szenenstatistiken bei der Lösung dieses Problems spielen können, wird dann auf einer allgemeinen Ebene betrachtet.

In Kapitel 3 wende ich mich der Frage zu, wie Vorhersagen darüber gewonnen werden können, welche konkreten Szenenstatisken für die Schätzung der Beleuchtung hilfreich sein können. Hierfür werden vier verschiedene idealisierte Modelle der chromatischen Umwelt, die verschiedene Annahmen über die spektralen Funktionen von Oberflächen und Beleuchtungen machen, vorgestellt und ihr Nutzen für die Ableitung möglicher Hinweisreize für die Beleuchtung wird bewertet. Für eines dieser Modelle, die sogenannte „Gaußsche Welt“, analysiere ich im Detail die Veränderungen, denen die retinalen Zapfenerregungen unterliegen, wenn die Beleuchtung wechselt. Einer der Vorteile der Gaußschen Welt gegenüber den anderen idealisierten Welten ist, dass diese beleuchtungsabhängigen Veränderungen in diesem Modell einer algebraischen Betrachtung zugänglich sind. Die Regularitäten dieser Veränderungen werden dann im Hinblick darauf diskutiert, welche Szenenstatistiken von der Beleuchtung abhängen und damit potentielle Hinweisreize für die Schätzung der chromatischen Eigenschaften der Beleuchtung darstellen.

In Kapitel 4 werden ebenfalls die Regularitäten der retinalen Veränderungen von Oberflächen unter einem Beleuchtungswechsel sowie die Szenenstatistiken, die folglich hilfreich für die Schätzung der Beleuchtung sein könnten, betrachtet, aber nun, indem anstelle einer idealisierten Welt Bilder realer, natürlicher Szenen unter verschiedenen simulierten Beleuchtungen analysiert werden. Wie sich zeigt, liefert diese Analyse der realen Welt Ergebnisse, die gut mit den Vorhersagen der Gaußschen Welt übereinstimmen.

In Kapitel 5 gehe ich dann zu der Frage über, ob die potentiellen Hinweisreize, die in den zwei vorangegangenen theoretischen Analysen gefunden wurden, nämlich die Korrelation zwischen Luminanz und Rötlichkeit im retinalen Bild,

die Korrelation zwischen Luminanz und Bläulichkeit und die Korrelation zwischen Rötlichkeit und Bläulichkeit, tatsächlich vom visuellen System ausgenutzt werden. Hierfür werden sieben psychophysikalische Experimente vorgestellt und diskutiert.

Kapitel 6 fasst die theoretischen und experimentellen Befunde der Kapitel 3, 4 und 5 zusammen und diskutiert, wie diese Ergebnisse in Beziehung stehen und welche Fragen unbeantwortet bleiben.

In dieser deutschen Zusammenfassung werde ich nach der eingehenderen Darstellung des der Farbwahrnehmung zugrunde liegenden Problems die Methoden der theoretischen und experimentellen Untersuchungen erläutern und deren wichtigste Ergebnisse dann in einem gemeinsamen Abschnitt darstellen.

Die Wahrnehmung von Oberflächen und Beleuchtungsfarbe

Das Problem

Das Wesen des Problems, das der Wahrnehmung von Oberflächen- und Beleuchtungsfarbe zugrunde liegt, kann anhand von Abbildung 2.1 (S. 9) verdeutlicht werden. In diesem Beispiel wird ein grünes Farbplättchen von einem Menschen betrachtet. Oberflächen werden durch ihre Reflektanzfunktion $S(\lambda)$ charakterisiert, d.h. durch die Funktion der reflektierten Strahlungsenergie im Verhältnis zur Strahlungsenergie, die auf die Oberfläche trifft, in Abängigkeit von der Wellenlänge. Wie in Abbildung 2.1(c) ersichtlich, reflektiert die Oberfläche in unserem Beispiel mittlere Wellenlängen besser als kurze und lange Wellenlängen.

Wird diese Oberfläche nun nacheinander von zwei verschiedenen Beleuchtungen (ein bläuliches Tageslicht mit 16 200 K Farbtemperatur, dessen spektrale Energieverteilung $I(\lambda)$ in Abbildung 2.1(a) dargestellt ist, und ein eher oranges Tageslicht mit 4500 K Farbtemperatur, dargestellt in Abbildung 2.1(b)) beleuchtet, so resultieren zwei substantiell unterschiedliche reflektierte Lichter: Die spektrale Energieverteilung des Lichtes, das von der Oberfläche unter der ersten Beleuchtung reflektiert wird (Abbildung 2.1(d)), hat mehr Energie in Bereich kurzer Wellenlängen und weniger Energie im Bereich langer Wellenlängen als das Licht, das unter der zweiten Beleuchtung reflektiert wird (Abbildung 2.1(e)), weil schon die beiden Beleuchtungen selbst in dieser Hinsicht unterschiedlich sind. Die spektrale Energieverteilung eines reflektierten Lichtes $R(\lambda)$ ergibt sich als (Wellenlänge für Wellenlänge berechnetes) Produkt der spektralen Reflektanzfunktion $S(\lambda)$ der

Oberfläche und der spektralen Energieverteilung $I(\lambda)$ der Beleuchtung. Infolgedessen unterscheiden sich auch die retinalen Zapfenerregungen, die von dieser Oberfläche unter den beiden Beleuchtungen resultieren (Abbildung 2.1(f) und (g)).

