

Durham E-Theses

The time-course of colour vision

ROBERT JAMES LEE

How to cite:

LEE, ROBERT JAMES (2010) The time-course of colour vision. Doctoral thesis, Durham University.

Use policy

The full-text may be used and/or reproduced, and given to third parties in any format or medium, without prior permission or charge, for personal research or study, educational, or not-for-profit purposes provided that:

- a full bibliographic reference is made to the original source
- a <https://etheses.durham.ac.uk/id/eprint/199/> is made to the metadata record in Durham E-Theses
- the full-text is not changed in any way

The full-text must not be sold in any format or medium without the formal permission of the copyright holders.

Please consult the [full Durham E-Theses policy](#) for further details.

R. J. Lee

The time-course of colour vision

Four experiments are presented, each investigating temporal properties of colour vision processing in human observers. The first experiment replicates and extends an experiment by Stromeyer et al. (1991). We look for a phase difference between combined temporal modulations in orthogonal directions in colour space, which might null the often-claimed latency of signals originating from the short-wavelength sensitive cones (S-cones). We provide another estimate of the magnitude of this latency, and give evidence to suggest that it originates early in the chromatic pathway, before signals from S-cones are combined with those that receive opposed L- and M-cone input. In the second experiment we adapt observers to two stimuli that are matched in the mean and amplitude of modulation they offer to the cone classes and to the cardinal opponent mechanisms, but that differ in chromatic appearance, and hence their modulation of later colour mechanisms. Chromatic discrimination thresholds after adaptation to these two stimuli differ along intermediate directions in colour space, and we argue that these differences reveal the adaptation response of central colour mechanisms. In the third experiment we demonstrate similar adaptation using the same stimuli, measured with reaction times rather than thresholds. In the final experiment, we measure the degree to which colour constancy is achieved as a function of time in a simulated stimulus environment in which the illuminant changes periodically. We find that perfect constancy is not achieved instantaneously after an illuminant chromaticity shift and that constancy of colour appearance judgements increases over several seconds.

The time-course of colour vision

Robert James Lee

Ph.D.

Department of Psychology

Durham University

2010

Contents

Acknowledgements.....	8
Outline	9
1 Introduction	10
1.1 Colour vision	10
1.1.1 Mechanisms of colour vision	10
1.2 Specification of colour	22
1.2.1 Cone-excitation space.....	23
1.2.2 The Macleod-Boynton (1979) Chromaticity Diagram	24
1.2.3 Derrington, Krauksopf & Lennie (1984) colour space (DKL)	25
1.3 Adaptation	26
1.4 Temporal issues	28
1.4.1 Classical S cone signal delay.....	28
1.4.2 Justification and Aims	32
1.5 Colour constancy.....	32
1.5.1 Measuring constancy	33
1.5.2 Constancy models	33
1.5.3 Estimating the illuminant.....	35
1.5.4 Model implementation	36
1.5.5 The time-course of colour constancy.....	37
1.5.6 Justification and Aims	39
2 Hue circle stimuli.....	40
2.1 Mathematical construction.....	40
2.1.1 Generating circular loci.....	40
2.1.2 Cardinal axes	42
2.1.3 Intermediate axes	44
2.2 Discrimination model.....	46

2.3	More complex models	52
2.4	Alternative models.....	57
2.5	Relationship to measured data.....	58
3	Common methods	60
3.1.1	Calibration.....	60
3.1.2	Individual observer calibration – 8-channel.....	65
3.1.3	Individual observer calibration – CRT	67
4	Experiment 1: intermediate axes experiment.....	71
4.1	Abstract.....	71
4.2	Introduction	71
4.2.1	Early chromatic pathways.....	72
4.2.2	Delay of the S-cone signal.....	73
4.2.3	The Peake Effect.....	74
4.2.4	Predictions	75
4.3	Methods.....	79
4.3.1	Apparatus.....	80
4.3.2	Calibration for individual observers.....	81
4.3.3	Stimulus characteristics	81
4.3.4	Experimental procedure	82
4.3.5	Experimental conditions	83
4.4	Results.....	83
4.4.1	‘Cardinal’ axes	85
4.4.2	‘Intermediate’ axes	86
4.5	Models	86
4.5.1	Model description.....	86
4.5.2	Model fits.....	88
4.5.3	Cardinal axes conditions	88

4.5.4	Intermediate axes conditions	89
4.6	Discussion.....	89
4.6.1	Do S-cone signal delays explain the Peake Effect?	89
4.6.2	Comparison to Stromeyer et al. (1991)	90
4.6.3	Comparison to other estimates of the magnitude of the S-cone signal delay	91
4.6.4	Alternative models.....	93
4.6.5	Summary	94
4.7	Appendix	95
4.7.1	Model equations.....	95
4.7.2	Model fit parameters.....	97
4.7.3	Further comments	98
5	Experiment 2: adaptation of central chromatic mechanisms.....	99
5.1	Introduction	99
5.2	Methods.....	102
5.3	Results.....	105
5.4	Discussion.....	109
6	Experiment 3: Adaptation of central and peripheral mechanisms measured with reaction times	112
6.1	Introduction	112
6.2	Methods.....	114
6.3	Results.....	115
6.3.1	Analysis of response times.....	115
6.3.2	Analysis of error rates.....	121
6.4	Discussion.....	125
7	Experiment 4: The time-course of colour constancy	127
7.1	Introduction	127
7.2	Methods.....	129

7.2.1	Observers	129
7.2.2	Apparatus.....	129
7.2.3	Stimuli	130
7.2.4	Generation of chromaticities	132
7.2.5	Procedure and task	133
7.3	Results.....	135
7.3.1	Classification boundaries and achromatic points	135
7.3.2	Colour constancy index.....	137
7.3.3	Linear regression.....	140
7.3.4	Nonlinear model fits	140
7.3.5	Projected colour constancy values	143
7.3.6	Within observer differences between illumination conditions	147
7.3.7	Logistic regression on chromaticity coordinates	148
7.3.8	Logistic regression on distance from boundary.....	150
7.4	Discussion.....	152
8	Conclusion.....	157
8.1	Summary of hue-circle experiments.....	157
8.2	Summary of colour constancy experiment.....	159
8.3	Further experiments	160
9	References	163

Statement of Copyright

The copyright of this thesis rests with the author. No quotation from it should be published without the prior written consent and information derived from it should be acknowledged.

Acknowledgements

I would like to thank all the observers who have given up many hours of time to participate in the experiments in this thesis: Hannah Smithson, Keira Ball, Kathryn Dawson and Wayne Smith. Also, two individuals who participated in pilot experiments: Susan Scrafton and Caroline Jones.

I am grateful for the supportive academic environment of the Psychology department at Durham University, where I have had several interesting and helpful discussions with several members of staff.

My interest for the subject matter began during my undergraduate studies under John Mollon. Without his encouragement and guidance I would never have considered a higher degree and would never have discovered my interest in the field. He has continued to show a great interest in this work and his input to discussions is invaluable.

Qasim Zaidi began much of the work that led to these experiments. His ideas have often been refreshingly different from the suggestions of others, and his alternative viewpoints have always been welcome. I am very glad that he continues to be enthusiastic about these studies.

My doctorate was supervised by Hannah Smithson. I consider myself lucky to have worked with such a brilliant, dedicated scientist. She always gives constructive advice as well as encouragement, and is never afraid to find the humour in any situation. I have thoroughly enjoyed this research, and I am confident I would not have done, to the same degree, with a different supervisor.

Outline

The first chapter of this thesis gives an introduction to the field of colour vision research, and reviews various studies. Specific attention is paid to the issues investigated in later in the thesis. In particular these are: the mechanisms involved in chromatic processing in the early visual pathways, latencies in signal transmission in these pathways, and colour constancy.

Chapter 2 describes some of the technical details of the stimuli we use in the first three experiments. Chapter 3 outlines methods of stimulus generation and apparatus setup that are common to all of the experiments.

Chapters 4 to 7 each describe an experiment. Each chapter contains its own introduction, methods, results and conclusions in a similar format to a journal article. The first, second and fourth experiments are self-contained, while the third is an extension of the second.

In the final chapter, the results and conclusions of all four experiments are summarised and related to the issues described in the introduction. Suggestions for further work to build upon the results presented in this thesis are discussed.

1 Introduction

1.1 Colour vision

The visual system allows us to detect radiation in a range of the electromagnetic spectrum, and use this to interpret our surroundings. Colour is a property of our visual perception, and is derived from the distribution of energy across wavelengths from approximately 380nm to 720nm.

In the 19th century, Thomas Young proposed that there were three types of receptor in the eye, each responsive to different coloured light. This theory was developed by Herman von Helmholtz who proposed that the three receptors were each sensitive to different ranges of wavelengths of light, and different relative excitations of these receptors gave rise to sensations of different colours. Ewald Hering proposed that colour was the result of three opposed mechanisms, in which increased detection of one colour was equivalent to decreased detection of its opposite. An introduction to the history of colour vision research is given by Mollon (2003). We now know both theories are correct: there are three classes of receptor with different wavelength sensitivities and their excitations are compared by opponent mechanisms. While the standard colour vision model has three receptor classes and three second-stage mechanisms, and the physiology corresponding to these mechanisms has been discovered, recent evidence suggests that the model may be over-simplified. All these mechanisms will be summarised in section 1.1.1. This thesis contains several experiments designed to investigate the time-course of information transmission and processing in the human visual system.

1.1.1 Mechanisms of colour vision

Research over the last few decades, often driven by psychophysical results, has provided insight into the physiological mechanisms that form the first stages of colour vision from the retina to the primary visual cortex. Excellent reviews of this research are by Solomon and Lennie (2007) and by Conway (2009). Masland (2001) gives a good review of retinal mechanisms in particular. The important stages relevant to this thesis will be summarised in the following sections. It should be noted here that much of the physiological research leading to our understanding of these mechanisms was not performed on humans. Electrophysiology measurements and in-vitro measurements are invasive and therefore cannot be performed on humans for ethical reasons. Most of the studies referenced below were performed on anaesthetised non-human primates, or retinal tissue extracted from primates. The similarity between the visual systems of old-world monkeys and humans allows us to apply findings from these experiments to human colour vision research.

1.1.1.1 Retina

The physiological mechanisms responsible for the first stage of colour vision are located in the retina. The retina is the tissue on the inner surface of the eye and contains layers of neurons of different types. The outer-most layer of these types contains two types of photoreceptors – cells that signal the absorbance of photons with neural signals. The two types of photoreceptors present in this layer are the rods and cones, although they are not the only type of photoreceptive cells present in the retina (see later).

Rods and cones have different dynamic sensitivity ranges. The responses of rods are saturated at daylight light levels so it is assumed that, under these conditions, the cone (photopic) system is dominant and rods do not contribute to daylight vision. At low light intensities cones are below threshold, so only the rods produce responses relative to their quantal catch. Rods share signal pathways with cones, however, so it is possible that at mesopic light levels (where both rods and cones are within their operating ranges) or that rod signals contribute to colour vision. Some interaction between rod and cone signals has been shown (see Buck, 2004 for review; e.g. Cao, Zele, & Pokorny, 2005) however the degree of cone signal contribution to colour perception is uncertain. While the contribution of rods is not investigated here, rod excitations are controlled whenever possible to minimise any unwanted contribution from their signals.

The central region of the retina, called the macula, contains a yellow pigment. The density of this pigment differs between individuals, and within individuals it changes with age and can be affected by diet (Hammond, et al., 1997; Moreland & Bhatt, 1984). The differing level of filtering by the differing concentrations of this pigment has the effective result of a difference in chromatic sensitivity between individuals at any given point in time.

Other individual differences in chromatic sensitivity arise from differences in the spectral sensitivities of the cones (see below). These factors need to be controlled for when designing experiments. These are some of the reasons why an individual-observer calibration stage is present in most of the experiments that will be described in this thesis.

Other important layers in the retina, moving progressively towards the inner surface, are the outer nuclear layer, outer plexiform layer, inner nuclear layer, inner plexiform layer and the ganglionic layer. The important tissue in each layer will be described next.

1.1.1.2 Cones

Cone cells contain proteins called opsins. These proteins are bound to chromophores, and these chemicals undergo a chemical change when they absorb light. This chemical change triggers the phototransduction cascade, which results in hyperpolarisation of the cell. There are three classes of cone cell in the normal human retina. The main functional difference between the three types is their spectral sensitivity, which is determined by the opsin expressed in the cell. The spectral sensitivity functions, or the relative probability of absorbing a photon as a function of wavelength, measured with behavioural experiments, are shown in Figure 1. These are the estimates of Stockman & Sharpe (Stockman & Sharpe, 2000; Stockman, Sharpe, & Fach, 1999), available in electronic format from the Colour and Vision Research Laboratories at www.cvrl.org. The spectral sensitivity curves of the cones are often referred to as the cone fundamentals, and will be referred to as such in this thesis.




Figure 1. The spectral sensitivity curves of the human cone photoreceptors. Each curve corresponds to one of the three cone classes. The vertical axis represents probability of absorbing a photon at the given wavelength, relative to the probability at the wavelength to which it is most sensitive. The colours of the curves illustrate the chromaticity at the peak sensitivities of each cone.

We refer to the three cone classes as the long-, medium- and short-wavelength sensitive cones, or the L, M and S cones for short. This is due to the wavelengths at which the classes are maximally sensitive. The cone classes are sometimes, often in classical literature, referred to as red, green and blue. This is imprecise, as no real lights will excite only one cone class there is no simple relationship

between excitation of a single cone class and a perceived hue. Furthermore, the peaks of the sensitivity functions do not correspond with these spectral colours.

The feature of three classes of cone with different spectral sensitivities gives rise to trichromacy. We now know that the receptors central to the Young-Helmholtz theory are the three classes of cones with their different sensitivity functions (Dartnall, Bowmaker, & Mollon, 1983). Each cone is itself colour-blind. While the light incident on a cone may differ in intensity and wavelength, the output is the cell's hyperpolarisation, a single linear variable. So a cone cannot distinguish a difference in wavelength of a stimulus from a difference in intensity. For example, an L-cone might output a similar signal in response to a high-intensity 650nm monochromatic light as to a low-intensity 550nm light. The relative responses of the three cone classes of cone are compared to generate perception of hue, but at the level of cones colours are represented in only three signals. The full spectral distribution of the incoming light does not survive in the neural pathway past the cones. It is therefore possible specify a match to any given colour by adjusting the intensities of almost any (but see section 3) three lights, as long as these lights are of different colours and one of them cannot be made by mixing the other two.

Differences between individuals' cone spectral sensitivities are caused by a number of factors. The first of these is a common polymorphism at amino-acid residue 180 of the L-cone pigment (Winderickx, et al., 1992). The majority of Caucasian males have Serine at this location, but in just over a third this is replaced by Alanine. This genetic difference affects the spectral tuning of the L-cone sensitivity (shifting it slightly towards short-wavelengths) and explains individual differences in Rayleigh match settings (the level of a narrowband yellow light adjusted to match a mixture of red and green narrowband lights whose relative levels are also adjusted (Rayleigh, 1881)). Other polymorphisms exist, although they occur less frequently. A review of these genetic variations is given by Deeb (2005). Further differences in relative absorptions of different cone classes occur due to differences in pigment optical density within the cones.

The outer layer of the retina contains hundreds of thousands of cone cells, as well as rod cells. Light must pass through several layers of cells (see Figure 2) before reaching the photoreceptors, and these are packed into a two-dimensional sheet, with the long axis of each cell perpendicular to the sheet. There are almost no rods in the fovea – the region of the retina roughly corresponding to the central 2° of viewing angle. This area is packed with cones, however there are no S-cones in the foveola, the very centre of the fovea (Curcio, et al., 1991). The density of cones drops rapidly with increasing eccentricity (radial distance from the fovea), and the density of rods increases with eccentricity to a

maximum at approximately 20° eccentricity, after which it decreases (Curcio, Sloan, Kalina, & Hendrickson, 1990). The experiments in this thesis concentrate on central vision, so that the results are mainly dependent on cones. It is known that in general the density of S-cones is around 5-10% of all cones, much less than that of L and M-cones and that the ratio of L- to M-cones can vary greatly, although the ratio of L:M-cones is greater than 1:1 in normal individuals (Hofer, Carroll, Neitz, Neitz, & Williams, 2005). Mysteriously, differences in the relative numbers of L and M cones do not lead to differences in colour vision (Neitz, Carroll, Yamauchi, Neitz, & Williams, 2002). The exact spatial arrangement of the cones further than already described is the subject of some debate. The strongest evidence suggests that the position of each cone is random within the distributions described; there is no regular lattice arrangement of cone types (Mollon & Bowmaker, 1992; Roorda & Williams, 1999).




Figure 2. A cartoon representation of a cross section of the retina, showing layers of photoreceptors, bipolar cells and ganglion cells.

As discussed above, an important feature of the cones with regards to colour vision is that one cone or cone class is itself colour blind, it cannot distinguish between a change of intensity of a light source

and a change in wavelength: the principle of univariance. The only response a cone can output is its level of hyperpolarisation and so might respond in the same way to a dim light of a wavelength close to the centre of its sensitivity curve as to a bright light of a wavelength near the tail of the curve. To achieve colour vision, the responses of the different cell types must be compared. This is done by post-receptoral mechanisms, described below, which contact several cones in a small region of the retina. Some types of ganglion cells output signals which receive excitatory input from some cone classes and inhibitory input from others. The resolution at which chromatic information can be distinguished, therefore, is limited by the distribution of the different cone classes in the retina.

It is thought that the L- and M-cones and their pigments differentiated from a common ancestor, and that this happened, along with the differentiation of some later chromatic mechanisms, relatively recently in our evolutionary history (at least compared to the S-cone chromatic signal) (Mollon, 1989). The L- and M-cones are indistinguishable histologically and immunologically (Solomon & Lennie, 2007) and it can be seen from Figure 1 that they have similar spectral sensitivity curves with peaks close together. The S-cones, however, show several morphological differences from the other cone-types and can be distinguished by histological staining (Curcio, et al., 1991).