Dieses Beispiel zeigt, dass der proximale Reiz (das retinale Bild, das durch die spektrale Wechselwirkung von distalen Oberflächen und Beleuchtungen erzeugt wird) sich mit einem Wechsel der Beleuchtung substantiell verändert. Aber, das menschliche visuelle System hat wie bereits oben erwähnt die Fähigkeit, die Farbe von Oberflächen trotz wechselnder Beleuchtungen in einem hohen Maße als konstant wahrzunehmen. So würde z.B. die Oberfläche des obigen Beispiels unter beiden Beleuchtungen in nahezu demselben Grünton erscheinen (zumindest unter Sehbedingungen, die nicht hochgradig künstlich sind; siehe unten).

Trotz dieser Konstanz der wahrgenommenen Oberflächenfarbe würden sich die beiden Fälle sehr wohl unterscheiden, und zwar im Hinblick auf die wahrgenommene Beleuchtung, die auf die Oberfläche fällt: Unter dem 16 200 K Tageslicht würde die Oberfläche als grüne Oberfläche, beleuchtet von einem kalten, bläulichen Licht, erscheinen. Unter dem 4500 K Tageslicht jedoch würde sie als dieselbe Oberfläche erscheinen, nur nun beleuchtet von einem wärmeren, orangen Licht. Farbkonstanz bedeutet nicht, dass eine Oberfläche unter wechselnden Beleuchtungen in jeder Hinsicht identisch aussieht (wie z.B. schon von Katz (1911) und Gelb (1929) betont wurde). Mausfeld (2003) bietet einen theoretischen Rahmen, der es erlaubt, diese Einsicht in Termini der modernen Kognitionswissenschaft zu formulieren: Das visuelle System erzeugt (unter geeigneten Bedingungen) sowohl eine Repräsentation für die Oberfläche als auch eine Repräsentation für die Beleuchtung, die beide einen Parameter für Farbe aufweisen.

Das Problem, diese beiden Parameter zu bestimmen, kann nicht dadurch gelöst werden, den retinalen Reiz lokal zu analysieren, weil eine gegebene Zapfenerregung durch viele Kombinationen von distalen Oberflächen und Beleuchtungen erzeugt werden kann. Dass das visuelle System für die Wahrnehmung der Oberflächenfarbe sich auf globalere Informationen stützt als die lokale retinale Erregung, wird z.B. aus der Tatsache ersichtlich, dass die Wahrnehmung sich drastisch ändert, wenn dieselbe lokale Information isoliert dargeboten wird. Wenn man z.B. die Oberfläche aus dem obigen Beispiel als das einzige Objekt in einem ansonsten vollständig dunklen Raum, beleuchtet von dem 16 200 K Licht, betrachten würde, erschiene sie nicht in demselben Grün wie in einer reichhaltigeren Szene mit mehr Objekten, sondern eher türkis. Darüber hinaus erschiene diese Farbe gar nicht wie eine Oberflächenfarbe, sondern hätte eher eine Erscheinungsweise, die auch als „Lochfarbe“ oder „Flächenfarbe“ bezeichnet wird (Katz, 1911). Ein weiteres Beispiel ist der Mond am Nachthimmel, der nicht wie eine felsig-graue Oberfläche aussieht, die von der Sonne beleuchtet wird, sondern eher selbst wie

eine Lichtquelle erscheint. Es scheint so zu sein, dass eine perzeptuelle Zerlegung in zwei Repräsentationen, eine für die Oberfläche und eine für die Beleuchtung, in Situationen wie diesen nicht auftritt.

Was traditionell „das Farbkonstanzproblem“ genannt wird, kann daher als zwei zusammenhängende Probleme betrachtet werden:

1. Wie schätzt das visuelle System die chromatischen Parameter der Beleuchtung um den Farbparameter der Beleuchtungsrepräsentation zu bestimmen?

2. Wie legt das visuelle System, unter Berücksichtigung dieser Beleuchtungsschätzung, die Farbparameter der Oberflächenrepräsentationen fest?

Die vorliegende Arbeit widmet sich vorwiegend der ersten Frage, indem ich analysiere, welche globalen chromatischen Hinweisreize im retinalen Bild verfügbar sind, um die Beleuchtung zu schätzen (Kapitel 3 und 4), und experimentell untersuche, ob diese Hinweisreize in der Tat vom visuellen System genutzt werden (Kapitel 5). Die zweite Frage wird nur indirekt insofern behandelt, als die Analysen in Kapitel 3 und 4 Ergebnisse darüber liefern, wie die Zapfenerregungen von verschiedenen Oberflächen unter Beleuchtungswechsel transformiert werden. Diese Befunde können verwendet werden, um Vorhersagen darüber abzuleiten, auf welche Art das visuelle System möglicherweise die Oberflächenfarben auf der Basis des retinalen Bildes und der geschätzten Beleuchtung festlegt.

Das Schätzen der Beleuchtung

In der vorliegenden Arbeit verwende ich als Information für die Schätzung der Beleuchtung die beleuchtungsabhängigen chromatischen Eigenschaften des retinalen Bildes. Beleuchtungen beeinflussen die Verteilung von retinalen Erregungen, die von den Oberflächen in der Szene hervorgerufen werden, (dargestellt z.B. als Tripple in einem Raum der Zapfenerregungen). Welche Beleuchtung vorliegt, könnte daher anhand von Statistiken dieser chromatischen Verteilung im retinalen Bild, sogenannten „chromatische Szenenstatistiken“ erkannt werden. Diese Herangehensweise lässt andere potentielle Hinweisreize außer Acht, insbesondere solche, die von der dreidimensionalen Geometrie der Szene abhängen, wie z.B. Glanzlichter (D'Zmura & Lennie, 1986; Lee, 1986; Tominaga & Wandell, 1989; Maloney & Yang, 2003) und gegenseitige Reflektionen in Ecken (Funt et al., 1991; Funt & Drew, 1993; Bloj et al., 1999). Die Situationen, die hier betrachtet werden, sind nur solche in einer Welt von koplanaren, diffuse reklektierenden Oberflächen, mit einer einzelnen homogenen Beleuchtung in jeder Szene.

Der chromatische Mittelwert

Wenn die Beleuchtung einer Szene wechselt, z.B. in Richtung mehr Energie im Bereich langer Wellenlängen, dann verändern sich auch alle reflektierten Lichter entsprechend und die über das gesamte Bild gemittelte Chromatizität wird rötlicher. Eine einfache Herangehensweise könnte daher sein, den Mittelwert der Chromatizität einer Szene als Schätzer für die chromatische Eigenschaft der Beleuchtung zu verwenden und diesen Schätzer zu benutzten, um die Veränderungen der Chromatizitäten von Objekten bei Beleuchtungswechseln auszugleichen. Eine ähnliche Überlegung wurde bereits von Helmholtz (1896) vorgeschlagen.

Buchsbaum (1980) hat einen Farbkonstanz-Algorithmus vorgeschlagen, der den räumlichen Mittelwert der Zapfenerregungen verwendet, um die Beleuchtung einer Szene zu schätzen. Seine Schätzung der Beleuchtung basiert auf der Annahme, dass die über eine Szene gemittelte Reflektanz für alle Szenen einem internen Reflektanz-Standard $S_0(\lambda)$ gleicht. Solch eine Annahme wird auch „Grey World Assumption“ genannt. Für Szenen, in denen die mittlere Reflektanz von $S_0(\lambda)$ abweicht, attribuiert der Algorithmus eine nicht-neutrale mittlere Zapfenerregung jedoch fälschlicherweise auf eine spektral unausgeglichene Szene. Jede Veränderung der mittleren Zapfenerregung wird interpretiert als durch einen Beleuchtungswechsel verursacht. Dieser Algorithmus ist daher nicht in der Lage, Variationen der mittleren Chromatizität des retinalen Bildes aufgrund der Zusammensetzung der Oberflächen einer Szene von solchen aufgrund der Beleuchtung zu unterscheiden. Diese Schwäche ist allen Modellen eigen, die ein räumliches Mittel der Chromatizität einer Szene verwenden, um die Beleuchtung zu schätzen. Der Grund hierfür ist, dass dieses Maß ambiguär ist: Für ein gegebenes retinales Bild einer Szene könnte die mittlere Chromatizität z.b. rötlich sein sowohl wegen einer Dominanz rötlicher Oberflächen in dieser Szene als auch wegen einer rötlichen Beleuchtung.

Diese Ambiguität wird in der Farbabbildung I(a) auf Seite 149 veranschaulicht, indem eine Stichprobe von Szenen unter zwei Beleuchtungen (eine weiße und eine rote) betrachtet wird. Dieselbe mittlere Chromatizität kann entweder von einer rötlichen Szene unter weißer Beleuchtung resultieren oder auch von einer weißen Szene unter rötlicher Beleuchtung, da diese beiden Beleuchtungen, auf die Stichprobe der Szenen angewendet, sich überlappende Verteilungen der mittleren Rötlichkeit erzeugen. In Abschnitt 4.3.2 auf Seite 45 zeige ich, dass diese Ambiguität nicht nur ein theoretisches Problem ist, sondern dass die mittleren Chromatizitäten von natürlichen Szenen, die von derselben Beleuchtung beleuchtet werden, eine große Varianz aufweisen können.

Der Nutzen höherer Szenenstatistiken

Trotz dieser geschilderten Ambiguität sind menschliche Betrachter sehr wohl irgendwie in der Lage, eine weiße Szene unter rötlicher Beleuchtung von einer rötlichen Szene unter weißer Beleuchtung zu unterscheiden (Gilchrist & Ramachandran, 1992; Kraft & Brainard, 1999). Welche Informationen helfen ihnen möglicherweise, die Ambigutät des chromatischen Mittelwertes aufzulösen? Der Nutzen, den höhere Szenenstatistiken (Standardabweichungen, Korrelationen, Schiefen der chromatischen Verteilung im retinalen Bild) für die Schätzung der Beleuchtung liefern könnten, wenn sie zusätzlich zum chromatischen Mittelwert berücksichtigt werden, wird in den schematischen Konstellationen (b)-(d) in Farbabbildung I auf Seite 149 veranschaulicht.

In dem Beispiel (b) von Farbabbildung I hängt eine höhere Szenenstatistik systematisch von der Beleuchtung ab: Sie nimmt für alle Szenen unter der roten Beleuchtung zu. Daher ist die Überlappung der zwei Wahrscheinlichkeitsverteilungen der Szenen unter den beiden Beleuchtungen reduziert und die Ambiguität des chromatischen Mittelwertes könnte vom visuellen System aufgelöst werden, indem es diese höhere Szenenstatistik berücksichtigt. Ein hoher Wert dieser Statistik im retinalen Bild einer Szene wäre Evidenz dafür, dass die Beleuchtung rötlich ist.