Of particular interest to this thesis are the differences in the temporal properties of the cones. The L- and M-cones have been shown to respond with slightly different latencies (Tsujiura, Shioiri, Hirai, & Yaguchi, 2000), but there is a body of evidence from psychophysical experiments (Blake, Land, & Mollon, 2008; Smithson & Mollon, 2004) that suggests that the signals from S-cones are delayed relative to those from L- and M-cones. It is important to note here that this suggestion is that there is a latency in S-cone signals, not a reduction in temporal resolution which would mean that high-frequency flicker would be undetectable (see Brindley, Ducroz, & Rushton, 1966; Chichilnisky & Baylor, 1999; Wisowaty & Boynton, 1980 for discussion of this). This latency may be a result of the differences between the classes of cones themselves, or it could arise at a stage later in the pathway of chromatic information from retina to cortex. These temporal differences will be discussed later (see section 1.4.1) and the results of our experiments in section 4 suggest that it may indeed be the cones that introduce the delay, although this is inconsistent with electrophysiology data from Schnapf, Nunn, Meister, and Baylor (1990) who directly measured the responses of photoreceptors.

1.1.1.3 Bipolar Cells

Rod and cone photoreceptors synapse with bipolar cells. These cells, so named because of their morphology, transmit signals from the photoreceptors to ganglion cells. Like the rods and cones, their responses are in the form of graded potentials, as opposed to the action potentials of later

mechanisms. Each bipolar cell may contact just one photoreceptor, and this is often the case in the centre of a midge centre-surround receptive field in the fovea (Dacey, et al., 2000) but is less in the peripheral retina (Boycott & Wässle, 1991; Dacey, 1993). Alternatively, in the periphery of a receptive field, a bipolar cell may combine the responses of a number of photoreceptors.

There are a number of different classes of bipolar cell. The classes differ in the types of photoreceptors they receive signals from, and the classes of ganglion cells they contact. A full account of the many classes of bipolars and their functions is not relevant here. However it should be highlighted that there are distinct classes that receive input from S-cones, and that separate classes signal S-ON and S-OFF (Kouyama & Marshak, 1992; Mariani, 1984). It is possible that these distinct classes of bipolar cell are the source of the S-cone signal delay (see section 1.4.1).

Other cell types in the retina, with synaptic connections to photoreceptors and ganglion cells are the horizontal and amacrine cells, respectively. These types of cell affect the lateral interaction of signals from rods, cones and bipolars, and provide mechanisms for lateral inhibition and feedback mechanisms between photoreceptors and ganglion cells. Again, these are not directly relevant to this work, and will not be described further. Horizontal cells have been shown to not exhibit spectral opponency (Dacey, Lee, Stafford, Pokorny, & Smith, 1996) so are unlikely to form the basis of a chromatic pathway.

1.1.1.4 Ganglion Cells

Ganglion cells occupy the inner layers of the retina, and receive input mostly from amacrine cells and some direct input from bipolar cells. It is currently believed that some of the classes of ganglion cell are the first cells responsible for our ability to distinguish colours, and their outputs result in chromatic opponency. Axons of retinal ganglion cells form the optic nerve, and those of interest in the context of this thesis project to the lateral geniculate nucleus (LGN). While the LGN is often described as a 'relay station', it may be the case that further chromatic processing is thought to take place here (see later) before signals are transmitted to the primary visual cortex (V1). Ganglion cells, the LGN and V1 will be described in turn.

Recent decades have seen the discovery of several classes of ganglion cell, so we now know of more than the handful originally discovered. The classes of ganglion cell differ in their size, morphology, and the layers of the retina in which their dendritic radiations lie. In general, ganglion cells have dendritic trees that extend radially in one or two planes of the inner plexiform layer of the retina. The radius of these dendritic trees determines the spatial extent of the photoreceptors that make contact

with the ganglion cell and therefore the size of the cell's receptive field. The classical ganglion cell classes are the parasol, midget, and bistratified cells.

Parasol cells do not respond much to changes in chromaticity, but to changes in luminance. They receive excitatory input from L- and M-cones. Since their inputs from L- and M-cones are synergistic, these cells are not usually classed as opponent. There is still some debate around whether or not they receive input from S-cones ((Sun, Smithson, Zaidi, & Lee, 2006a, 2006b), but see Chatterjee & Callaway (2002) for further debate). Luminance is usually defined as the sum of the L- and M-cone signals and conventionally does not include an S-cone input (although see Stockman, MacLeod, & DePriest, 1987). Parasol cells, often referred to as m-cells (see later), have relatively small dendritic trees that extend into either inner or outer planes of the inner plexiform layer. Dacey et al. (2003) give a good overview of many ganglion cell types and their morphologies.

Throughout this document, the signs of cone inputs to ganglion and LGN cells, or units in models, are abbreviated. The sign preceding the letter L, M and S indicates the input from that cone type. For example $-S+L-M$ indicates a unit with inhibitory input from S-cones, excitatory input from L-cones, and inhibitory input from M-cones.

The midget ganglion cells (often referred to as p-cells, see later) are generally agreed to be the mechanisms of one fundamental colour (Lennie, Haake, & Williams, 1991) comparison, with both ON and OFF centre-surround types (Dacey, 1993) although this is disputed (Calkins & Sterling, 1999; Rodieck, 1991). Parasol cells have relatively broad dendritic trees that can be at various levels within the inner plexiform layer. The ON and OFF types usually have dendritic trees in the inner and outer portions, respectively, of the inner plexiform layer and with relatively large and small radii. These dendritic trees have relatively larger radii than midget cells, and receive cone input over a larger receptive field. The ON-centre cells receive excitatory input from the cones in the centre of their receptive field and inhibitory input from cones in the surround. In the central region of the retina, the centre of the receptive field often receives input from just one cone. Midget cells receive opposite input from L- and M-cones, and will increase their firing rate to increases in L-cone activity and decrease firing with increased M-cone activity, or vice-versa. It is unclear whether a single cone input to the centre of a receptive field is opposed by a surround field composed of inputs from a single cone type or multiple cone types (Lennie, et al., 1991; Reid & Shapley, 1992).

This L-M-cone opponency forms the basis of what is sometimes imprecisely called the Red-Green pathway (Wiesel & Hubel, 1966). The sensations elicited by stimuli designed to cause polarisations of this mechanism (i.e. large ratios of L:M activity) look red (high L-cone activity, low M-cone activity)

and green (low L-cone activity, high M-cone activity) but these are not perceptually significant reds or greens such as the perceptual 'unique hues' (see Mollon and Jordan (1997) for a discussion of unique hues).

As with bipolar cells, there is a distinct population of cells that carry signals that originate from S-cones, although there is a greater variety of these cells. A second colour opponent pathway begins with cells that oppose S-cone input with a combination of L- and M-cone input. The small bistratified cells have been associated with the S-ON pathway (Dacey & Lee, 1994), but many more types of ganglion cell have been discovered recently, some of which receive S-cone input (e.g. Dacey & Packer, 2003).

The large, sparse bistratified ganglion cells also have an S-ON characteristic response, and are inhibited by L+M input (Dacey et al 2003). These cells would appear to constitute a second S-ON opponent pathway, but the reason for the two different cell types is unknown.

Large, sparse monostратified ganglion cells (Dacey, Peterson, & Robinson, 2002) respond in an S-OFF, L+M-ON manner. They receive their S-cone input from the S-cone bipolars, which are believed to carry an S-ON signal, so this signal must be inverted somehow. An S-OFF specific bipolar cell has not been discovered.

Some ganglion cells are themselves photosensitive, containing the photopigment melanopsin. Not all of these have projections to the LGN, and they have a much slower temporal response and low spatial acuity. Their outputs are used to drive circadian rhythms and the pupil dilatory response (Berson, 2007). A subset of these cells, the giant melanopsin-expressing cells, also receive rod and cone input, show an S-OFF, L+M-ON response pattern, and project to the LGN. These cells may well, therefore, contribute to colour perception. If this is the case they are the second retinal pathway carrying S-OFF signals.

A simple neural correlate of the classical blue-yellow mechanism at this level is difficult to identify. It may be the case that the responses of all of the above S-cone-signal-carrying cells are combined to create a single mechanism, or that there are several mechanisms and the model of a single blue-yellow mechanism is too simple. This is explored more in the experiment in section 4. Further, there is psychophysical evidence for separate S-ON and S-OFF mechanisms (Krauskopf & Zaidi, 1986; Shinomori, Spillmann, & Werner, 1999).

1.1.1.5 Lateral Geniculate Nucleus

The Lateral Geniculate Nucleus (LGN) is a structure in the thalamus, to which axons from retinal ganglion cells project. The LGN has often been described as a 'relay station' between the eye and visual cortex but this view may be a little simplistic. The LGN has a layered structure containing three distinct classes of cells. The magnocellular layers contain cells with relatively large cell bodies, and that electrophysiological recordings have shown to respond quickly to changes in luminance. Their response to purely chromatic changes is of a much lower magnitude (Derrington, Krauskopf, & Lennie, 1984). The parvocellular layers contain cells that respond more slowly, and strongly to changes in chromaticity. In between the magno- and parvocellular layers are koniocellular cells. There is much less agreement on the function of these cells, however cells that respond to input from S-cones are reliably found in these layers.

The designations of these three cell types in the LGN extend to associated cell types in the rest of the visual pathway. So, parasol cells are often regarded as part of the magnocellular or m-cell pathway, midget ganglion cells are part of the parvocellular or p-cell pathway. The koniocellular or k-cell pathway is less well defined.

Many cells in the LGN show responses that are similar to those of ganglion cells in both their receptive fields and chromatic sensitivities. Cells in the parvocellular layer have been shown to respond maximally to similar chromatic changes as midget ganglion cells: to exchanges of L- and M-cone activity (Derrington, et al., 1984; Wiesel & Hubel, 1966) and there are cells which respond to these exchanges with opposite sign: some cells maximum responses are to increases in L cone activity (L-ON/M-OFF cells) while other cells show their maximum response to decreases in M-cone activity (L-OFF/M-ON cells). Cells with an S-ON response have been shown in the koniocellular regions and their response rates increase with increased S-cone activity and decrease with increased L- or M-cone activity (Derrington, et al., 1984; Hendry & Reid, 2000; Hubel & Wiesel, 1966). There are also cells with S-OFF responses in the parvocellular layer (Valberg, Lee, & Tigwell, 1986), but their responses are not always simple inversions of S-ON responses as with the L/M responsive cells. The chromatic tuning of S-OFF LGN cells is more heterogeneous, with some cells receiving inhibitory input from S-cones and excitatory input from L- and M-cones and some more recently discovered cells receiving inhibitory input from S- and M-cones and excitatory input from only L-cones (Tailby, Solomon, & Lennie, 2008). In fact, in this recent study, the $-S+L-M$ cells slightly outnumbered the $-S+L+M$ cells. There is also, when compared with cells with other chromatic tunings, a larger distribution in the weightings of inputs from different cone classes.

The asymmetry between S-ON and S-OFF mechanisms in terms of the morphologically different cells found in the retina and the response tuning found there and in the LGN is justification for re-examination of the classical colour-opponent model of signals from three cone classes combined into just two colour opponent mechanisms and a luminance mechanism. It is not yet clear whether the differences in physiological mechanisms at these stages result in differences in perception. Furthermore, since there are a range of morphological properties of the cell types carrying S-cone signals in particular, it stands to reason that these different cell types might have different temporal response characteristics which may affect perception. As already noted and as will be discussed further later, it has been suggested that there is a delay in S-cone signal transmission to the cortex relative to L- and M-cone signals. It is possible that any one or more of the cell types described in the retina or LGN might be responsible for this delay, and that asymmetric mechanisms have different temporal properties. Also, these different mechanisms may have different adaptation characteristics (adaptation is described in section 1.3), possibly leading to different magnitudes and time courses of adaptation to different chromaticities.

1.1.1.6 Primary visual cortex and beyond

Neurons in the lateral geniculate nucleus project to the primary visual cortex, often abbreviated to V1, or referred to as striate cortex because of its layered structure. Different layers appear to receive input from different sources. This area is located at the posterior pole of the occipital lobe. Although much has been learned about the chromatic tuning and organisation of cells in V1, there are still many open questions about the number, tuning and response characteristics of chromatic mechanisms in V1 and beyond.

Receptive fields and response tunings of cells in V1 have been studied in a similar manner to those in the LGN. The response magnitudes of V1 cells can be well described in terms of linear combinations of cone responses (DeValois, Cottaris, Elfar, Mahon, & Wilson, 2000). The different groups of LGN cells project to different layers in V1: magnocellular LGN cells project to V1 layer 4C α , parvocellular cells to layer 4C β and koniocellular cells to the lower region of layer 3 and to layer 4A. S-ON and S-OFF cells still appear to retain some segregation at this point, with most S-OFF cells found in layer 4A and S-ON cells showing a wider distribution extending into layer 3 (Chatterjee & Callaway, 2003). The segregation of chromatic mechanisms must end at some point in V1, however, since Lennie, Krauskopf & Sclar (1990) and DeValois et al. (2000) found a much broader distribution of chromatic preferences.

Cells in the upper layers of V1 are organised into ocular dominance columns – cells in a given column receive input from only one eye. Within these columns, groups of cells react differently to cytochrome oxidase staining and ‘blobs’ of cells with high staining reactivity appear in the centre of these columns, in between ‘interblobs’. Livingstone and Hubel (1984) showed that cells in these blobs were chromatically sensitive and project to thin striped in area V2 that stain in the same way. V2 has been shown to contain chromatically responsive cells (Solomon, Peirce, & Lennie, 2004), and one study even shows adjacent cells tuned to spectrally adjacent chromaticities (Xiao, Wang, & Felleman, 2003). Area V4 (in Macaques) was once considered an area dedicated to colour processing (McKeefry & Zeki, 1997), and Zeki et al. (1973) found that all the cells they recorded from in this area were chromatically tuned. However, the area has since been shown to be involved in other aspects of visual processing (Schiller & Lee, 1991). The anatomical organization and human homologue of V4 is debated, but it is located within the posterior inferior temporal cortex (PIT) and some studies refer to this instead of V4 specifically. Conway et al. (Conway, Moeller, & Tsao, 2007; Stoughton & Conway, 2008) found that most neurons in this area were responsive to colours corresponding to the perceptual ‘unique hues’. These chromatic directions are not aligned with the cardinal directions defined by the responses of cells in the retina and LGN, and so cells here must be receiving inputs from all three cone types and both of the early chromatic-opponent mechanisms. Cells in PIT are clustered in ‘globs’ and it is these that show the most chromatic tuning. Conway and colleagues found that most of the cells in globs were chromatically tuned independent of luminance, but ‘interglob’ cells were not. Hadjikhani, Liu, Dale, Cavanagh and Tootell (1998) show, with an fMRI study, that the corresponding colour sensitive area in humans is not anatomically correspond to V4, but to another retinotopic area that they term V8 and that this area also responds selectively to coloured after-images.

The receptive fields of many cells in V1 show evidence of convergence of cone signals. The majority show preference for edges, and show little response to spatially uniform stimulation. This is achieved by cells with double-opponent receptive fields – two cells that combine cone signals in opposition with opposite signs and have adjacent spatial receptive fields (Conway, 2001; Johnson, Hawken, & Shapley, 2001; Kentridge, Heywood, & Weiskrantz, 2007; Livingstone & Hubel, 1984). For example, a cell combining the responses of a +L–M cell adjacent to a –L+M cell would respond maximally to an edge running between the two receptive fields and defined by red-green contrast. Some of these cells show orientation selectivity. Double-opponent cells have been found that show combinations of cone signals that leave them tuned to chromatic directions other than the cardinal axes defined by the cells in the LGN shown by Derrington et al. (1984) and therefore requiring input from all three cone classes

(Horwitz, Chichilnisky, & Albright, 2005). They include cells with L vs. M+S input like LGN cells more recently described by Tailby et al. (2008).

1.1.1.7 Colour-opponent mechanisms

As mentioned earlier, Hering (translated in Hering, Hurvich, & Jameson, 1964) proposed a theory of colour vision in which responses three independent systems give rise to the perception of hue. In this theory, pairs of colours are associated with the extremes of these bipolar mechanisms. Red and green are opposed to form one of these mechanisms, blue and yellow are opposed to form another and black and white are opposed to form the third. This is based on perceptual observations: it is not possible to create, for example, a yellowish-blue or greenish-red hue. Mixtures of lights with chromaticities at opposite extremes of these dimensions result in neutral hues. Hurvich & Jameson (1957) measured, with colour-matching experiments, the amount of light at one wavelength that was required to cancel light of another and provided quantification of Hering's opponent colour theory. Boynton, Ikeda and Stiles (1964) give good evidence of excitatory and inhibitory interactions between chromatic detection mechanisms that support colour opponency. Opponent mechanisms are discussed from an information processing perspective by Buchsbaum and Gottschalk (1983).

Mechanisms that colour opponency have been discussed, but it is important to consider their explicit role in opponency. Cells in the LGN (see section 1.1.1.5) appear to have opponent response properties that are similar mechanisms predicted by to perceptual observations (DeValois, Abramov, & Jacobs, 1966; Hubel & Wiesel, 1966). Later work established distinct groups of neurons (Derrington, et al., 1984) that respond to different combinations of cone inputs in an opponent fashion. Cone opponency arises as early as the retina, where different classes of ganglion cells combine cone signals with different signs (see section 1.1.1.1). It is now clear that the chromatic directions to which these neurons are tuned are not perceptually significant – they do not correspond to the unique hues.

1.2 Specification of colour

Throughout this thesis, it will often be necessary to specify the colours presented to an observer, or the neural excitations we intend to produce by such stimuli. There are a number of different ways in which this could be accomplished.