Die Konstellation (c) in Farbabbildung I zeigt eine weitere Art, in der höhere Szenenstatistiken die Ambiguität des chromatischen Mittelwertes reduzieren könnten. Obwohl die höhere Szenenstatistik in diesem Fall für eine bestimmte Szene sich unter den beiden Beleuchtungen nicht unterscheidet, trennt sie die beiden Wahrscheinlichkeitsverteilungen, wenn sie zusätzlich zu der mittleren Rötlichkeit betrachet wird. Auch in diesem Fall wäre, bei gleichem chromatischen Mittelwert, ein größerer Wert dieser höheren Szenenstatistik im retinalen Bild Evidenz dafür, dass die Beleuchtung rötlich ist.

Aber nicht jede Statistik, selbst wenn sie systematisch von der Beleuchtung abhängt, ist hilfreich, um die Ambiguität des chromatischen Mittelwertes aufzulösen, wie in Farbabbildung I(d) veranschaulicht wird. Hier sind zwei Statistiken so miteinander korreliert, dass die Kombination dieser beiden nur so hilfreich für die Schätzung der Beleuchtung ist, wie jeder der beiden Statistiken alleine.

Theoretische Welten

Welche Heuristik ist für die Suche nach chromatischen Hinweisreizen geeignet?

Wie oben erwähnt, könnte das visuelle System chromatische Szenenstatistiken als Hinweisreize für die Wahrnehmung von Oberflächen- und Beleuchtungsfarbe verwenden, um die Ambiguität aufzulösen, die der mittleren Chromatizität des retinalen Bildes anhaftet. Aber welche chromatischen Szenenstatistiken hängen systematisch von den chromatischen Eigenschaften der Beleuchtung ab und stellen somit hilfreiche Hinweisreize für die Beleuchtung dar? Für eine Herangehensweise an diese Frage, die von computationalen Überlegungen im Sinne von Marr (1982) geleitet werden, ist ein idealisiertes Welt-Modell wünschenswert, das beschreibt, wie die chromatischen Eigenschaften im proximalen Reiz von den chromatischen Eigenschaften der distalen Oberflächen und Beleuchtungen abhängen. Basierend auf einem solchen Modell könnte man analysieren, wie die proximalen Chromatizitäten sich verändern, wenn die Beleuchtung sich verändert. Jedes realistische Modell, das systematische Veränderungen der proximalen Chromatizitäten bei Beleuchtungswechseln vorhersagt, stellt eine aussichtsreiche Heuristik für die Suche nach Hinweisreizen für die Beleuchtung dar: Eine Statistik der chromatischen Verteilung im proximalen Reiz, die eine von dem Modell vorhergesagte beleuchtungsabhängige Regularität einfängt, kann als Kandidat für einen Hinweisreiz betrachtet werden, der vom visuellen System zur Schätzung der Beleuchtung benutzt werden könnte.

Im Allgemeinen ist es nicht möglich, die chromatischen Eigenschaften eines reflektierten Lichtes zu bestimmen, ohne die spektralen Funktionen der entsprechenden Oberfläche und Beleuchtung zu kennen. Anders als die Ergebnisse additiver Lichtmischungen kann die Wechselwirkung von Oberflächen und Beleuchtungen nicht auf der Ebene von Zapfenerregungen (oder anderen Farbrohwerten) gesetzmäßig beschrieben werden, sondern, wie auch die Ergebnisse subtraktiver Lichtmischungen, nur auf der spektralen Ebene. Daher muss jedes idealisierte Welt-Modell explizite Annahmen über die spektralen Eigenschaften von Oberflächen und Beleuchtungen machen. Diese spektralen Annahmen eines Modells bestimmen dann, ob (und welche) beleuchtungsabhängigen Regularitäten durch das Modell vorgesagt werden.

Die folgenden vier idealisierten Welten werden in Kapitel 3 betrachtet:

1. In der *Chaotischen Welt* werden keine Einschränkungen für die spektralen Eigenschaften von Oberflächen und Beleuchtungen gemacht. Ein einfaches

Szenario mit einer bösartigen Wahl von Oberflächen- und Beleuchtungsspektren (siehe Abbildung 3.1 auf Seite 21) zeigt, dass in einer solchen uneingeschränkten Welt das Wechselspiel zwischen Oberflächen und Beleuchtungen chaotisch sein kann: Im Allgemeinen liegt keine Systematik in den Veränderungen der Zapfenerregungen unter einem Wechsel der Beleuchtung vor.

2. Eine Ordnung in den beleuchtungsabhängigen Veränderungen der Zapfenerregungen kann es nur unter geeigneten Einschränkungen der spektralen Eigenschaften von Oberflächen und Beleuchtungen geben. Die Einschränkungen der *Drei-Band-Welt* führen zu einem der höchsten Grade von Ordnung: Alle Zapfenerregungen einer Szene werden durch einen Beleuchtungswechsel der gleichen multiplikativen Transformation unterworfen. Die Einschränkungen, die diese idealisierte Welt macht, sind jedoch als Beschreibungen der spektralen Eigenschaften der natürlichen Welt sehr unrealistisch.

3. In der *Linearen Welt* wird eine realistischere Annäherung an natürlich auftretende Spektren angestrebt, indem Reflektanzfunktionen von Oberflächen und Energieverteilungen von Beleuchtungen durch gewichtete Summen von spektralen Basisfunktion modelliert werden. Zwar kann ein solches Modell jeden Satz von spektralen Funktionen mittels einer ausreichend großen Anzahl geeignet gewählter Basisfunktionen perfekt beschreiben, ob jedoch eine befriedigende Annäherung auch mit einer Zahl von Basisfunktionen erreicht werden kann, die klein genug ist, um das Modell auch als Heuristik für die Farbkonstanz bei Menschen plausibel zu machen, wird gegenwärtig noch heiß diskutiert. Trotz einer Vielzahl von Arbeiten, die die Lineare Welt als analytisches Werkzeug verwendet haben, hat dieses Modell bis heute nicht zu Vorhersagen darüber geführt, welche chromatischen Szenenstatistiken vom visuellen System als Hinweisreize für die Beleuchtung verwendet werden könnten.

4. In der *Gaußschen Welt* werden alle spektralen Funktionen (Oberflächen Reflektanzen, Energieverteilungen von Beleuchtungen, aber auch Sensitivitätskurven der Zapfen) durch Gaußsche Funktionen beschrieben. Neben vielen anderen Vorteilen dieser Art der Modellierung liefert die Gaußsche Welt etwas, das in den anderen idealisierten Welten nicht verfügbar ist: Eine einfache algebraische Lösung für die aus der Wechselwirkung der Oberflächen und Beleuchtungen resultierenden retinalen Zapfenerregungen. Mit Hilfe dieses Werkzeuges lassen sich eine Vielzahl von Regularitäten der beleuchtungsabhängigen Veränderungen der Zapfenerregungen vorhersagen. Diese Vorhersagen werden dann in Kapitel 4 durch eine Analyse von Bildern natürlicher Szenen überprüft. Die drei vielversprechendsten potentiellen Hinweisreize, die sich hieraus ergeben, werden dann in Kapitel 5 darauf-

hin getestet, ob das visuelle System diese tatsächlich für die Wahrnehmung von Oberflächen- und Beleuchtungsfarbe nutzt. Die Ergebnisse der Analyse der Gaußschen Welt werde ich in dieser deutschen Zusammenfassung daher weiter unten zusammen mit den Ergebnissen von Kapitel 4 und 5 darstellen.

Die reale Welt

Welche chromatischen Hinweisreize sind in der natürlichen Umwelt verfügbar?

In Kapitel 3 wird analysiert, welche Regularitäten der retinalen Veränderungen von Oberflächen unter einem Beleuchtungswechsel sich in der natürlichen Umwelt finden lassen. Ziel ist weiterhin, Szenenstatistiken zu finden, die hilfreich für die Schätzung der Beleuchtung sein könnten. Anders als in Kapitel 2 werden jetzt jedoch nicht theoretische Modelle idealisierter Welten betrachtet, sondern Bilder von real existierenden natürlichen Szenen, für die verschiedene Beleuchtungen simuliert werden. Hierbei wird besonderes Augenmerk auf die Vorhersagen der Gaußschen Welt gerichtet.

Um die chromatischen Eigenschaften retinaler Bilder derselben Szenen unter verschiedenen Beleuchtungen zu bestimmen, habe ich sogenannte hyperspektrale Bilder von natürlichen Szenen verwendet. Hyperspektrale Bilder geben, anders als die Bilder üblicher Digitalkameras, für jeden Bildpunkt die gesamte Reflektanzfunktion der entsprechenden Oberfläche in der Szene an. Ein hyperspektrales Bild ermöglicht es daher, das retinale Bild der entsprechenden Szene unter verschiedenen Beleuchtungen zu simulieren, indem man verschiedene Beleuchtungsspektren auf die Reflektanzspektren der Oberflächen des hyperspektralen Bildes anwendet. Ich habe dabei auf einen Satz von zwölf hyperspektralen Bildern natürlicher Szenen zurückgegriffen, die Ruderman et al. (1998) aufgenommen haben. Als Beleuchtungsspektren habe ich vier verschiedene Tageslichter von 4000 K bis 20 000 K Farbtemperatur verwendet. Die Ergebnisse der Analyse der beleuchtungsabhängigen Veränderungen in diesen so simulierten retinalen Bildern werden, wie bereits erwähnt, weiter unten gemeinsam mit den Ergebnissen der anderen Untersuchungen dargestellt.

Experimente

Welche chromatischen Hinweisreize werden vom menschlichen visuellen System genutzt?

In diesem Kapitel werden die drei vielversprechendsten Kandidaten für Hinweisreize für die Beleuchtung, die sich aus den theoretischen Analysen in Kapitel 3 und 4 ergeben haben, daraufhin getestet, ob das visuelle System diese potentiellen Hinweisreize tatsächlich für die Wahrnehmnung von Oberflächen- und Beleuchtungsfarbe nutzt. Diese drei Kandidaten sind:

1. die Korrelation zwischen Luminanz und Rötlichkeit im retinalen Bild,
2. die Korrelation zwischen Luminanz und Bläulichkeit im retinalen Bild,
3. die Korrelation zwischen Rötlichkeit und Bläulichkeit im retinalen Bild.