As discussed above, sensations of colour depend on the spectral composition of the light reaching the eye as well as the state of adaptation of the mechanisms responsible for encoding this information. However, a representation of this spectral distribution does not survive past the first stage of vision at the photoreceptor level. As discussed earlier, the three cone classes (long-, medium- and short-wavelength sensitive, or L, M and S) have different spectral sensitivities, but their output is univariant

and each cone is inherently colour-blind: each cone outputs a one-dimensional neural potential and cannot signal the difference between a change in intensity and a change in wavelength. For example a cone might respond identically to a high intensity light of a wavelength at either tail of its sensitivity curve and a low intensity light of a wavelength at the peak of its sensitivity. Changes in wavelength composition and changes in intensity are ambiguous to a single cone type. Information is carried as excitations of these and later mechanisms and it is the comparison of relative cone excitations that allows us to distinguish hues. It is therefore often unnecessary and impractical to specify the whole spectral distribution when discussing an experimental procedure (although this is required when considering stimuli representing an illuminated surface, where the chromaticity is dependent on the spectral power distribution of the illuminant and reflectance function of the surface, where both may change, as in the experiment in section 7). One consequence of having only three univariant cone outputs is that two different spectral distributions may result in the same levels of relative excitation in the three cone classes: these distributions are said to be metameric. A further important consequence is that it is possible to generate colours by specifying the desired relative excitations of the three cone classes by adjusting the relative output levels of just three primary lights. This is how coloured stimuli are generated throughout these experiments (details of stimulus generation will be given later) and the stimulus could be represented by three numbers: the power output levels of the primaries with given emission spectra. This is a common way of specifying colour on computer display systems which use red, green and blue phosphors with known spectral outputs as their primaries: any different combination of red, green and blue (RGB) outputs leads to a different colour.

Many colour spaces, in which each colour is represented by a coordinate in a (two- or three-dimensional) space, have been defined and each is suited for a particular application such as display systems or printing. For experiments in human colour vision such as the ones reported in this thesis, where we seek to investigate how information is encoded and represented in mechanisms of the visual system, a colour space whose coordinates are based on responses in the human colour vision mechanisms is most appropriate. In such a space, each point represents a different combination of cone signals and therefore a different colour. These systems will be used to represent chromatic stimuli throughout the rest of the document, so it is important to define them now. Coordinates in these spaces are converted into output levels of primary lights in order to generate stimuli (see section 3).

1.2.1 Cone-excitation space

Perhaps the simplest way of specifying colour in terms of neural mechanisms is to specify the excitations of the three cone types. This means that any visible light can be represented as a point in

a three dimensional space whose axes are defined as the responses of the three cone types. The distance from the origin to the projection of the point along a given axis represents the quantal catch rate of the corresponding cone class. All points on a straight line passing through the origin will have the same ratio of cone excitations and therefore the same chromaticity. The distance from the origin along this line represents the absolute quantal catch in the cones. Points in this space having equal luminance (equal summed L and M cone catches) lie in a plane parallel to the S-cone axis and passing through the L and M axes at the same value (i.e. $L + M = k$, where k is a constant).

1.2.2 The Macleod-Boynton (1979) Chromaticity Diagram

MacLeod and Boynton (1979) chose such a plane of equal luminance in which to represent colours. This space has become widely used in vision science to define chromaticity independent of luminance. In the Macleod-Boynton diagram, the horizontal axis represents the ratio of L and M cone activity: $L/(L+M)$, which could also be described as the proportional L cone contribution to luminance. The vertical axis represents S cone activity, normalised by luminance: $S/(L+M)$. When cone excitations are calculated, the L- and M- cone fundamentals are linearly scaled so that their sum is equal to the photopic luminosity sensitivity function, $V(\lambda)$. This will be discussed further in section 3.1.1.

The limits of the horizontal axis in this diagram are defined from 0 to 1 since the axis represents the fractional contribution of the L-cone signal to luminance, however there is no such definition on the scaling of the vertical axis. A common convention, and one that is used throughout this thesis is to scale the vertical axis so that the point on the spectrum locus with the greatest $S/(L+M)$ value lies at $S/(L+M) = 1.0$.

A diagram of this space is shown in Figure 3, with the locus of monochromatic lights plotted, as calculated from the Stockman & Sharpe (Stockman & Sharpe, 2000; Stockman, et al., 1999) cone fundamentals. The chromaticity of all real lights can be represented by a point on or inside this curve. Also plotted is the chromaticity of a light metameric to equal-energy white (EEW). That is, the coordinates in this space of a light which produces the same relative excitation in the three cone types as a light with a flat spectral distribution in the visible range. This is the definition of 'white' that is used throughout these experiments.




Figure 3 The MacLeod-Boynton (1979) Chromaticity diagram. The black curve represents the chromaticities of monochromatic lights, and the cross indicates the position of equal-energy white.

A useful property of this diagram is that a mixture of any two lights will have a chromaticity which falls on a line joining the two chromaticities of the two component lights. The distance of the mixture from each of the components is proportional to the inverse of the relative luminance of the component.

The attractive feature of this representation is that the axes correspond to the excitations in the classical opponent colour mechanisms discussed earlier. Stimuli which only differ in the amount of excitation they produce in the $L/(L+M)$ or 'red-green' mechanism will lie along a horizontal line in this space, and stimuli only differing in the amount of excitation they produce in the $S/(L+M)$ or 'violet-yellow' mechanism will lie along a vertical line. Many of the experiments in this study involve stimuli intended to modulate the responses in these cardinal mechanisms so it is intuitive to be able to conceptualise these modulations as orthogonal modulations in a two-dimensional space.

1.2.3 Derrington, Krauksopf & Lennie (1984) colour space (DKL)

Derrington, Krauksopf & Lennie (1984) expanded the chromaticity diagram of MacLeod & Boynton (MacLeod & Boynton, 1979) into a three-dimensional space. Their experiment involved electrophysiological measurements of cells in the LGN of macaques, and showed that the responses of cells in this structure were tuned to chromatic directions parallel the axes of this space. They defined two of their axes to be a 'constant blue' axis, $L-M$, equivalent to the $L/(L+M)$ axis of the MacLeod-Boynton Diagram, a 'constant red and green' axis, $S-(L+M)$, equivalent to the $S/(L+M)$ axis of

the MacLeod-Boynton Diagram. They added a third axis corresponding to luminance, and specify that in any plane orthogonal to this axis, the sum of signals from the L- and M-cones is constant. Any such plane is equivalent to the MacLeod-Boynton Diagram.

Since the definition of the L-M axis in this space is not a ratio, there are no defined limits on the axes, as in the MacLeod-Boynton diagram. In this space, the origin of the L-(L+M) and S-(L+M) axes is defined to be achromatic, and coloured lights are positive or negative excursions from this point in an equiluminant plane. Increasing luminance moves a light higher in the space, decreasing luminance moves it lower. When stimuli are described in this space in this document, white is assumed to be equal-energy-white as described above. This configuration is shown in Figure 4.




Figure 4 The colour space proposed by Derrington, Krauskopf & Lennie (1984) in which the three axes correspond to the tuning of cells in the macaque LGN. An example of an equiluminant plane is shaded.

1.3 Adaptation

The concept of adaptation (often called habituation) is a property of many sensory systems, and indeed affects many aspects of vision. It is the reduction in response of a mechanism in an inverse relationship to the activity in that mechanism, which then affects later perception that relies on that

mechanism. It can be measured in a variety of ways including direct measurement of response amplitudes of neurons through electrophysiology, or psychophysical measurements of increased thresholds or altered perception. Classic examples of this effect in the visual system include adaptation to spatial frequency (Blakemore & Campbell, 1969) and motion (Barlow & Hill, 1963) and these often result in easily-visible after effects.

Adaptation in colour vision has been widely studied, and has revealed details of the colour vision system. Von Kries adaptation (von Kries, 1878) can be likened to applying a gain control to each receptor mechanism at each retinotopic position. This could easily be accomplished by the cones themselves, which might adjust their output based on their quantal capture rates, however it could be achieved by any mechanism or mechanisms later in the system. It has been suggested (Ives, 1912) that Von Kries adaptation is a component of colour constancy (see section 1.5.2).

Simple, short-term adaptation to colour can easily be seen by viewing a stimulus for several seconds, and then replacing that stimulus with an achromatic field. The field will appear to be the complimentary colour to that of the first stimulus, and this can last several minutes. Perhaps the most important experiment involving chromatic adaptation with regard for sections 4, 5 and 6 of this thesis is that of Krauskopf, Williams & Heeley (Krauskopf, Williams, & Heeley, 1982) who showed that adaptation to one of what we now consider the cardinal directions of colour space (Derrington, et al., 1984) reduced sensitivity primarily to chromaticities in that direction only and this likely occurs in early independent opponent mechanisms which combine responses of cone types. This study is often given in support of the model of colour vision based on two opponent mechanisms, although their data do imply (Krauskopf, Williams, Mandler, & Brown, 1986) that other mechanisms tuned to other directions are adapted as well. Webster & Mollon (1994), adapt observers to similar chromatic stimuli in a matching experiment with suprathreshold test stimuli, and also conclude that there are more than two adaptable chromatic mechanisms (plus luminance).

Some aspects of visual adaptation are retinotopically local – one area of the retina can have a different adaptation state to another (He & MacLeod, 1998; MacLeod & He, 1993). Other effects can be seen in the vision of either eye after only one was adapted, implying that the adapted mechanisms are at some cortical site after the images from both eyes are combined.

The McCollough effect (McCollough, 1965) is an example of a contingent adaptation effect and given as evidence for mechanisms that are selective to more than one attribute of the image. In this case, the adaptation is to a colour-orientation combination, and can make achromatic gratings appear coloured if and only if they have the same orientation as the adapting stimulus.

1.4 Temporal issues

A theme of the experiments in this thesis is that of the time-course of the mechanisms underlying colour vision. The first experiment (section 4) investigates a difference in the latency of signals originating from the S-cones. This latency is exploited in the second and third experiments (sections 5 and 6), which use temporally modulated stimuli in an attempt to adapt central mechanisms that receive their input from the classical opponent mechanisms. In the final experiment (section 7) we investigate whether colour constancy (see section 1.5) relies on mechanisms which require a measureable amount of time to achieve constancy.

1.4.1 Classical S cone signal delay

As previously discussed, there are many physiological differences between the neural mechanisms responsible for the transmission of signals from short-wavelength sensitive cones (S-cones). It has often been suggested that there is a latency in the S-cone signal pathway, meaning that signals from S-cones arrive later, relative to signals from L- and M-cones, at some site at which the signals from the three cone types are combined. The various types of evidence for and against this delay will be discussed.

1.4.1.1 Reaction times

Much of the evidence for delayed S-cone signals comes from psychophysical studies. In particular, several studies have demonstrated that reaction times (RTs) to stimuli that isolate the S-cone pathway are longer than reaction times only modulating other chromatic mechanisms. Mollon & Krauskopf (1973) demonstrated that stimuli encoded by Stiles' π_1 mechanism (see Wyszecki & Stiles (2000) for details of these mechanisms) generate longer RTs.

McKeefry, Parry & Murray (2003) measured reaction times for stimulus onsets in several chromatic directions in the equiluminant plane. They found longer RTs for chromaticities closer to the tritan direction and a difference between S-opponent-isolating and L-M-opponent isolating stimuli of up to 300ms, depending on cone contrast. They also found that RTs to stimuli in the +S direction were longer than those in the -S direction. RTs for chromaticities in the +L-M and -L+M directions were comparable. This not only suggests that S-cone signals are delayed, but also provides further evidence for asymmetric S-ON and S-OFF mechanisms as discussed in section 1.1.1.

Smithson & Mollon (2004) measured reaction times for stimuli varying along S-opponent and L-M opponent directions while specifically attempting to isolate each observer's tritan line and to avoid intrusion from luminance signals. They found an S-cone signal latency of 20-30ms.

All of these studies show relative latency differences between signals from stimuli isolating S-cones and stimuli of other chromatic directions. However, they all provide different estimates of the magnitude of this latency difference. Furthermore the relative delay can be different for manual, saccadic or perceptual responses (Bompas & Sumner, 2008), indicating that delays are injected after the divergence of the pathways that support these different tasks.

1.4.1.2 Other psychophysics

Stockman, MacLeod & Lebrun (1993) suggest that a sluggish temporal response arises in a post-receptoral mechanism. Using a psychophysical task that bypasses the need for direct perception of the S-cone signal, they have shown that, under extreme L- and M-cone adaptation, rapid S-cone flicker can produce visible beats if superimposed on L- or M-cone flicker. Under these conditions, they suggest that the S-cone signal is transmitted by a “luminance” pathway and the magnocellular pathway is a possible, though controversial, candidate. Stockman, MacLeod & Depriest (1991) give a similar explanation for their finding that S-cone signals can null L-cone flicker at high frequencies, and at high levels of S-cone adaptation. These studies measure frequency response as opposed to phase response, or the ability to encode a high frequency flicker as opposed to a delay in signal transmission, but suggests that the S-cones have similar temporal response characteristics to the L- and M-cones. This is inconsistent with the results of the experiment in section 4 of this thesis, where we conclude that the relative delay between S-cone signals and L- and M-cone signals arises early in the pathway, possibly in the receptors themselves.

Psychophysical measurement of temporal impulse response functions (IRF) reveals a slower time course for S-cone signals than for luminance signals. Consistent with the physiological differences in S-ON and S-OFF pathways, the IRF for S-cone increments and decrements differs, with faster responses to increments (Shinomori & Werner, 2008).

Stromeyer et al. (1991) cite delayed transmission of the S-cone signal as an explanation as to why modulations in different directions around identical loci in colour space were not indistinguishable. They were able to counteract this, making the stimuli indistinguishable, by adjusting the relative phase difference of the S-cone modulation in the stimulus. This study inspires many of the experiments in this thesis and will be discussed in more detail later (see section 4).

In all of these experiments, the measured temporal responses depend on the adaptation state of the mechanisms involved. Stockman, Langendorfer, Smithson & Sharpe (2006) show that cone sensitivity is traded for speed, which results in faster flicker responses at higher adaptation levels. In Stromeyer

et al.'s (1991) experiment they measure phase delay at different adapting intensities and show a relationship.

1.4.1.3 Electrophysiology

In a variety of electrophysiology studies, measurements of photoreceptors, ganglion cells, LGN and cortical units have been made. However, few of these give clear evidence for a delayed S-cone signal. Schnapf, Nunn, Meister, and Baylor (1990) made direct measurements of the temporal response of macaque photoreceptors. While they did not study a large number of S-cones, they didn't find any temporal or sensitivity differences between them and the L- and M-cones.

Using silent substitution, Yeh, Lee and Kremers (1995) measured the temporal characteristics of signals of the three cone-types at the ganglion cell level. They found similar temporal modulation transfer functions for +L–M, +M–L, and +S–(L+M) cells, for both excitatory and inhibitory cone inputs. These results imply that the temporal response of the S-cones themselves is not particularly poor and that limitations must be imposed by post-receptoral processing. This is consistent with the conclusions of Stockman et al. above, and with Chatterjee and Callaway (2002) who claim a small (10%) but consistent S-cone input to magnocellular neurons in Macaque LGN. However, Sun, Smithson, Zaidi & Lee (2006a, 2006b) find no significant S-cone input to magnocellular and parvocellular ganglion cells, at least under neutral adaptation.

Tailby et al. (2008) measured the temporal frequency selectivity of S+ and S- LGN cells that were driven either by S-cone isolating modulation or by achromatic modulation. The peak sensitivity for both cell-types was around 3Hz when driven by S-cone isolating modulation, but around 7Hz when driven by achromatic modulation. They argue therefore that the poor temporal resolution of the S-cone pathways arises early in the retina. However, they find no evidence for an accompanying *phase delay* of the S-cone signal, relative to the L- and M-cone signals provided by parvocellular cells.

It could be that the S-cone signal delay arises at a stage after the LGN, and recordings from cortical neurons may reveal this. Cottaris and De Valois (1998), recording from cortical area V1 in the Macaque, found that S-opponent signals were available only after 96–135 ms, whereas L/M-opponent signals were available after 68–95ms. They additionally found that cells with late S-cone inputs dynamically change their chromatic tuning over time. They suggest that the weak S-cone signal is recursively amplified and delayed in V1 before being combined with L/M-opponent signals, and this process results in the sluggishness of the signal and change in chromatic tuning.

1.4.1.4 VEPs

Studies of visual evoked potentials (VEPs) also suggest longer latencies of tritan-aligned signals, relative to L/M signals. Robson and Kulikowski (1998) measured a mean difference of 55ms and Rabin, Switkes, Crognale, Schneck, & Adams (1994) make a similar measurement. In both studies, S-cone isolating stimuli generate EEG waveforms whose shape is distinct from that of L/M stimuli, which might be taken as further evidence for functionally different processing channels.

1.4.1.5 The Peake Effect

The temporal difference described above has been used to explain an interesting phenomenon, first demonstrated with stimuli that are early examples of the type used in experiments in this thesis (sections 4, 5 and 6). We refer to this phenomenon as the Peake Effect since it was originally discovered by a Miss Olive E. Peake, and reported by Hartridge (1949). It was also later reported by Davidoff, Aspinall, & Hill (1978) and Hill, Rodger, & Smalridge (1980), and again by Spillmann and Neumeier (1984) and summarised again in Stromeyer et al. (1991).

In the original description, coloured papers were placed around the outside of a tube, so that when the tube was rolled along a tabletop, the colours appeared in spectral sequence. When the tube is rolled fast enough so that the individual colours are indistinguishable, then the predominant colours seen are different when the tube is rolled in one direction or another.

In Spillmann and Neumeier's (1984) version, they constructed a disk with twenty radially divided segments. Each segment was filled with a Munsell paper of constant chroma and value, but different hue. The hues were evenly distributed around the Munsell colour sphere. This is represented in Figure 1 in Stromeyer et al. (1991). When the disc was spun in different directions, the effect was similar to that of Peake's tube: the predominant colours seen are different. They report that the effect is present when the disc is spun at 7-9Hz and illuminated with white light. They report that the colours seen flicker between green and purple when the disc is spun clockwise, and orange and turquoise for counterclockwise rotation. They discount motion as a cause for this effect, as the effect remains when a small portion of the tube is viewed through a tube. Therefore, the reversal of the sequence of chromaticities is the important feature of this stimulus.