Die prinzipielle Idee der experimentellen Untersuchung ist es, in einem separaten Experiment für jeden dieser drei Szenenstatistiken den Wert dieser Statistik bei der Darbietung vielfarbiger Reize systematisch zu variieren und dabei alle anderen Statistiken der chromatischen Verteilung konstant zu halten. Versuchspersonen müssen ein Testfeld in diesen Reizen so einstellen, dass es grau aussieht. Hängt dabei die Chromatizität der Lichter, die von den Versuchspersonen als grau ausgewählt werden, in der erwarteten systematischen Weise von der Szenenstatistik ab, die in dem jeweiligen Experiment getestet wird, dann wird dies als Evidenz dafür betrachtet, dass das visuelle System diese Szenenstatistik für die Wahrnehmung von Oberflächen- und Beleuchtungsfarbe berücksichtigt und dass es diese Szenenstatistik also in der Tat ausnutzt.

Für die Korrelation zwischen Luminanz und Rötlichkeit sowie die Korrelation zwischen Rötlichkeit und Bläulichkeit wurde jeweils neben einem Experiment mit Reizen, die im Mittel chromatisch neutral waren, auch noch ein Experiment durchgeführt, in denen Reize mit nicht-neutralen (bläulichen, rötlichen, gelblichen und grünlichen) chromatischen Mittelwerten verwendet wurden. Für die Korrelation zwischen Luminanz und Rötlichkeit sowie die Korrelation zwischen Luminanz und Bläulichkeit wurde außerdem jeweils ein weiteres Experiment durchgeführt, in denen statt des nicht-luminanzgewichteten Mittelwertes der Chromatizität im Reiz der lumininanz-gewichtete Mittelwert für alle Korrelationsbedingungen konstant gehalten wurde. (Für diese beiden Szenenstatistiken ist es prinzipiell nicht möglich, beide Arten von Mittelwerten der Chromatizität gleichzeitig konstant zu halten.) Somit ergiebt sich eine Gesamtzahl von sieben Experimenten.

Überblick über die wichtigsten Ergebnisse

Im Folgenden werde ich die wichtigsten Befunde der vorliegenden Arbeit kurz darstellen, indem ich für jeden der drei wichtigsten Kandidaten für Hinweisreize für die Beleuchtung, die sich aus den theoretischen Analysen ergeben haben und die daher experimentell untersucht wurden, die Ergebnisse der Analysen und Experimente zusammenfasse und diskutiere, wie diese Ergebnisse in Beziehung stehen. Bei der Darstellung der Ergebnisse der theoretischen Analysen wird nochmals unterschieden in folgende drei Teile:

1. Die Analyse der Vorhersagen der Gaußschen Welt für die Veränderung der Zapfenerregung von Oberflächen unter Beleuchtungswechseln (Kapitel 3).
2. Die Analyse der Veränderung der Zapfenerregung von Oberflächen unter Beleuchtungswechseln in der realen Welt (Kapitel 4).
3. Die Analyse der Szenenstatistiken im gesamten retinalen Bild natürlicher Szenen in der realen Welt (Kapitel 4).

Während die ersten beiden Teile Regularitäten auf der Ebene einzelner Oberflächen bzw. Bildpunkte betrachten, bezieht sich der dritte Teil der theoretischen Analysen auf die Ebene von Szenenstatistiken des gesamten retinalen Bildes.

Zusammenfassung für die Korrelation zwischen Luminanz und Rötlichkeit

Alle drei Teile der theoretischen Analyse deuten darauf hin, dass die Korrelation zwischen Luminanz und Rötlichkeit im retinalen Bild hilfreich für die Schätzung der chromatischen Eigenschaften der Beleuchtung ist:

1. Die algebraische Lösung für die Zapfenerregungen in der Gaußschen Welt macht die Vorhersage, dass Oberflächen, die der Beleuchtung chromatisch ähnlich sind, relative zu Oberflächen, die der Beleuchtung unähnlich sind, als etwas heller im retinalen Bild wiedergegeben werden (siehe Seite 29). Wenn die Beleuchtung z.B. rötlicher wird, sollten daher die rötlichen Oberflächen heller werden, relativ zu z.B. grauen oder grünlichen Oberflächen.
2. Diese beleuchtungsabhängige Regularität ist in der Simulation der realen Welt als für die Bildpunkte von natürlichen Szenen zutreffend festgestellt worden (siehe Seite 41).

3. Die Analyse der natürlichen Szenen auf der Ebene von chromatischen Szenenstatistiken hat gezeigt, dass diese beleuchtungsabhängige Regularität die Korrelation zwischen Luminanz und Rötlichkeit zu einem hilfreichen Hinweisreiz für die Beleuchtung macht. Dies trifft zu, obwohl diese beleuchtungsabhängige Regularität mit einer zweiten Regularität interagiert, so dass der Wert der Korrelation einer bestimmten Szene nahezu unbeeinflusst von Beleuchtungswechseln ist. Trotzdem ist diese Statistik in der Simulation der realen Welt von diagnostischem Wert für die Beleuchtung: Beispielsweise zeigt eine hohe Korrelation zwischen Luminanz und Rötlichkeit eine rötliche Beleuchtung an (siehe Seite 49).

In Übereinstimmung mit diesen theoretischen Befunden haben die experimentellen Untersuchungen in Kapitel 5 starke Evidenz dafür geliefert, dass das visuelle System die Korrelation zwischen Luminanz und Rötlichkeit für die Schätzung der Beleuchtung verwendet:

1. Die Ergebnisse aller elf Versuchspersonen in Experiment 1 weisen die erwartete Abhängigkeit der Graueinstellungen von der Korrelation zwischen Luminanz und Rötlichkeit im Reiz auf. Diese Effekte sind individuell signifikant und substantiell für alle getesteten elf Versuchspersonen. Diese Ergebnisse bestätigen die Befunde eines früheren Experimentes, das ich in Zusammenarbeit mit Donald I. A. MacLeod durchgeführt habe (Golz & MacLeod, 2002).