Several explanations of this effect have been given, and these assume that the neural signals transmitting chromatic information are distorted between photoreceptors and a stage at which a perceptual comparison can be made, in such a way as to render the two loci of colours different.

One explanation comes from asymmetries in temporal masking between opponent mechanisms that are unequally distributed around the hue-circle, so that the colour mechanisms that are masked are different when viewing CW or CCW stimuli (Spillmann, 1990; Spillmann & Neumeier, 1984). In this model, two of four mechanisms are masked when viewing the stimuli, but the pair that are masked are different for CW and CCW stimuli. This results in two colours seen more than others, and these colours are different for opposite modulation directions.

Another popular explanation comes from differential latencies between colour mechanisms, possibly at a stage as early as the cones (Hill, et al., 1980). Stromeyer et al. (1991) explained the effect by suggesting a temporal delay in the mechanism carrying signals from S-cones and provided evidence for this with their experimental results. Changing the rotational direction of the stimulus reverses the phase difference between modulations of the mechanisms carrying S and L-M signals, and a delay in the S-cone signal mechanism (evidence for this is discussed above) would add an additional phase difference which would render the sequence of hues different. We accept this explanation for the purposes of this thesis, and base one experiment on that of Stromeyer et al. (1991) we use a stimulus designed to have similar temporal and chromatic properties that are used in these examples. Stromeyer et al.'s experiment is described in section 4.2, and the hue-circle stimuli that they, and others, use are described more in section 2.

1.4.2 Justification and Aims

This delay of the S-cone signal gives rise to the first three experiments in this thesis. Firstly (section 4), we replicate Stromeyer et al.'s (1991) experiment, making our own measurements of the delay. As discussed, the stage in the S-cone signal pathway at which the latency is introduced is a subject of debate, and with a modification to the experiment we aim to pin-down the neural locus of the delay. We conclude that the latency arises early in the pathway, and then (section 5) design an experiment exploiting this delay to investigate another debated subject: colour vision processing after the peripheral opponent mechanisms. Using a stimulus described later (section 2) we seek adaptable mechanisms that must be sited after the combination of signals from the cardinal opponent mechanisms, since they are tuned to chromatic directions other than those of the cardinal mechanisms. We measure adaptation, assumed to be in these mechanisms, in two ways: with threshold measurements (section 5) and reaction-time measurements (section 6).

1.5 Colour constancy

In the natural world, the spectral composition of light reflected from the surface of an object to the eye often changes. Usually, the spectral reflectance functions of the surface remains constant, but

the spectral composition of the light incident on it varies for a number of reasons. The light reflected from the surface has a spectral composition determined by the surface reflectance and the illuminant, and will also vary with direction. However, in most cases, the colour appearance of the object remains the same, and it is this aspect of perceptual constancy that is the focus of the experiment in section 7. The experiment in section 7 of this thesis is primarily concerned with the time course of colour constancy, which has not been the subject of much investigation since Helson (1947) postulated the level-of-reference theory, in which stimuli are judged against an average of those seen in the period before and which applies to psychophysical stimuli in general. A very brief summary of the relevant theory and the problems associated with how constancy might be achieved will be given here, before temporal issues are discussed. Comprehensive reviews of the problems and possible solutions to colour constancy are given by Smithson (2005), Maloney (1999) and Hurlbert (1998).

1.5.1 Measuring constancy

There are a number of suggestions of ways to measure colour constancy in experimental situations. Firstly, the observer could be tasked to match colours seen under different illuminants (asymmetric colour matching) when these stimuli are presented either side by side (simultaneous matching (Arend & Reeves, 1986; Arend, Reeves, Schirillo, & Goldstein, 1991; Brainard, Brunt, & Speigle, 1997)) or one after the other (successive matching (Brainard & Wandell, 1992)), or even in the visual fields of different eyes (McCann, McKee, & Taylor, 1976). Secondly, the observer's task could be to adjust a test reflectance in a scene with a given illuminant until it appears white (or some other internal standard such as "neither red nor green" (Brainard, 1998; Fairchild & Lennie, 1992)). Thirdly, the observer can be asked what name a reflectance is given under different illuminants, and the probability of a surface being assigned the same name can be measured (Troost & Deweert, 1991). Most of these experiments can be carried out with real stimuli or ones generated artificially on a computer screen. It has been suggested that all of these methods are not actually measuring colour constancy – specifically how similar the same surfaces appear under different illuminants – but a simpler measure such as the relative differences between surfaces, or between surfaces and illuminant (Foster, 2003). Different results can be obtained by giving the observer different instructions, particularly in matching experiments: either to match surfaces so that they appear to belong to the same material, or to match hue and saturation (Arend & Reeves, 1986).

1.5.2 Constancy models

There are several plausible models of the mechanisms by which constancy is achieved. Most assume that the visual system recovers some information about the reflectances of surfaces, and discards

contributions from the illuminant, or estimates the illuminant spectrum in order to discount the effect on reflectances. To do this, the transformation in neural signals must be understood. Since chromatic information is reduced to ratios of three cone signals at the photoreceptor stage (see section 1.1.1.2) it is interesting to examine the effects of changing the illuminant on the cone excitations of reflectances. For each of the sets of L-, M- and S-cone signals elicited by a collection of surface reflectance, changing the illuminant has the effect of applying a roughly constant multiplication factor to each coordinate (Dannemiller, 1993; Foster & Nascimento, 1994; Nascimento, Ferreira, & Foster, 2002; Zaidi, Spehar, & DeBonet, 1997). This consistent scaling extends to later stages in the traditional colour processing model where chromaticity can be expressed in coordinates of the Macleod-Boynton chromaticity diagram, which represents opponent chromatic signals in retina and LGN. The $S/(L+M)$ coordinates undergo a similar multiplicative scaling with a change in illuminant, however it has been suggested that the $L/(L+M)$ coordinate is translationally transformed: an approximately constant amount is added to each coordinate (Zaidi, et al., 1997), although this may be over-simplified.

The set of multiplication factors that transform the cone coordinates of a given illuminant to the coordinates of an equal-energy illuminant also transform the coordinates of surface reflectances to their values under the equal-energy illuminant, and this is referred to as the Ives transform (Brill, 1995). Mathematically, each vector of three cone coordinates is multiplied by the same diagonal matrix, whose three elements are determined by the cone coordinates of the illuminant. This is a similar concept to many models of chromatic adaptation, however the same transformation would need to be applied across the whole scene, or area of the scene under a given illuminant, and so the transformation matrix must contain information about the illuminant chromaticity. This is therefore slightly different from von-Kries adaptation, as discussed in section 1.3, in which cone signals at each retinal location are adjusted.

The rank order of any of the cone signals is not significantly altered under an illuminant change. The relative magnitudes of any of the coordinate values are also approximately preserved, such that any given coordinates retain their relative rank. So, colour constancy could also be achieved by encoding the rank orders of cone excitations, as these are invariant under an illuminant shift, and are equivalent to encoding the rank order of surface reflectance chromaticities.

Alternatively, encoding cone excitation ratios would at least allow an observer to distinguish between a change in illuminant, which doesn't change relative cone excitation ratios, and a change in material reflectance, which does (Craven & Foster, 1992; Foster & Nascimento, 1994) and it has been shown

that observers can make this distinction, even when both changes are simultaneous (Foster, et al., 2001). Comparing cone excitation ratios to a spatial average might allow colour constancy even when the spatial arrangement of materials in a scene changes (Amano & Foster, 2004).

1.5.3 Estimating the illuminant

Some models of colour constancy, in particular those implementing an Ives transform, require that cone coordinates of the illuminant, or the relationship between test and reference cone coordinates, is estimated. There are several theories of how this may be accomplished. Firstly, this could be done by comparing a reference surface, or set of reference surfaces, under two illuminants. If one surface known to be constant is seen under the two illuminants, the matrix required for the Ives transform can be estimated directly, assuming the surface can be identified in a complex scene and represented across an illuminant change. This could also be achieved with sets of surfaces seen under the two illuminants, when either the whole set or a subset of surfaces remain the same (Zaidi, 1998).

A second suggestion is that the chromaticity of the illuminant is estimated from the mean chromaticity of the whole scene (Buchsbaum, 1980; Land, 1983). This is commonly known as the 'Grey world' hypothesis and assumes that the mean of the reflectances in the scene is spectrally uniform, so the mean chromaticity corresponds to that of the illuminant. This is essentially an extension of von-Kries adaptation to include the whole scene. A scene with uniform mean spectral reflectance is unlikely in the real world (Webster & Mollon, 1997), but constancy could be achieved if the mean reflectance remained constant.

Another suggested model is that correlations between luminance and chromaticity could provide an estimate of the illuminant, although this would be most reliable for only the brightest objects (Golz & MacLeod, 2002).

A suggestion that originated with Land's retinex algorithm (Land, 1964) is to assume that the brightest point of a scene at each point of the spectrum is reflecting 100% of the illuminant at that spectral point. The spatially averaged cone signals correlate with the maximum signal in each class, and also with the illuminant cone responses, and either cue could be used to determine the Ives transform parameters.

The theories discussed so far assume only Lambertian (isotropically reflecting) surfaces. The chromatic convergence theory relies on the fact that real surface reflections have specular highlights or 'interface reflections', which have the same spectral power distribution as the illuminant. These are combined with 'body reflections', which are the product of the illuminant and material spectra.

The combination of these is determined by the geometry of the scene, but all chromaticities from the surface will fall on a line in colour space between the material and illuminant chromaticities. Therefore, these lines corresponding to several surfaces in a scene will intersect at the coordinate of the illuminant (D'Zmura & Lennie, 1986; H. C. Lee, 1986). Specular highlights have been demonstrated to influence colour constancy (Yang & Maloney, 2001), and constancy is improved when stimuli create accurate binocular disparity, however these two factors don't seem to interact (Yang & Shevell, 2002).

It has been suggested that surface reflectance functions, at least in terms of combinations of basis functions, could be estimated by having multiple views of the same scene under different illuminants. An equivalent to this would be to use shadow boundaries, which provide a view of the same surface under two effective illuminants (D'Zmura, 1992; D'Zmura & Iverson, 1993a, 1993b). An alternative would be to use the distribution of macular pigment density across the retina, which would allow views of the same scene filtered differently at different retinal locations, equivalent to differing illuminants (Broackes, 1992).

Finally, it has been suggested that the correlations between colours in a scene and prior knowledge of the probability of seeing different colours under different illuminants, could be used to estimate the illuminant (Finlayson, Hordley, & Hubel, 2001). Or, at least, it would allow several solutions to be obtained from which the best one is chosen. This would require a calibration process to obtain this prior knowledge, and there is evidence that suggests that experience during infancy with broadband lights and surfaces is necessary for colour constancy to develop in monkeys (Brenner & Cornelissen, 2005; Sugita, 2004).

A number of higher-level factors might constrain the amount to which constancy can be achieved, and different models might be affected in different ways. These factors include the relative weighting given to the various cues that might be available in a complex scene, and segmenting the retinal image of a scene to identify surfaces, shadows, specular reflections, and different illuminant sources.

1.5.4 Model implementation

The type of mechanism by which colour constancy is achieved may imply the neural locus at which the transformations are carried out. For example, the simple cone coordinate multiplicative transformations mean that it is possible for signals early in the visual system at the level of the photoreceptors to be adjusted by adaptation to achieve constancy. Similarly this could happen in the second-stage opponent mechanisms. Adaptation is a widely accepted phenomenon that occurs in these mechanisms (see section 1.3) but most evidence suggests that adaptation here affects local

retinal regions only (He & MacLeod, 1998; MacLeod & He, 1993). However, eye movements do serve to convert spatially distributed samples into a temporal stream (D'Zmura & Lennie, 1986; Fairchild & Lennie, 1992) and so could facilitate adaptation to general scene statistics provided the time constants of the mechanisms are great enough. It should be made clear, however, that the types of signal transformation discussed here do not explicitly require that the transformations take place in the peripheral mechanisms. Alternatively, it has been suggested that the cone-rich edge of the retina, the ora-serrata, may sample scattered light not directly used for perception, and may have a role in extracting spatially integrated information from the scene (Mollon, Regan, & Bowmaker, 1998).

Mechanisms of constancy requiring more cognitive function such as those using prior knowledge about the diagnostic colours of objects, or mechanisms that require information to be extracted from and compared across well-separated elements of a scene, would presumably require higher-level cortical mechanisms. There is a suggestion that simple gain adjustments in the early visual system are not an appropriate solution since it would make illumination changes imperceptible (Katz, 1935, p. 265). Instantaneous changes in illuminant chromaticity are noticeable and observers have been shown to have access to the surface colour and the illuminant colour (D'Zmura & Iverson, 1993a, 1993b). Further, it has been shown that the set of materials to be judged can affect colour appearance, not just the temporal and spatial properties of the stimuli (Smithson & Zaidi, 2004). Although the evidence for a 'colour area' in the brain is not conclusive (see section 1.1.1.6), let alone a 'colour constancy area', there is some evidence for cortical mechanisms necessary for constancy. These include studies of split-brain patients (Land, Hubel, Livingstone, Perry, & Burns, 1983; Ruttiger, et al., 1999), recordings of macaque area V4, and humans with damage to area V4 (Zeki, Aglioti, McKeefry, & Berlucchi, 1999). Some studies of patients with damage to area V1 have shown that retinal mechanisms can't support constancy (Barbur, de Cunha, Williams, & Plant, 2004), yet other studies suggest that achromatopic patients can discriminate cone excitation ratio changes and do display constancy (Hurlbert, Bramwell, Heywood, & Cowey, 1998).

1.5.5 The time-course of colour constancy

Most of the studies on colour constancy use simulated environments that are static, or do not consider the amount of time needed to achieve constancy. In section 7 of this thesis, we present an experiment that measures colour constancy over time, in a simulated environment where the illuminant changes. There are several reasons why this may be interesting. In the natural world we move around the environment and this affects our viewpoint towards surfaces, as well as allowing objects and illuminant sources to become occluded. Natural illuminants may change their chromaticity and geometry over time, the most obvious case of this is the transition of the sun during

the day, or the movement of clouds. Eye movements over a scene may allow distributed spatial information to be converted to a temporal stream for the purposes of estimating the average chromaticity, so it may be beneficial to allow the colour constancy corrections to take place over time, rather than making the correction instantaneously. Also, information on the time course of constancy may allow us to infer the type of mechanisms involved. For example, we know that the effects of local retinal adaptation can be measured after fractions of a second (Crawford, 1947) but those involving post-receptoral mechanisms may take longer (Pugh & Mollon, 1979). Models relying on previous experience with coloured objects during development might predict that constancy in simulated environments, such as the ones we use, is impossible because the stimuli do not contain objects that could have been seen previously with consistent reflectances.

Much of the relevant work on the time course of constancy has studied lightness constancy, not specifically colour constancy. A well-known phenomenon is the Gelb effect (See Kentridge (2005) for an example and discussion of this effect), in which a black object appears white when it is the only thing illuminated by a bright spotlight. Moving a white object close to it makes the black surface immediately appear dark, and it returns to white when the real white object is removed. This has been given as evidence against a model based on prior knowledge of object reflectances, but has been disputed (Newson, 1958). The most prevalent model of lightness constancy is an anchoring one, where the surface with highest reflectance in a scene is treated as white, and other surfaces take on brightness relative to this (e.g. Gilchrist, et al., 1999; Land & McCann, 1971). It has been demonstrated that the strength of the effect can be affected by the complexity (or articulation) of the scene (Cataliotti & Bonato, 2003).

Cataliotti & Gilchrist (1995) introduced light grey and white patches into an artificial Mondrian environment, consisting of dark grey patches (the brightest of which appeared white), changing the average lightness. They found that the new patches initially appeared self-luminous and that it took two minutes for the original dark grey patches to appear darker. Annan & Gilchrist (2004) conclude that this is due to anchor persistence – the assumption of surfaces that are luminance-constant across the scene – not to lightness persistence or retinal adaptation. They varied the size of this effect by varying the number of patches that did not change, and they conclude that the assumption by the visual system that these patches were constant drives the perception that the illuminant has not changed, so re-anchoring is not immediate.

Cataliotti & Bonato (2003) make an analogy between spatial anchoring and temporal anchoring in lightness constancy. They show that the brightness of even a spot-in-a-void can affect the brightness

of another, after several seconds, and they demonstrate that this is not due to adaptation in peripheral mechanisms.

It has been shown by Hansen, Walter & Gegenfurtner (2007) that removing temporal contrast information reduces colour constancy (as measured by a constancy index similar to those described in section 7.3.2). In this case they define temporal contrast as replacing the test stimuli to be categorized with the illuminant chromaticity (which also covered the surround), and show that constancy is better in this condition than when the test stimuli is replaced by black. Removing spatial contrast, by limiting the eccentricity of the surround illuminant field within a larger black field, has a similar effect.

Schultz, Doerschner, and Maloney (2006) found that constancy was greater for a series of isolated simulated reflectances when these reflectances are presented in blocks in which the illuminant was constant, compared to blocks in which the illuminant and reflectance were randomised on each presentation. They conclude that judgements must have been dependent on information integrated temporally over preceding samples, and that this process could be explained by a process of slow chromatic adaptation. A similar process has also been proposed to explain the results of Smithson and Zaidi (2004) who found strong influences of the chromatic statistics of preceding stimuli, even when these were pitted against spatial cues.

1.5.6 Justification and Aims

As the previous section has described, few studies have investigated the time course of colour constancy. Most of the studies in which temporal context is considered are of lightness constancy only, not colour constancy. They do, however, suggest that several seconds or more are required to achieve lightness constancy so we ask whether colour constancy is also not instantaneous. In the experiment in section 7, we aim to continue the theme of the temporal nature of processing in the colour vision system and apply it to colour constancy. We use a stimulus environment that changes regularly and includes shifts of illuminant chromaticity, and we make instantaneous measurements of colour appearance so that we can investigate the degree of colour constancy as a function of time.

2 Hue circle stimuli

The experiments in sections 4, 5 and 6 use a stimulus which, for brevity, we often refer to as a Hue Circle, although in many cases this is somewhat inaccurate. As described in section 4.2.3, the stimulus is a periodic sequence of colours in spectral sequence, or in order of continuously changing hue around an elliptical locus in colour space. The history and versions of this stimulus are discussed in section 1.4.1.5. At a frequency of 10Hz (i.e. 10 cycles of the periodic stimulus per second), the stimulus appears to flicker between two predominant colours, and these colours differ when the temporal order is reversed. For the purposes of these and previous experiments (Stromeyer, et al., 1991; Zaidi & Halevy, 1993), it is appropriate to consider the sequences of colours to be following a locus in colour space. As discussed in sections 1.2.2 and 1.2.3, the colour spaces that are most convenient in this case have their axes defined by the mechanisms thought to convey chromatic signals, from either psychophysical or physiological experimental findings. These loci will also be elliptical in other colour spaces that can be described by linear transformations of the spaces we discuss (e.g. the CIE 1931 diagram). Specifying stimuli in a physiologically based colour space allows easy comparison between the properties of the stimulus loci and signals in the neural mechanisms we are attempting to manipulate.

2.1 Mathematical construction

At the very simplest level, the stimuli discussed here are repeating sequences of hues, with the hues in either forward or reverse order. As such, they could be defined in various ways including sequences of cone excitation ratios or contrast ratios, primary light outputs of a given apparatus that would produce those cone excitation ratios, or spectral power distributions that are metameric to those cone excitation ratios. Indeed, the stimuli are converted to primary light outputs in order to present them experimentally, and since any digitally controlled apparatus can only present discretely sampled values of the stimulus, it is also possible to store the stimulus definitions as a discrete list of numeric values.

However, it is much more convenient when discussing theory to represent the stimuli as mathematical functions and manipulate them analytically. There are situations, however, when mathematical representations do not simplify nicely, and are therefore unwieldy and difficult to solve. In these cases numerical methods are used. Both approaches will be discussed.

2.1.1 Generating circular loci

An ellipse in a plane with axes x and y can be described by the parametric equations:

$$x = \sin(t) \quad (1)$$

$$y = \sin(t - \theta) \quad (2)$$

A circular locus results when the phase difference θ between these two oscillations is $\pm\pi/2$ (or equivalently $\pm3\pi/2$). Reversing the sign of the phase difference reverses the direction in which the point specified at any given value of the parameter moves around the locus as the parameter t increases. Since we want to generate temporal modulations of chromaticity around a locus, the parameter t we use corresponds to time. t has the value of zero at the beginning of stimulus presentation and increases linearly by 2π during one cycle around the locus. We use a stimulus that cycles at 10Hz for most of the experiments, so one unit of t corresponds to $1/(20\pi)$ seconds. When the component modulations expressed in equations 1 and 2 are along the two axes of the equiluminant plane of MB-DKL colour space, the locus traces a circular sequence of colours such as those in the hue circle stimuli described above and we can relate the appearance of the stimuli to the signals in the underlying neural mechanisms. The coordinates specified by the values of the component modulation functions at a given value of t represent the chromaticity of the stimulus at the corresponding time. In a similar way to Derrington, Krauskopf and Lennie (1984), we set the origin of this space to be white, so the distance (radius) of any point from the origin corresponds to saturation. For the rest of this section, to simplify mathematical representation, $\Delta[L-M]$ is abbreviated to l and $\Delta[S-(L+M)]$ is abbreviated to s .

Phase differences other than $\pm\pi/2$ will result in loci that are not circular. Component modulations that are in phase with one another will result in a loci collapsed to a straight line along the positive diagonal of the plane, and over time the chromaticity will move sinusoidally back and forth along this line. If the components are in antiphase, that is they have a phase difference of π , the locus will collapse to a straight line along the *negative* diagonal. Phase differences between $-\pi/2$ (or $3\pi/2$) and $\pi/2$ will result in an elliptical locus oriented with its major axis along the positive diagonal and phase differences between $\pi/2$ and $3\pi/2$ will result in an elliptical locus oriented with its major axis along the negative diagonal (where a positive phase difference indicates the vertical or $S-(L+M)$ modulation lags the $L-M$ modulation). Examples of relationships between phase and ellipse orientation and shape are shown in Figure 5.

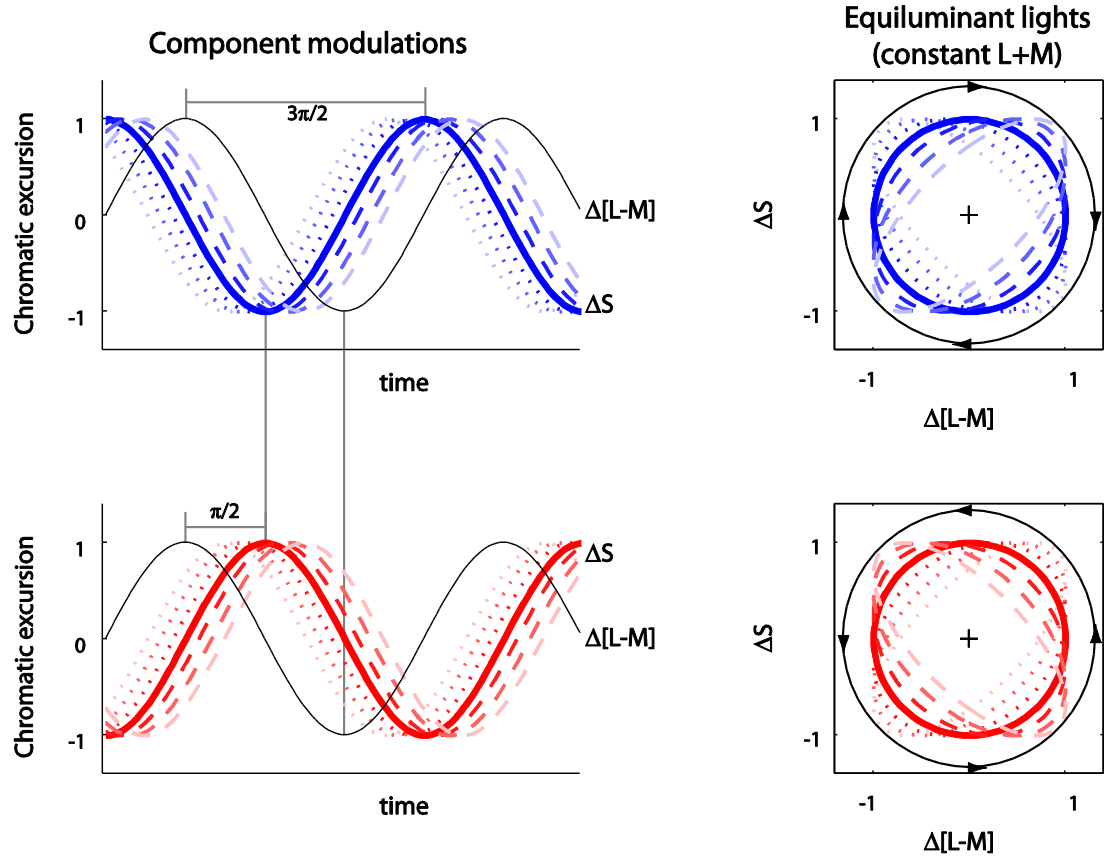


Figure 5. Illustration of how phase differences between l and s modulations affect the shape and orientation of the resultant ellipse.

2.1.2 Cardinal axes

The stimuli we wish to create in sections 4, 5 and 6 are constructed from sinusoidal modulations along the cardinal axes of MB-DKL space. Additional stimuli for experiment 1 are generated from components in directions intermediate to the cardinal axes, and they will be discussed below. In both cases, the component modulations have a relative phase difference of θ . In experiment 5, θ is fixed at $\pi/2$. In both experiments we transform a clockwise (CW) stimulus into a counter-clockwise stimulus by adding a phase difference of π to one component.

We define one component as a sinusoidal modulation of the signal in the L-M mechanism (here represented by l) as a function of t :

$$l = \sin(t) \quad (3)$$

and the second component as a sinusoidal modulation of the signal in the S-(L+M) mechanism (here represented by s) as a function of t with a phase difference of θ for a CW stimulus :

$$s_{CW} = \sin(t - \theta) \quad (4)$$

or $\theta + \pi$ for a CCW stimulus:

$$s_{CCW} = \sin(t - \theta - \pi) \quad (5)$$

When we simulate another phase difference φ due to a delay inherent to the S-(L+M) mechanism then the simulated modulations arriving at the site at which signals in the two opponent mechanisms are combined become

$$l' = \sin(t) \quad (6)$$

$$s'_{CW} = \sin(t - \theta + \varphi) \quad (7)$$

$$s'_{CCW} = \sin(t - \theta + \varphi - \pi) \quad (8)$$

Example loci with a range of θ and φ values are plotted in Figure 6.

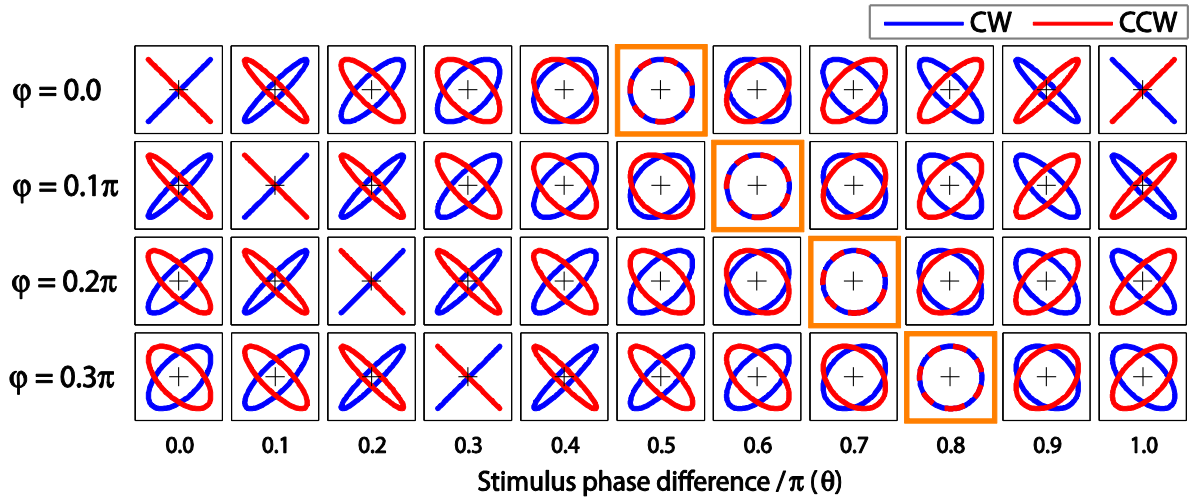


Figure 6. Simulated loci reaching a processing stage at which modulations in the two cardinal axes of equiluminant space are combined. Each square plot represents a section of the equiluminant plane, as described in section 1.2.2. Blue lines represent the loci plotted by l and s_{CW} , and as t increases, chromaticity moves in a CW direction around the ellipse. Red lines represent the loci plotted by l and s_{CCW} , where chromaticity moves in a CCW direction. Each row represents a different simulated neural delay φ , as indicated on the left. The top row, $\varphi=0$, represents no simulated neural delay, and also shows the loci of the physical stimulus presented. The cases when CW and CCW ellipses are identical are highlighted with an orange surround.

Using Pythagoras' theorem, the radius r of the ellipse, or distance from the origin of the space (in this case, the chromaticity of EEW), is given by:

$$r_{CW} = (l'^2 + s'_{CW}{}^2)^{\frac{1}{2}} \quad (9)$$

$$r_{CCW} = (l'^2 + s'_{CCW}{}^2)^{\frac{1}{2}} \quad (10)$$

Therefore:

$$r_{CW} = (\sin^2 t + \sin^2(t - \theta + \varphi))^{\frac{1}{2}} \quad (11)$$

$$r_{CCW} = (\sin^2 t + \sin^2(t - \theta + \varphi - \pi))^{\frac{1}{2}} \quad (12)$$

2.1.3 Intermediate axes

In the intermediate axes condition, we initially specify our stimuli as sinusoidal modulations of our component axes which are rotated 45° or $\pi/4$ relative to the cardinal axes. Referring to these modulations as a and b :

$$a = \sin(t) \quad (13)$$

for a CW stimulus :

$$b_{CW} = \sin(t - \theta) \quad (14)$$

or $\theta+\pi$ for a CCW stimulus:

$$b_{CCW} = \sin(t - \theta - \pi) \quad (15)$$

However, in order to perform any modelling and make predictions, we will need to represent these modulations as functions of coordinates in the cardinal axes of colour space, or in terms of l and s . This is particularly important since the phase shift caused by a neural delay in the S-(L+M) mechanism must only be applied to signals in that mechanism, or to s . The relationship between the intermediate axes modulations and their l and s coordinates, before considering neural phase differences, is obtained with a simple rotation matrix and multiplication:

$$\begin{pmatrix} l \\ s \end{pmatrix} = \begin{pmatrix} \cos(\alpha) & -\sin(\alpha) \\ \sin(\alpha) & \cos(\alpha) \end{pmatrix} \begin{pmatrix} a \\ b \end{pmatrix} \quad (16)$$

Giving:

$$l = a\cos(\alpha) - b\sin(\alpha) \quad (17)$$

$$s = a\sin(\alpha) + b\cos(\alpha) \quad (18)$$

Where α is the angle between the intermediate axes and the cardinal axes and positive values indicate the intermediate axes are rotated counter-clockwise from the cardinal ones. In the cardinal axes condition, $\alpha=0$ so $\sin(\alpha) = 0$ and $\cos(\alpha) = 1$. In this case, the equations 17 and 18 simplify to

the equations for the cardinal axes conditions (equations 3, 4 and 5). In the intermediate axes condition of the experiment in section 4 in this thesis, $\alpha = \frac{\pi}{4}$, so $\cos(\alpha) = \sin(\alpha) = 2^{-\frac{1}{2}}$.

Expanding equations 17 and 18 gives:

$$l_{CW} = \sin(t) \cos(\alpha) - \sin(t - \theta) \sin(\alpha) \quad (19)$$

$$s_{CW} = \sin(t) \sin(\alpha) + \sin(t - \theta) \cos(\alpha) \quad (20)$$

$$l_{CCW} = \sin(t) \cos(\alpha) - \sin(t - \theta - \pi) \sin(\alpha) \quad (21)$$

$$s_{CCW} = \sin(t) \sin(\alpha) + \sin(t - \theta - \pi) \cos(\alpha) \quad (22)$$

and inserting the neural phase delay into s gives:

$$l'_{CW} = \sin(t) \cos(\alpha) - \sin(t - \theta) \sin(\alpha) \quad (23)$$

$$s'_{CW} = \sin(t + \varphi) \sin(\alpha) + \sin(t - \theta + \varphi) \cos(\alpha) \quad (24)$$

$$l'_{CCW} = \sin(t) \cos(\alpha) - \sin(t - \theta - \pi) \sin(\alpha) \quad (25)$$

$$s'_{CCW} = \sin(t + \varphi) \sin(\alpha) + \sin(t - \theta + \varphi - \pi) \cos(\alpha) \quad (26)$$

Substituting into the equations giving radius r :

$$r_{CW} = \left(\frac{(\sin(t) \cos(\alpha) - \sin(t - \theta) \sin(\alpha))^2 + (\sin(t + \varphi) \sin(\alpha) + \sin(t - \theta + \varphi) \cos(\alpha))^2}{(\alpha)^2} \right)^{\frac{1}{2}} \quad (27)$$

$$r_{CCW} = \left(\frac{(\sin(t) \cos(\alpha) - \sin(t - \theta - \pi) \sin(\alpha))^2 + (\sin(t + \varphi) \sin(\alpha) + \sin(t - \theta + \varphi - \pi) \cos(\alpha))^2}{(\alpha)^2} \right)^{\frac{1}{2}} \quad (28)$$

Example loci are shown in Figure 7, with the same example values of θ and φ as above in Figure 6. Compare the two figures, and note the different patterns in ellipse shape and orientation.

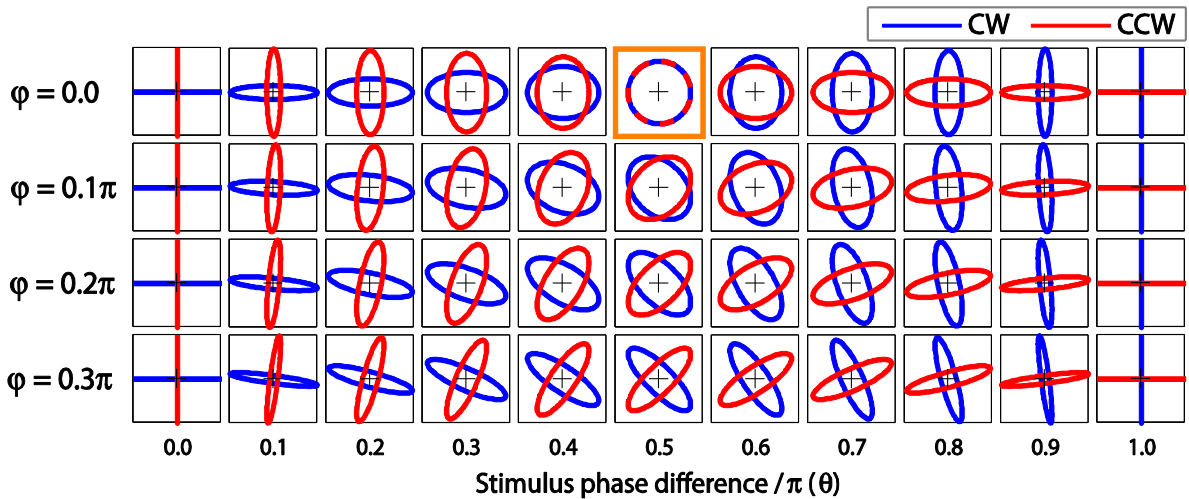


Figure 7. Simulated loci reaching a processing stage at which modulations in two directions, oriented 45° ($\alpha=\pi/4$) from the two cardinal axes of equiluminant space, are combined having been converted to representations in the cardinal axes with a phase delay of φ added to the s modulation. Figure labelling is the same as in Figure 6.