2. Experiment 2 erweitert diesen Effekt der Korrelation zwischen Luminanz und Rötlichkeit, der in Experiment 1 für Reize mit neutralem chromatischem Mittelwert gefunden wurde, auf Reize mit nicht-neutralem chromatischem Mittelwert. Die erhaltenen Effekte sind individuell signifikant und substantiell für alle vier chromatischen Mittelwerte bei allen vier Versuchspersonen, die in Experiment 2 getestet wurden.

3. Experiment 3 zeigt, dass dieser Effekt auch zutrifft, wenn (anstelle des nicht-luminanzgewichteten Mittelwertes der Chromatizität im Reiz) der luminanzgewichtete Mittelwert der Chromatizität für alle Bedingungen der Korrelation zwischen Luminanz und Rötlichkeit konstant gehalten wird. Die gefundenen Effekte sind individuell signifikant und substantiell für alle acht Versuchspersonen, die in Experiment 3 getestet wurden.

Zusammenfassend kann festgestellt werden, dass die robusten und substantiellen Effekte der Korrelation zwischen Luminanz und Rötlichkeit, die in diesen Experimenten gefunden wurden, in bester Übereinstimmung mit den Vorhersagen der theoretischen Analysen sind und darauf hinweisen, dass das menschliche visuelle System diese chromatische Szenenstatistik als Hinweisreiz für die Wahrnehmung von Oberflächen- und Beleuchtungsfarbe ausnutzt.

Zusammenfassung für die Korrelation zwischen Luminanz und Bläulichkeit

Die theoretische Evidenz dafür, dass die Korrelation zwischen Luminanz und Bläulichkeit im retinalen Bild hilfreich für die Schätzung der chromatischen Eigenschaft der Beleuchtung ist, ist durchwachsener als in dem Fall der Korrelation zwischen Luminanz und Rötlichkeit, da nur zwei der drei Teile der theoretischen Analyse positive Ergebnisse liefern:

1. Ähnlich der beleuchtungsabhängigen Regularität, die der Korrelation zwischen Luminanz und Rötlichkeit zugrunde liegt, macht die Gaußsche Welt die Vorhersage, dass die bläulichen Oberflächen an Luminanz gewinnen, verglichen mit z.B. grauen oder gelblichen Oberflächen, wenn die Beleuchtung bläulich wird (siehe Seite 29).

2. Diese beleuchtungsabhängige Regularität ist in der Analyse der Simulation der realen Welt auf der Ebene der Bildpunkte für alle verwendeten natürlichen Szenen als zutreffend festgestellt worden (siehe Seite 41).

3. Überraschenderweise ist jedoch die Korrelation zwischen Luminanz und Bläulichkeit in der Analyse der Simulation der realen Welt auf der Ebene der Szenenstatistiken nahezu unabhängig von der Beleuchtung (siehe Seite 51).

Mit Hinblick auf diese gemischten theoretischen Befunde ist es sehr interessant, dass Experiment 4 und 5 starke Evidenz dafür liefern, dass das menschliche visuelle System die Korrelation zwischen Luminanz und Bläulichkeit für die Schätzung der Beleuchtung berücksichtigt:

1. Die Ergebnisse aller sechs Versuchspersonen in Experiment 4 weisen die erwartete Abhängigkeit der Graueinstellungen von der Korrelation zwischen Luminanz und Bläulichkeit im Reiz auf. Diese Effekte sind individuell signifikant und substantiell für alle getesteten sechs Versuchspersonen.

2. Dieselben sechs Versuchspersonen zeigen ebenfalls individuell signifikante und substantielle Effekte in Experiment 5, in dem (anstelle des nichtluminanzgewichteten Mittelwertes der Chromatizität im Reiz) der luminanzgewichtete Mittelwert der Chromatizität für alle Bedingungen der Korrelation zwischen Luminanz und Bläulichkeit konstant gehalten wird.

Aufgrund der Übereinstimmung der Ergebnisse, die zwei der drei Teile der theoretischen Analyse mit den Experimenten aufweisen, könnte man geneigt sein zu schließen, dass das menschliche visuelle System die Korrelation zwischen Luminanz und Bläulichkeit berücksichtigt, weil es dadurch einen diagnostischen Wert dieser Statistik für die Schätzung der Beleuchtung ausnutzen kann. Es bleibt jedoch die Frage, warum der dritte Teil der theoretischen Analyse, die Analyse der Simulation der realen Welt auf der Ebene der chromatischen Szenenstatistiken, für diese Statistik keinen solchen diagnostischen Wert vorhersagt. Es könnte natürlich sein, dass die in der Simulation der realen Welt verwendeten Szenen in dieser Hinsicht schlicht nicht repräsentativ sind für die chromatische Umwelt, an die sich das menschliche visuelle System angepasst hat. Dieser Punkt bedarf jedoch weiterer Erforschung und bis die Ursache für the Diskrepanz zwischen den Befunden für die Korrelation zwischen Luminanz und Bläulichkeit nicht gefunden ist, kann keine definitive Schlussfolgerung gezogen werden.