There are important contrasts between the two patterns, and these have theoretical implications for the experiment in section 4. The differences are summarised below.

	Cardinal	Intermediate
Symmetry	$\theta=\pi/2+\varphi$	$\theta=\pi/2$
Collapsed ellipses	Oriented at $\pm 45^\circ$	Oriented at 0° and 90°
Ellipse orientation	Oriented at $\pm 45^\circ$	Orientation changes with θ
Identical (circular)CW and CCW loci	$\theta=\pi/2+\varphi$	$\theta=\pi/2$, only when $\varphi=0$

2.2 Discrimination model

In the experiment in section 4 (described in much more detail later), observers are tasked to discriminate between two stimuli such as the ones described here. One has a physical phase difference of θ , the other has a phase difference of $\theta-\pi$. This corresponds to CW or CCW stimuli as described above. We measure discrimination threshold as a function of θ , and we wish to model the results and to generate a template to fit to the data.

Our measure of threshold is obtained by reducing the amplitude of one of the component modulations until the CW and CCW are indistinguishable. As the stimulus amplitude ratio decreases, the loci collapse to straight lines in the direction of the modulation that is not altered by the staircase. This results in the major axes of the ellipses being rotated towards this line. This is illustrated in Figure 8.

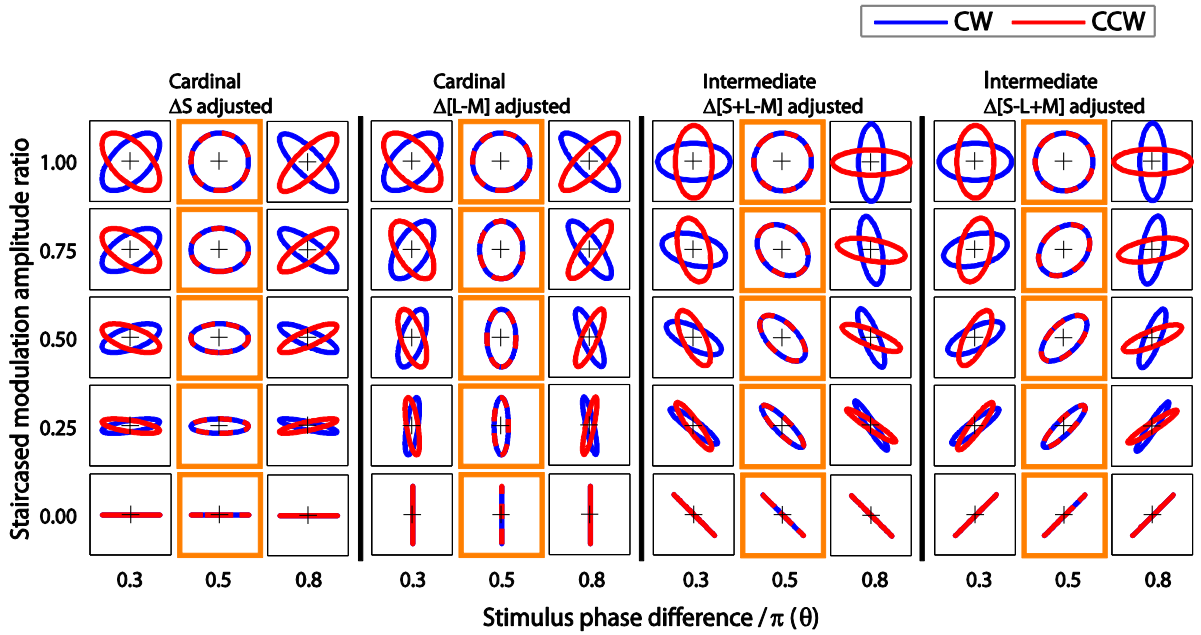


Figure 8. Simulated loci at various examples of stimulus phase difference, when the stimuli are constructed from modulations in the cardinal axes (left panels) and intermediate axes (right panels), showing how the loci are changed by reducing the amplitude of one of the components. Symbology is the same as the preceding figures. The top row of boxes shows loci at the beginning of each session of the experiment in section 4 – modulation amplitudes are equivalent. Subsequent rows show loci as one component amplitude is reduced, as in the section 4.

This experiment repeats and extends that of Stromeyer et al. (1991). In their experiment, they use stimuli similar to those in our cardinal axes condition. They fit their data with a curve of the form $|\cos(\theta-\varphi)|^{-1}$ (symbols have been replaced with those corresponding to the equivalent notation in this thesis), and this curve captures the basic features of their data. The form of the curve is shown in Figure 9.




Figure 9. The curve $T = |\cos(\theta-\varphi)|^{-1}$ used by Stromeyer et al, at several values of φ .

This curve tends to infinity at $\theta = \varphi + \pi/2$, and decreases symmetrically either side of this point to a minimum at $\theta = \varphi$ and $\theta = \varphi + \pi$.

For the purposes of this thesis we wish to undertake detailed modelling with manipulation of more variables, and so we derive a function predicting threshold from θ , from scratch. We also wish to include more variables, including φ and α . The function must display the following features, based on the qualitative predictions:

- The predicted threshold should be maximal when the resultant CW and CCW loci have the most similar shape
- If the CW and CCW loci have the same shape, the stimuli should be predicted to be indiscriminable
- The predicted threshold should be minimal when the resultant CW and CCW loci have the most different shape

The simplest metric that fits these criteria is obtained in two stages: First we represent the difference between the loci by calculating the Weber contrast between the maximum of one locus (say, the CW locus) and the minimum of the other (the CCW locus) at any given values of θ , φ , etc, and subtract one so that the value of this difference is zero when the loci completely overlap. We then take the inverse of this to predict threshold. This is represented in equation 29, showing the general form of the model predicting discrimination threshold T :

$$T = \left(\frac{\max(r_{CW})}{\min(r_{CCW})} - 1 \right)^{-1} \quad (29)$$

Similar results are obtained by taking the difference between $\max(r_{CW})$ and $\min(r_{CCW})$ instead of the ratio, although this can lead to complications when the relative scaling of the signals is considered.

To find $\max(r_{CW})$ and $\min(r_{CCW})$ we take the derivatives of equations 27 and 28, in the case of the cardinal axes, which are:

$$\frac{dr_{CW}}{dt} = (\sin^2 t + \sin^2(t - \theta + \varphi))^{-\frac{1}{2}} (\sin t \cos t + \sin(t - \theta + \varphi) \cos(t - \theta + \varphi)) \quad (30)$$

$$\frac{dr_{CCW}}{dt} = (\sin^2 t + \sin^2(t - \theta + \varphi - \pi))^{-\frac{1}{2}} (\sin t \cos t + \sin(t - \theta + \varphi - \pi) \cos(t - \theta + \varphi - \pi)) \quad (31)$$

At the maxima and minima of r_{CW} and r_{CCW} :

$$\frac{dr}{dt} = 0$$

therefore, for CW stimuli:

$$(\sin^2(t) + \sin^2(t - \theta + \varphi))^{-\frac{1}{2}} = 0$$

OR

$$(\sin(t) \cos(t) + \sin(t - \theta + \varphi) \cos(t - \theta + \varphi)) = 0 \quad (32)$$

and for CCW stimuli:

$$(\sin^2(t) + \sin^2(t - \theta + \varphi - \pi))^{-\frac{1}{2}} = 0$$

OR

$$(\sin t \cos t + \sin(t - \theta + \varphi - \pi) \cos(t - \theta + \varphi - \pi)) = 0 \quad (33)$$

When $0 < (\theta - \varphi) < \frac{\pi}{2}$ and $\frac{3\pi}{2} < (\theta - \varphi) < \pi$:

r_{CW} is maximum when

$$t = \frac{\theta - \varphi}{2} + \left(k + \frac{1}{2}\right)\pi$$

r_{CW} is minimum when

$$t = \frac{\theta - \varphi}{2} + k\pi$$

r_{CCW} is maximum when

$$t = \frac{\theta - \varphi + \pi}{2} + k\pi$$

r_{CCW} is minimum when

$$t = \frac{\theta - \varphi + \pi}{2} + \left(k + \frac{1}{2}\right)\pi$$

When $\frac{\pi}{2} < (\theta - \varphi) < \frac{3\pi}{2}$:

r_{CW} is maximum when

$$t = \frac{\theta - \varphi}{2} + k\pi$$

r_{CW} is minimum when

$$t = \frac{\theta - \varphi + \pi}{2} + \left(k + \frac{1}{2}\right)\pi$$

r_{CCW} is maximum when

$$t = \frac{\theta - \varphi}{2} + \left(k + \frac{1}{2}\right)\pi$$

r_{CCW} is minimum when

$$t = \frac{\theta - \varphi + \pi}{2} + k\pi$$

where $k \in \mathbb{Z}$.

Substituting this into above gives:

$$T = \left(\frac{\left(\sin^2\left(\frac{\theta - \varphi}{2} + \left(k + \frac{1}{2}\right)\pi\right) + \sin^2\left(\frac{\theta - \varphi}{2} + \left(k + \frac{1}{2}\right)\pi - \theta + \varphi\right)\right)^{\frac{1}{2}}}{\left(\sin^2\left(\frac{\theta - \varphi + \pi}{2} + \left(k + \frac{1}{2}\right)\pi\right) + \sin^2\left(\frac{\theta - \varphi + \pi}{2} + \left(k + \frac{1}{2}\right)\pi - \theta + \varphi - \pi\right)\right)^{\frac{1}{2}} - 1} \right)^{-1} \quad (34)$$

$$\text{for } 0 < (\theta - \varphi) < \frac{\pi}{2} \text{ and } \frac{3\pi}{2} < (\theta - \varphi) < \pi, k \in \mathbb{Z}$$

$$T = \left(\frac{\left(\sin^2\left(\frac{\theta - \varphi}{2} + k\pi\right) + \sin^2\left(\frac{\theta - \varphi}{2} + k\pi - \theta + \varphi\right)\right)^{\frac{1}{2}}}{\left(\sin^2\left(\frac{\theta - \varphi + \pi}{2} + k\pi\right) + \sin^2\left(\frac{\theta - \varphi + \pi}{2} + k\pi - \theta + \varphi - \pi\right)\right)^{\frac{1}{2}} - 1} \right)^{-1} \quad (35)$$

$$\text{for } \frac{\pi}{2} < (\theta - \varphi) < \frac{3\pi}{2}, k \in \mathbb{Z}$$

which simplify to:

$$T = \left(\frac{\left(\sin^2\left(\frac{\theta - \varphi + (2k + 1)\pi}{2}\right) + \sin^2\left(\frac{\varphi - \theta + (2k + 1)\pi}{2}\right)\right)^{\frac{1}{2}}}{\left(\sin^2\left(\frac{\theta - \varphi + (2k + 2)\pi}{2}\right) + \sin^2\left(\frac{\varphi - \theta + 2k\pi}{2}\right)\right)^{\frac{1}{2}} - 1} \right)^{-1} \quad (36)$$

$$\text{for } 0 < (\theta - \varphi) < \frac{\pi}{2} \text{ and } \frac{3\pi}{2} < (\theta - \varphi) < \pi$$

$$T = \left(\frac{\left(\sin^2 \left(\frac{\theta - \varphi + 2k\pi}{2} \right) + \sin^2 \left(\frac{\varphi - \theta + 2k\pi}{2} \right) \right)^{\frac{1}{2}}}{\left(\sin^2 \left(\frac{\theta - \varphi + 3k\pi}{2} \right) + \sin^2 \left(\frac{\varphi - \theta + (2k - 1)\pi}{2} \right) \right)^{\frac{1}{2}} - 1} \right)^{-1} \quad (37)$$

for $\frac{\pi}{2} < (\theta - \varphi) < \frac{3\pi}{2}$

Example curves, showing T plotted against θ , for three values of φ , are shown in Figure 10.




Figure 10. Examples of the model predicting thresholds as a function of θ in the cardinal axes condition, as defined by equations 36 and 37. The curves are only shown for the range $\theta < \theta < \pi$. The curve is periodic over π and the curve in the range $\pi < \theta < 2\pi$ is a repeat of that shown.

In the intermediate axes condition, as mentioned above, the derivatives of these equations do not simplify elegantly. They are unwieldy and difficult to solve analytically. Therefore, templates, plots and fits using these expressions were generated by finding $\max(r_{CW})$ and $\min(r_{CCW})$ numerically.

Example curves for the intermediate axes condition, again showing T plotted against θ , for the same three values of φ as above and with $\alpha = \pi/4$, are shown in Figure 11.




Figure 11. Examples of the model predicting thresholds as a function of θ in the intermediate axes condition, generated numerically. As in Figure 10, the curve repeats in the range $\pi < \theta < 2\pi$.

There are important similarities and differences between the curves for the cardinal and intermediate axes conditions, which mirror those between the sets of loci in Figure 6 and Figure 7. The important differences are summarised below:

	Cardinal	Intermediate
Symmetry	$\theta=\pi/2+\varphi$	$\theta=\pi/2$
Peak threshold	$\theta=\pi/2+\varphi$	$\theta=\pi/2$
Threshold tends to infinity	<i>Yes</i>	Only when $\varphi=0$

Perhaps the most important characteristic of the cardinal axes curves is how predicted threshold tends to infinity and increasing the simulated delay φ in the S-opponent mechanism shifts the position of this peak to the right (higher θ). In the intermediate axes condition, the peak threshold only tends to infinity when $\varphi=0$ (and in this case the loci and threshold curves in the cardinal and intermediate conditions are identical), and increasing φ reduces the magnitude of the peak, which remains at $\theta=\pi/2+\varphi$.

2.3 More complex models

The model presented above and summarised in Figure 12 assumes the classical two-opponent-mechanism theory of colour vision processing, and may be too simplistic. The model makes several implicit assumptions about the visual system. These are:

- That chromatic information is carried in two, and only two, independent mechanisms that linearly combine cone signals: $\Delta[L-M]$ and $\Delta[S-(L+M)]$.
- These mechanisms are bipolar, responding to both increases and decreases in signal (relative to the current adaptation state), and with opposite sign.
- Mechanisms tuned to chromatic directions other than the cardinal axes are not simulated
- The simulated phase delay φ in the S-cone mechanism is not specific to a particular stage (e.g. cones, ganglion cells, LGN) and is simply assumed to apply to all signals in mechanisms with input from S-cones. This is simplified by the fact that we constrain stimuli to lie in the equiluminant plane, where L+M is held constant, so our $\Delta[S-(L+M)]$ modulation can be thought of (and is generated experimentally) as simply ΔS .

These assumptions have been challenged by recent physiological data and some psychophysical results described in section 1.1.1.4.




Figure 12. The stages in the simple model of colour processing in which we assume signals from cones are carried in two opponent chromatic mechanisms. Phase differences θ applied to the S-cone isolating stimulus modulation are added linearly to the single neural phase difference φ . φ could be modelled in the S-cones or in the S-opponent pathway.

Since we are attempting to investigate possible influences of mechanisms that combine cone signals in other ways, we also constructed a model to include some of these features. This model assumes modulations of the cone signals that are the same as in the model above. There are a number of possible configurations of this model, and alternatives will be discussed below, however we believe the following outline of the stages (summarised in Figure 13), is the most appropriate:

1. Lights are modulated sinusoidally to generate sinusoidal modulations of cones that would generate the elliptical loci in MB-DKL space. L- and M-cones are modulated in antiphase, and L+M is held constant. S-cones are modulated with a phase difference of θ .
2. The signals in cones are in response to the stimulus in stage 1. An additional phase difference φ_1 is added to the S-cone signal to simulate a delay introduced at a pre-opponent stage such as S-cones or S-cone-bipolar cells.
3. Five more signals are constructed, each differently combining mechanisms from stage 2. Each mechanism is rectified with a transducer function that which smoothly flattens negative responses while leaving positive responses linear (see Smithson, Henning, MacLeod and Stockman (2009) for recent review). The function we use takes the form:

$$f(x) = 0.1 \left(\log_{10} \left(1 + e^{\frac{x}{0.1}} \right) - \log_{10}(2) \right) \quad (38)$$

These units are designed to simulate opponent mechanisms such as ganglion cells or LGN cells (without modelling a specific stage), and respond to the following combinations of cone signals: +L-M, -L+M, +S-L-M, -S+L+M and -S+L-M. The first four of these units simulate ON and OFF mechanisms aligned to the cardinal axes, while the last simulates off-axis cells, in particular those demonstrated by Tailby et al. Independent phase delays can be added to each of the mechanisms with S-cone input: φ_2 , φ_3 and φ_4 are added to the signals in S-L-M, -S+L+M and -S+L-M respectively, to simulate delays arising in these mechanisms.

4. Signals from the five units in the previous stage are combined to simulate two orthogonal mechanisms: $(+L-M)-(-L+M)$ and $(+S-L-M)-(-S+L+M)-(-S+L-M)$. A 5th phase shift, φ_5 , is added to the signal in the second of these simulated mechanisms. This is to simulate delays originating after stages that directly compare cone outputs. This could be at the LGN or cortical level.
5. Thresholds are predicted from the ratio of maximum and minimum of the CW and CCW stimulus loci, respectively, as in equation 29.

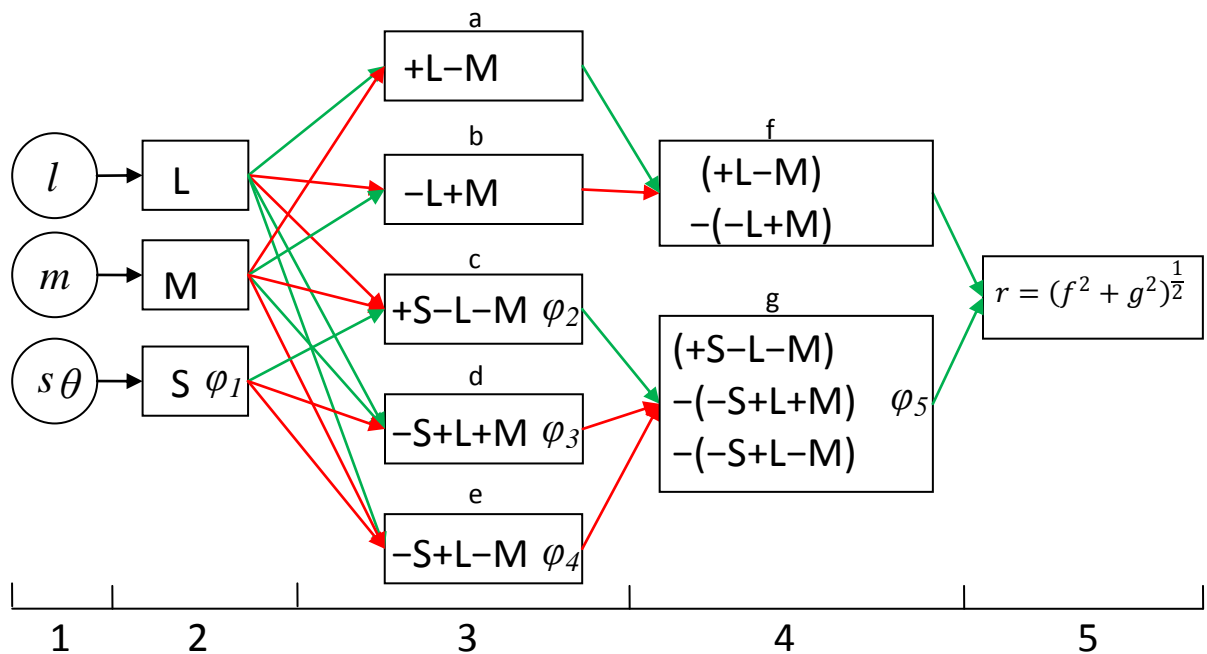


Figure 13. Block diagram representing the model described above incorporating rectified opponent mechanisms, and mechanisms responding to chromatic directions not aligned to the cardinal axes. Green arrows indicate input with a positive sign, and red arrows indicate input with a negative sign.

As with the intermediate axes case of the simple model, threshold predictions were calculated numerically using a computer program, rather than analytically.

The magnitudes of the responses in each of the units at each stage could be adjusted independently, to simulate differently weighted inputs of one stage to another, or to 'turn off' individual units to see the difference between a particular unit being present or not.

An important feature to note is that phase delays in successive aligned mechanisms add linearly. For example, the phase difference between the outputs of the $+S-L-M$ unit and the s signal will be $\varphi_1 + \varphi_2$. Each delay, or combination of delays, should be considered relative to the original modulations of the lights, not to modulations in parallel mechanisms (those at the same stage of processing). However, this becomes more complicated when parallel mechanisms have different phase delays, for example

if φ_2 , φ_3 and φ_4 are not all equal. In this case, the differently-delayed signals are added by the next unit, before φ_5 is applied to the resulting signal.

The loci are now excursions in a plane defined by the signals in units f and g in the above figure, and these signals are still treated as if they are orthogonal. When $-S+L-M$ units are active, loci are rarely elliptical, and are quite complicated shapes often including sharp corners due to the rectification stage. Further, the CW and CCW ellipses are also rarely the same shape, regardless of orientation. There are a few occasions when the loci overlap, although they are still not elliptical in these occasions. In fact, the only conditions when the loci overlap are when φ_1 is the only nonzero simulated delay. The shape of the loci means that finding a general characterisation of the difference between the two, and therefore calculating predicted discrimination threshold, is difficult. The maximum and minimum excursions of the loci are not always in orthogonal directions from the centre of modulation, and the simple metric we use takes no account of the difference in shape that may exist between the two loci.

With this in mind, we adjusted parameters in the model to generate examples of the shapes of threshold, as a function of θ for plausible values of parameters, drawing conclusions only when the shapes of loci suggested that our difference metric might be sensible. In particular, we were interested in whether the $-S+L-M$ responsive cells (Tailby et al. (2008), and see discussion in section 4.2.1) contribute to performance in our task, so compared predicted threshold curves with these units 'turned on' or 'turned off'.

If the magnitudes of all the units except the $-S+L-M$ units are set to 1.0, and φ_2 and φ_3 are equal, the model is equivalent to the simple model above, since, for example, the $(+L-M)-(-L+M)$ unit opposes the rectified outputs of the $+L-M$ and $-L+M$ units in the previous stage, resulting in a unit that responds like the $+L-M$ opponent unit in the simple model. The same occurs for mechanisms with S-cone input. We compared the simple model with the configuration including $-S+L-M$ units, and find some important differences, summarised in Table 1 below:

Table 1. Summary of important features of the model described in the text, given different magnitudes of phase differences.

	-S+L-M not present		-S+L-M present							
	$\varphi_2 = \varphi_3$		$\varphi_1 > 0$ only		$\varphi_2 = \varphi_3, \varphi_2, \varphi_3 > 0$ only		$\varphi_4 > 0$ only		$\varphi_5 > 0$ only	
	Cardinal	Intermediate	Cardinal	Intermediate	Cardinal	Intermediate	Cardinal	Intermediate	Cardinal	Intermediate
Symmetry	$\theta = \pi/2 + \varphi_1 + \varphi_2 + \varphi_3 + \varphi_5$	$\theta = \pi/2$	None in range $0 < \theta < \pi$	None	None	None	None	None	None	None
Peak threshold	$\theta = \pi/2 + \varphi_1 + \varphi_2 + \varphi_3 + \varphi_5$	$\theta = \pi/2$	$\theta = \pi/2 + \varphi_1$	$\theta \approx \pi/2$	$\theta \approx \pi/2 + \varphi_2 + \varphi_3$	Varies	Varies	Varies	Varies	Varies
Threshold tends to infinity (Identical loci)	Yes	Only when $(\varphi_1 + (\varphi_2, \varphi_3) + \varphi_5) = 0$	$\theta = \pi/2 + \varphi_1$	No	No	No	No	No	No	No

In summary, adding $-S+L-M$ units to the model destroys the symmetry in the predicted threshold curves, and unless the only phase delay present is φ_1 , reduces the height of the curves in the cardinal axes condition so that they do not tend to infinity.

2.4 Alternative models

There are a number of alternative configurations that this model could take. Firstly, to represent the way in which $-S+L-M$ cells are tuned to chromaticities along the negative diagonal in the equiluminant plane (i.e. in the bottom-right quadrant when white is defined as the origin) we could feed the horizontal and vertical components of this vector to the units in the next stage. This would make the units in the next stage $(+L-M)-(-L+M)+(-S+L-M)\cos(45^\circ)$ and $(+S-L-M)-(-S+L+M)-(-S+L-M)\cos(45^\circ)$.

Alternatively, the traditional S-OFF opponent units $(-S+L+M)$ could be completely removed, and replaced by the $-S+L-M$ units. Rationale for this could be that S-OFF mechanisms are notoriously difficult to find (see section 1.1.1.4). However, Tailby et al (2008) do show that the proportions of cells responding to $-S+L+M$ and $-S+L-M$ chromatic directions are roughly equal, so we reject this as implausible.

A further way of combining $-S+L-M$ mechanisms with other opponent ones at the same stage is to treat them as a completely separate signal, akin to a third chromatic dimension (although different from the luminance mechanism which is represented as the third dimension in DKL colour space (see section 1.2.3)). Radii are calculated as the distance from the origin as before, but now in three dimensions. Instead of

$$r = (f^2 + g^2)^{\frac{1}{2}} \quad (39)$$

we take

$$r = ((a - b)^2 + (c - d)^2 + (e)^2)^{\frac{1}{2}} \quad (40)$$

In this case, we don't explicitly define stage 4 and so do not have a separate phase difference applied to the S-cone signalling mechanism after the stage 3 units are combined.

The calculation of the discrimination threshold metric is another part of the model that could be done in different ways. Instead of taking the ratio of maximum and minimum radii, as in equation 29, we could use the difference instead:

$$T = (\max(r_{CW}) - \min(r_{CCW}))^{-1} \quad (41)$$

Another option is to sum the differences between the summed magnitudes of differences between CW and CCW radii at each radial direction from the origin, and then take the inverse:

$$T = \left(\int_{\beta=0}^{\beta=2\pi} |r_{CW(\beta)} - r_{CCW(\beta)}| d\beta \right)^{-1} \quad (42)$$

where $r_{CW(\beta)}$ and $r_{CCW(\beta)}$ are the radii of the loci at a given value of β , the angular radial direction from the origin. This fulfils the first two of the requirements of the model (see beginning of section 2.2) but in many circumstances fails on the third: the threshold curve peaks and tends to infinity when resultant loci are circular, indicating the CW and CCW stimuli are indiscriminable, but peaks again when loci collapse to straight lines. At this point, the loci are in fact most different, so the threshold should be minimal.

Measurements of actual thresholds in experiment show that in the cardinal axes condition, peak thresholds are displaced from $\theta=\pi/2$ and we find the highest thresholds impossible to measure. We believe this is akin to the predicted thresholds tending to infinity in our model. In the intermediate axes condition, the maximum thresholds occur at $\theta\approx\pi/2$ but are easily measureable. The curves we see are most similar to the ones generated by a model in which $-S+L-M$ cells are not present, whatever additional assumptions we choose to make in implementing the model. The conclusions we draw from this are discussed thoroughly in section 4.6.

2.5 Relationship to measured data

The only predictions that we can make from these models, when considering them in order to predict the results of the experiment in section 4 are the shape of the plots of threshold against phase difference θ . We cannot draw any conclusions from, for example, the heights of the curves or differences between these heights in predicted and measured data. The predictions are derived from ratios of the maximum and minimum excursions of the predicted loci, whereas the measurements we make in the experiment are of the ratio of amplitudes of the component modulations when the CW and CCW stimuli are indistinguishable. We believe that these two metrics should be correlated, but will not be numerically equivalent. Experimental limitations might also constrain our measured ratios at threshold: we cannot, within the conditions of the experiment, produce a physical stimulus with a ratio of modulation amplitudes approaching infinity because one S modulation amplitude is fixed and

there is an upper limit on the amplitude of the other. We can, however, compare the relative predicted thresholds at different phase differences with the same relative differences in measured thresholds. Most importantly, we can compare the phase differences at which thresholds are predicted to be minimal and maximal (or approaching infinity) with the phase differences at which measured thresholds are minimal and maximal (or out of gamut). We can compare the general shape of the plots to test the predictions about the symmetry of the curves.

3 Common methods

The experiments in this thesis, those in sections 4, 5, and 6 in particular, require stimuli which generate specified cone excitations, in order to create desired modulations in later mechanisms. Such stimuli can be generated to a high degree of precision with apparatus that combines the outputs of three or four primary lights with known spectral output distributions (four primaries are needed to independently control rod as well as cone stimulation). This can be done with several types of equipment: the '8-channel' optical system is particularly suited for the first experiment, while the other experiments are carried out with a cathode-ray-tube (CRT) monitor. Both are computer controlled, and require calibration procedures.

3.1.1 Calibration

The four-primary Maxwellian-view photostimulator (the '8-channel') is an arrangement of optical components that can generate coloured visual stimuli with a centre-surround arrangement. Stimuli are generated by mixing eight primary lights (four for the centre and four for the surround). The primaries are light-emitting-diodes (LEDs) with peak outputs at 459nm (blue), 516nm (cyan), 561nm (green) and 664nm (red). The LEDs are behind neutral-density and interference filters, which provide narrow-band emission spectra. The radiant output of each LED is determined by the input voltage to the system which is converted to a pulse modulation and passed through the LEDs. Input voltage is controlled by computer, allowing precise levels and timing to be specified. A comprehensive description of this apparatus, including pictures, specifications and calibration procedures is given by Pokorny et al. (2004). Calibration allows the experimenter to specify the desired retinal photoreceptor stimulations, and have these converted into appropriate outputs of the primaries.

The principles of stimulus generation with the CRT, which is attached to a Cambridge Research Systems VSG 2/5 or ViSaGe, is similar. The outputs of the primaries are controlled by the voltages applied to the guns, although the CRS system provides built-in routines to perform some of the calibration. The CRT allows generation of stimuli with more complex spatial arrangements, however uses only three primaries. An overview of photoreceptor excitations in colorimetric systems such as these is given by Shapiro et al. (1996).

A simple calibration would involve measuring the emission spectra of the primary lights at their maximum outputs, and calculating the relative stimulations of the photoreceptors of a standard observer. Several factors complicate this, however, and a more involved procedure is needed. Firstly, the luminance output of the primaries as a function of input voltage is not linear across the whole range for either apparatus. Secondly, there are significant individual differences in pre-receptoral

filtering which mean that the standard observer sensitivity functions are not sufficient to allow adequate independent receptor stimulation.

The first stage is to measure the output of the primaries and construct a conversion matrix to allow conversion between the standard observer functions and LED outputs.

Emission spectra of the primary lights should be measured (at the eyepiece in the case of the 8-channel, or at a position close to the monitor in the case of the CRT) with a spectroradiometer with the instrument aiming at the illuminated stimulus. For the experiments in this thesis, a Gamma Scientific RadOMA GS-1271 was used, and this itself was calibrated with sources with NIST calibration. Each primary's emission spectrum is measured individually at maximum output. Measuring output at the eyepiece, in the case of the 8-channel, means that effects of the filters and other optical components of the system are taken into account. Example emission spectra of the 8-channel primaries are shown in Figure 14. Spectra of CRT phosphors are more broadband, and the spectrum of the red phosphor has an irregular form. Examples are shown in Figure 15.

When calibrating the 8-channel, the radiant output power of each LED, as a function of input voltage, should also be measured at the eyepiece with a photometer, clamping the instrument in front of the artificial pupil. Power output was measured in μW . This can be converted to trolands if desired by taking the product of the luminance (obtained by multiplying the corrected emission spectrum by $V^*(\lambda)$, see below) and pupil area (Wyszecki & Stiles, 2000, p. 101). Measurements are taken from, say, 20% to 100% of maximum output in steps of 10%, and from zero to 20% in steps of 1%, as the input-output function is unlikely to be linear in this region. These measurements can be used to generate a function relating input voltage to output power. This is necessary as the relationship is not linear. The CRS VSG system performs this step automatically for the CRT, measuring outputs at up to 256 input voltages with the CRS OptiCAL device.




Figure 14. Spectral output functions (in arbitrary units) of the four primary lights of the photostimulator. The curves are colour coded red, green, cyan and blue to represent the subjective colour of each light.




Figure 15. Spectral output functions (in arbitrary units) of the three primary phosphors of a CRT. The curves are colour coded red, green, and blue to represent the subjective colour of each light.

The spectroradiometer used is often inaccurate when measuring absolute radiance or luminance. For this reason, the output spectra of each LED should be linearly scaled so that the integrals of the spectra are in the same ratio as the radiance measurements taken at maximum output.

In order to obtain a measure of the relative stimulation of each photoreceptor by each primary light, sensitivity functions (or fundamentals) of the photoreceptors are needed. Many receptor sensitivity functions are available, and have been measured by a number of researchers. Many different

examples can be obtained in electronic form from Colour and Vision Research Laboratories database (www.cvrl.org). In most cases, the receptor functions are published normalised so that they peak at unity. They are often also available in terms of energy or quanta, and on linear or logarithmic scales. For the current purpose, energy measurements on linear scales are needed. The sensitivity functions should be scaled appropriately, and often this means scaling the long- and middle- wavelength sensitivity functions so that their sum corresponds to luminous efficiency function appropriate to those fundamentals. Throughout this document, the Stockman & Sharpe (2000) 2° fundamentals will be used (Stockman & Sharpe, 2000; Stockman, et al., 1999), so in this case, functions should be scaled so that $L(\lambda) + M(\lambda) = V^*(\lambda)$ (Sharpe, Stockman, Jagla, & Jägle, 2005, see Figure 15). Often, scaling factors are published with the functions. The S fundamental is most usefully scaled so that the function $S(\lambda)/(L(\lambda) + M(\lambda))$ peaks at unity. This allows colours to be specified in terms of L, M and S cone activations on the equiluminant plane in cone-receptor space, with axes corresponding to the cone opponent colour channels ($L/(L+M)$ and $S/(L+M)$) where $L+M$ is always equal to unity, and the maximal $S/(L+M)$ coordinate of the spectrum locus plots at 1.0. This is the diagram specified by Macleod and Boynton (1979) and in widespread use in vision science. Scaling the spectral functions in this way allows hues to be later specified in terms of these axes. The CIE standard scotopic luminous efficiency function, $V'(\lambda)$, can be used as the rod sensitivity function.




Figure 16. Photoreceptor photosensitivity functions and CIE standard luminance efficiency functions. $V^*(\lambda)$ is a corrected photopic luminosity function based on the Stockman & Sharpe (2000) spectral sensitivity functions, which are individually plotted as L, M, and S. $V'(\lambda)$ is a scotopic luminance efficiency function, equivalent to rod spectral sensitivity. The L-, M- and S-cone curves are scaled so that $L+M = V^*(\lambda)$ and the maximum value of $S/(L+M)$ is 1.0, as described in the text.