Zusammenfassung für die Korrelation zwischen Rötlichkeit und Bläulichkeit

Während der potentielle Nutzen für die Schätzung der Beleuchtung im Falle der beiden obigen Korrelationsstatistiken aus einer Steigerung der Intensität für Oberflächen, die der Beleuchtung chromatisch ähnlich sind, resultiert, basiert der diagnostische Wert der Korrelation zwischen Rötlichkeit und Bläulichkeit auf einer beleuchtungsabhängigen Regularität, die durch Unterschiede in der spektralen Bandbreite von Oberflächenreflektanzen vermittelt wird. Während diese Regularität auf der Ebene von individuellen Oberflächen in den ersten beiden Teilen der theoretischen Analyse klar gefunden wurde, ist es weniger klar, ob die Ergebnisse für die Korrelation zwischen Rötlichkeit und Bläulichkeit in dem dritten Teil der theoretischen Analyse anzeigen, dass diese Statistik hilfreich für die Schätzung der Beleuchtung ist:

1. Die Gaußsche Welt prophezeit, dass das reflektierte Licht von Oberflächen mit engen Reflektanzbandbreiten unter Beleuchtungswechseln einer geringeren Veränderung unterworfen ist als das reflektierte Licht von Oberflächen mit breiteren Reflektanzspektren (siehe Seite 32).

2. Entsprechend ist die beleuchtungsabhängige Veränderung in der Simulation der realen Welt größer für bläuliche und graue Oberflächen als für gelbliche Oberflächen, die aufgrund ihrer kleineren Bandbreite veränderungsresistenter sind (siehe Seite 42).

3. Dies führt in der Analyse der Simulation der realen Welt auf der Ebene von Szenenstatistiken zu einem Anstieg der Korrelation zwischen Rötlichkeit und Bläulichkeit, wenn die Beleuchtung sich in Richtung zum rötlicheren 4000 K Tageslicht verändert. Leider hängt diese Statistik ebenfalls von der (beleuchtungsunabhängigen) mittleren Rötlichkeit der Oberflächenreflektanzen der Szene ab, und zwar in einer Art und Weise, die den Nutzen dieser Statistik als Hinweisreiz für die Beleuchtung schwächt: Die Konstellation dieser Statistik zur mittleren Rötlichkeit im retinalen Bild (siehe Abbildung 4.13 auf Seite 53) ähnelt zu einem gewissen Grade der schematischen Konstellation (d) in der Farbabbildung I auf Seite 149, in der zwei Hinweisreize so miteinander korreliert sind, dass die Kombination dieser beiden Hinweisreize nur so hilfreich ist, wie jeder Hinweisreiz alleine (siehe Abschnitt 2.3.2). Auf der Basis der Simulation der realen Welt bleibt es also offen, ob die Korrelation zwischen Rötlichkeit und Bläulichkeit in der natürlichen Umwelt hilfreich ist für die Schätzung der Beleuchtung.

Die experimentelle Evidenz, dass die Korrelation zwischen Rötlichkeit und Bläulichkeit im retinalen Bild vom menschlichen visuellen System für die Wahrnehmung von Oberflächen- und Beleuchtungsfarbe berücksichtigt wird, ist gemischter als die sehr klaren experimentellen Ergebnisse für beide der vorherigen Szenenstatistiken:

1. Der erwartete Einfluss der Korrelation zwischen Rötlichkeit und Bläulichkeit auf die Graueinstellungen wurde nur für vier der sieben Versuchspersonen von Experiment 6 gefunden. Und für drei von diesen vier Versuchspersonen mit signifikanten Ergebnissen ist der Effekt eher schwach.

2. Der Effekt der Korrelation zwischen Rötlichkeit und Bläulichkeit unterscheidet sich wesentlich für die Reize mit verschiedenen chromatischen Mittelwerten, die in Experiment 7 verwendet wurden. Für grüne sowie für gelbe Reize wurde bei allen sieben Versuchspersonen ein individuell signifikanter und substantieller Effekt gefunden, während für blaue und rote Reize kein substantieller Effekt gefunden wurde.

Die Ergebnisse für die grauen Reize, die in Experiment 6 verwendet wurden, können als zwischen den beiden Ergebnismustern aus Experiment 7 liegend betrachtet werden (insbesondere, da dieselben sieben Versuchspersonen an beiden Experimenten teilnahmen). Es scheint daher so, dass der Effekt der Korrelation zwischen Rötlichkeit und Bläulichkeit von der mittleren Chromatizität im Reiz abhängt. Weitere Untersuchungen sind daher wünschenswert, um diese Beziehung genauer auszumachen und mögliche Gründe hierfür zu klären. Auf der

experimentellen Ebene sollte eine größere Anzahl von mittleren Chromatizitäten für die Reize getestet werden. Auf der theoretischen Ebene könnte man z.B. nach möglichen Gründen für diese Beziehung suchen, indem die Simulation der realen Welt aus Kapitel 4 dahingehend erweitert wird, dass Beleuchtungen mit Chromatizitäten außerhalb des Bereichs typischer direkter Tageslichter einbezogen werden. Von besonderem Interesse wären hierbei Lichter, die in der natürlichen Umwelt vorkommen, aber kleinere Bandbreiten haben als die Tageslichter, die in Kapitel 4 verwendet wurden, z.B. durch Blätter gefiltertes Licht in Wäldern.