Once the output spectra of the primary lights and the photoreceptor sensitivity functions have been obtained and properly scaled, the emission and sensitivity functions must be individually multiplied and integrated. In practice, rather than calculus, the easiest way to perform these integrals is to multiply the emission spectra data by the sensitivity data wavelength-by-wavelength, and total the results. This is easily done for all combinations of spectra by matrix-multiplying the cone sensitivity and primary output spectra. This results in a matrix, which could have column headings L,M,S,R (for the long-, middle- and short-wavelength sensitive cones and the rods) and row headings R,G,B,C (for the red, green, blue and cyan primaries) and specifies relative receptor stimulations by each primary output. The CRT system lacks a cyan primary, and so only three receptor types can be controlled and the rods are omitted. The values of this matrix will need to be scaled so that the least-powerful primary is used to its full potential but no more. This can be easily done by dividing each value by the luminous efficiency of the least-powerful primary, or the scaled output spectra weighted by the luminous efficiency function. This results in the maximum output that each primary output is capable of being specified by a value of 1.0, which is very suitable for the purposes of programming software to specify output levels. This matrix can be used to convert from primary outputs to receptor

stimulations. This matrix should then be inverted, and transposed (remembering to transpose row/column headings). This new matrix is used to convert from intended values for receptor excitation to required primary output values. Values outside the range 0.0 - 1.0 are out of gamut.

Due to the receptor sensitivity fundamentals overlapping, it will never be possible to completely silence all but one class of photoreceptor. In particular, since the scotopic luminosity function overlaps all three cone sensitivity functions, it will not be possible to stimulate cones but not rods. It should be possible, however, to maintain rod stimulation at a constant level, while modulating cones, within a certain range. If an allowable tolerance is defined for rod modulation, a lower tolerance will result in a greater range of modulation for the cones. For the CRT system where only three primaries are available, we cannot specify rod modulation. Stimuli are chosen to be of sufficiently high luminance that the rods are effectively saturated, and can be assumed not to contribute to perception (Shapiro, et al., 1996).

Since these calibration values have been determined from the output spectra of the primaries at their maximum, and the input/output functions are not linear, it will be necessary to correct for this when specifying outputs in software. Fitting a function to plots of output power vs. input voltage (as measured earlier) will allow actual required input voltages to be specified from theoretically determined values when the system is assumed to be linear. This linearization process can be done either with the fitted function or a look-up table generated from it.

3.1.2 Individual observer calibration – 8-channel

As discussed above, there are significant differences between individuals' photosensitivity functions. These arise due to differences in prereceptoral filtering by the lens and macular pigment, as well as photopigment absorption spectra differences due to polymorphisms in the genes coding photopigments. We adjusted the calibration of the 8-channel individually for each observer.

The individual differences in sensitivities can be corrected by linear transformations of the standard functions. Previously, in experiments requiring similar calibration procedures, the difference between individual cone sensitivities and linear transformations of the standard functions has been shown to be very small (Pokorny, et al., 2004; Sun, Pokorny, & Smith, 2001). The standard observer sensitivity functions are simply scaled to correspond to the relative luminance values obtained by a calibration procedure, such as those outlined below.

Scotopic luminance colour matches can be used to measure differences in pre-receptoral filtering (Sun, et al., 2001). This is an exploitation of the fact that there are no known polymorphisms in

rhodopsin (the rod photopigment) in normal observers (Sung, et al., 1991), and it is these polymorphisms that lead to differences in receptor sensitivity functions. Any differences between real observer's scotopic luminance matches and those predicted by the standard observer functions are due only to differences in pre-receptoral filtering. This procedure cannot be used, however for primaries of long wavelength, however, where the scotopic and cone sensitivities are similar. This precludes the possibility of using a scotopic luminance match for the red primary in these systems.

For the red primary, a photopic colour match can be used. Sun, Pokorny and Smith et al. (2001) did this by asking observers to make a match between mixtures of pairs of primaries. The blue and green lights were mixed to the cyan and red lights. The relative outputs of the blue, cyan and green primaries were fixed at values that the standard observer data would predict to give a colour match, after being corrected for pre-receptoral filtering differences determined by the scotopic match, while the observer was allowed to adjust the red primary output. Any difference between the matching values for the red primary observed here and the theoretical values are likely to be due mainly to pre-receptoral filtering differences at this wavelength, but possibly due, to a small degree, to receptor photosensitivity differences.

In several previous studies (Pokorny, et al., 2004; Sun, et al., 2001), all the observers have been able to make colour matches with primaries set to theoretical values predicted for the standard observer, it is not necessary to perform the above procedure, which is time-consuming. A simpler procedure, from Pokorny, Smithson et al. (2004) is described next. They showed, by calculation, that this procedure produces corrected photoreceptor sensitivity functions that resulted in individual photoreceptors being stimulated individually with a very small uncertainty.

A most efficient way of obtaining the colour matching data is to match a mixture of the blue (459nm) and green (561nm) lights with a mixture of the cyan (516nm) and red (664nm) lights. There is a unique set of relative outputs from the four guns that will allow these pairs to match. The green primary is used as a reference and its luminance is held constant. The green primary was chosen as the reference as at its maximum, it is the least-luminous of the four. The observer adjusts three parameters: the luminance of the blue primary, the luminance ratio of the cyan and red primaries, and the total luminance of the cyan and red primaries, while toggling between the two mixtures. Comparing the relative outputs of each of the primaries after colour-matching, individually for each photoreceptor class, with the theoretical standard observer values for these colour matches enables a set of correction factors to be calculated for each observer. These specify the relative sensitivity of each of that observer's photoreceptor classes relative to the standard observer.

Once a set of scaling factors have been obtained, an individual conversion matrix can be calculated for each observer. The first matrix generated in the calculations above (with rows R, G, B, C and columns L, M, S, R) specifies relative receptor excitations by each primary light. Each column, therefore, can simply be multiplied by the observer's correction factor for that receptor. This is simpler than, but equivalent to, multiplying the receptor sensitivity function by this scaling factor before integration or matrix-multiplication. This matrix can then be inverted as previously described, resulting in an individual LMSR-to-RGBC conversion matrix for each individual observer.

The procedure above was performed with the LEDs that correspond to the centre field of the stimulus. Since there is variation in pre-receptor filtering with retinal eccentricity, a similar procedure on the surround field would likely result in different corrections. We wanted to reduce contrast at the border between the centre and surround fields, so to correct the surround each observer matched output of each surround LED to the corrected centre LED.

3.1.3 Individual observer calibration – CRT

As with the previous procedure for the 8-channel, we adjusted the calibration of the VSG and CRT system for the second and third experiments in this thesis.

We sought to counteract individual differences from the standard observer and hence to reduce any residual luminance signals generated by our stimuli. We used a modification of the minimum-motion procedure used by Cavanagh, Anstis & MacLeod (1987), in a similar way to Smithson & Mollon (2004), to equate the luminance of stimuli of different chromaticities. The procedure was performed with the same apparatus and viewing conditions as the experiments in sections 5 and 6 for which the calibration was performed. Our stimulus consisted of four frames presented on the monitor, each an annulus subtending a visual angle of 7° , divided radially into a square-wave grating (8 cycles per annulus) varying around EEW in either chromaticity (in the equiluminant plane) in the first and third frames, or luminance in the second and fourth frames. The phase of the spatial frequency grating in one frame was $\pi/2$ ahead of that of the previous, so the two luminance frames were in counterphase, as were the two chromatic frames. The observer was required to indicate the direction of rotational motion seen in the stimulus as the sequence of four frames was repeated at 5Hz (the whole sequence was repeated five times each second) for 4000ms. This special and temporal arrangement is shown in Figure 17.




Figure 17. The sequence of frames making up the minimum-motion stimulus. Note the $\pi/2$ phase relationship between each frame, highlighted with the dashed white line. This sequence was repeated at 5Hz for 4 seconds.

This procedure is not new, so a complete description of the rationale is not necessary but will be briefly summarised. If the two chromaticities differed in luminance, then the higher-luminance segments of in each of the four frames would be $\pi/2$ out of phase. The higher-luminance segments of the chromatic grating would be perceptually grouped with the higher luminance segments of the luminance grating, and vice versa. This would result in apparent rotation of the annulus in a direction determined by the direction of the apparent luminance difference of the two chromaticities.

The luminances of the chromatic components were adjusted in opposite directions, in accordance with the responses given by the observer, in a staircase procedure. The run ended after the direction of rotation reported changed from one trial to the next eight times. At this point, the luminance difference between the two chromaticities was assumed to have been nulled so that the direction of rotation was ambiguous. Two staircases were interleaved in each run, one beginning with the luminance offset in one direction, the other with the luminance offset in the opposite direction. The luminance offsets at the last six of the eight reversals, in both staircases of a run, were averaged and recorded. We used four pairs of chromaticities, spaced evenly around EEW, each in a separate run.

Two of the pairs were chromaticities lying in the cardinal directions from EEW, the other two pairs were along directions intermediate to the cardinal directions.

Adjusting the luminances of the two chromaticities in opposite directions is equivalent to tilting a line joining the two chromaticities (in DKL space) about EEW. We therefore adjusted the rotation of a plane around the L-M and S-(L+M) axes of DKL space to obtain a least-squares best-fit to a plane through the luminance-adjusted chromaticities determined by the procedure.




Figure 18. Example of adjusted luminances of pairs of chromaticities opposite EEW. The red circle lies in the equiluminant plane of the standard observer, and the black dots show the positions of chromaticities adjusted for subjective equiluminance with the minimum-motion task. The green grid represents a plane fitted to the points by rotation around the $L/(L+M)$ and $S/(L+M)$ axes and taken to be the corrected equiluminant plane of the observer.

This allowed us to estimate an amount and direction in which an observer's equiluminant plane was tilted relative to that of the standard observer. This correction was applied to all further stimulus generation for that observer for the experiments in sections 5 and 6.

4 Experiment 1: intermediate axes experiment

The material in this chapter has been published as a journal article. The full the citation is:

Lee, R. J., Mollon, J. D., Zaidi, Q., & Smithson, H. E. (2009). Latency characteristics of the short-wavelength-sensitive cones and their associated pathways. *Journal of Vision*, 9(12):5, 1-17, <http://journalofvision.org/9/12/5/>, doi:10.1167/9.12.5

The text and figures are reproduced in their entirety here as part of this thesis. Since the text has not been altered for inclusion here, some information in the introduction here will have already been discussed in the general introduction, section 1.

4.1 Abstract

There are many distinct types of retinal ganglion and LGN cells that have opponent cone inputs and which may carry chromatic information. Of interest are the asymmetries in those LGN cells that carry S-cone signals: in S-ON cells, S+ signals are opposed by (L+M) whereas, in many S-OFF cells, L+ signals are opposed by (S+M), giving $-S+L-M$ (Tailby, et al., 2008). However, the S-opponent pathway is traditionally modeled as $\pm[S-(L+M)]$. A phase lag of the S-cone signal has been inferred from psychophysical thresholds for discriminating combinations of simultaneous sinusoidal modulations along $\pm[L-M]$ and $\pm[S-(L+M)]$ directions (Stromeyer, et al., 1991). We extend this experiment, measuring discrimination thresholds as a function of the phase delay between pairs of orthogonal component modulations. When one of the components isolates the tritan axis, there are phase delays at which discrimination is impossible; when neither component is aligned with the tritan axis, discrimination is possible at all delays. The data imply that the S-cone signal is delayed by approximately 12ms relative to (L-M) responses. Given that post-receptoral mechanisms show diverse tuning around the tritan axis, we suggest that the delay arises before the S-opponent channels are constructed, possibly in the S-cones themselves.

4.2 Introduction

Are the signals of the short-wavelength sensitive (S-) cones delayed in their transmission to a central site where perceptual decisions are made? And, if so, where does the delay arise? These issues gain fresh interest from the discovery that there is a plurality of chromatically opponent pathways that carry signals originating in the S-cones.

To investigate the delays in the S-cone pathway, we measured thresholds for discriminating stimuli that offered different temporal modulations to the three cone classes, using a method introduced by Stromeyer, Eskew, Kronauer & Spillmann (1991). In our experiments the stimulus was modulated

concurrently along two axes of color space. In one condition, these axes were the cardinal axes identified by Krauskopf, Williams and Heeley (1982) and in the other they were intermediate to the cardinal axes. In one interval of a two-alternative temporal forced-choice (2ATFC) trial, the two modulations were in a phase relationship of ϑ and in the other interval they were in the relationship $\theta - \pi$. The observer's task was to distinguish these phase relationships. When $\theta = \pi/2$, the stimuli can always be distinguished – whether the axes are cardinal or intermediate. But in the case of the cardinal axes, and not the intermediate axes, it is possible to introduce a phase advance that renders the two stimuli indistinguishable. We use this result to estimate the delay of the S-cone signal and to consider the site at which the delay arises.

4.2.1 Early chromatic pathways

To extract chromatic information, signals from photoreceptors with different spectral sensitivities must be compared. These comparisons begin in the neural circuitry of the retina, and recent work has revealed fresh details of these circuits (e.g. Dacey & Packer, 2003). The number and diversity of cells that perform color opponent comparisons are greater than previously thought.

The biological substrate of the color mechanism that compares signals from the long- and middle-wavelength sensitive (L- and M-) cones is generally accepted to be the ON and OFF-midget ganglion cells (but see Calkins & Sterling, 1999; Rodieck, 1991). The midget ganglion cells receive opposed L and M input and project to parvocellular layers of the lateral geniculate nucleus (LGN) (Wiesel & Hubel, 1966). For the S-opponent chromatic pathway, the S-ON signal has for several years been identified with the small bistratified ganglion cell, which draws excitatory inputs from S-cones and inhibitory inputs from L- and M- cones (Dacey & Lee, 1994). But, in the retina, the S-OFF pathway has been difficult to identify and remains controversial (Klug, Herr, Ngo, Sterling, & Schein, 2003; Lee, Telkes & Grunert, 2005). The recent discovery of several additional, low-density, LGN-projecting ganglion cells has identified a possible substrate for the S-OFF signal and has additionally revealed further chromatically opponent S-ON ganglion cells. In brief summary, these are the intrinsically photosensitive, melanopsin-containing ganglion cells, which have an S-OFF opponent receptive field (Dacey, et al., 2005), the large sparse monostратified ganglion cells (S-OFF), and large sparse bistratified ganglion cells (S-ON) (Dacey, et al., 2002).

In the LGN, cells with S-OFF responses have been found reliably for many years. Valberg, Lee & Tigwell (1986) report that the L-cone input to such cells, if present, is synergistic with the S-cone input, and opposed to an excitatory M-cone input. In contrast, Tailby, Solomon & Lennie (2008) report that the S-OFF signal most commonly has the same sign as the M-cone signal and is opposed

by L-signals. In other S-OFF cells, the S-signal is antagonistic to both L- and M-signals. Thus it appears that there is some heterogeneity in the chromatic tuning of the S-OFF population.

The separability of pathways that respond to S-increments and S-decrements has also been demonstrated psychophysically (Krauskopf & Zaidi, 1986; Shinomori, et al., 1999). The inferred S-ON and S-OFF pathways additionally show several asymmetries: They differ in the ratio of L- to M- inputs (McLellan & Eskew, 2000) and in spatial summation areas (Vassilev, Mihaylova, Racheva, Zlatkova, & Anderson, 2003).

4.2.2 Delay of the S-cone signal

Schnapf, Nunn, Meister, and Baylor (1990) made direct measurements of the temporal response of Macaque photoreceptors. Only three S-cones were studied in detail, but their kinetics and sensitivities were roughly comparable to those of the L- and M-cones. Using silent substitution, Yeh, Lee and Kremers (1995) measured the temporal characteristics of signals of the three cone-types at the ganglion cell level. They found similar temporal modulation transfer functions for +L-M, +M-L, and +S-(L+M) cells, for both excitatory and inhibitory cone inputs.

Tailby et al. (2008) measured the temporal frequency selectivity of S+ and S- LGN cells that were driven either by S-cone isolating modulation or by achromatic modulation. The peak sensitivity for both cell-types was around 3Hz when driven by S-cone isolating modulation, but around 7Hz when driven by achromatic modulation. They argue therefore that the poor temporal resolution of the S-cone pathways arises early in the retina. However, they find no evidence for an accompanying *phase delay* of the S-cone signal, relative to the L- and M-cone signals provided by parvocellular cells.

So, perhaps the S-cone delay arises only after the LGN. In support of this, Cottaris and De Valois (1998), recording from cortical area V1 in the Macaque, found that S-opponent signals were available only after 96–135 ms, whereas L/M-opponent signals were available after 68–95ms. They additionally found that cells with late S-cone inputs dynamically change their chromatic tuning over time, which they suggest implies that the S-cone signal is amplified and delayed in V1 before being combined with L/M-opponent signals.

Consistent with the suggestion that S-cone signals arrive late at a central site, several authors have found longer reaction times to S-cone chromatic signals, than to L- and M-cone chromatic signals (McKeefry, et al., 2003; Smithson & Mollon, 2004). Analysis of visual evoked potentials (VEPs) also suggests a relative delay of the S-cone signal (Rabin, et al., 1994; Robson & Kulikowski, 1998). These studies all find a latency difference, but they differ in their estimates of the magnitude of that

difference. Furthermore the relative delay can be different for manual, saccadic or perceptual responses (Bompas & Sumner, 2008), indicating that delays are injected after the divergence of the pathways that support these different tasks.

Psychophysical measurement of the temporal impulse response functions (IRF) reveals a slower time course for S-cone signals than for luminance signals. Consistent with the physiological differences in S-ON and S-OFF pathways, the IRF for S-cone increments and decrements differs, with faster responses to increments (Shinomori & Werner, 2008).

4.2.3 The Peake Effect

Our experiment exploits a phenomenon first described by Miss Olive E. Peake (Hartridge, 1949, p. 148), and which we refer to as the Peake Effect. A rapidly presented sequence of hues appears noticeably different when the sequence is ordered in a clockwise (CW) sense around the hue circle (decreasing dominant wavelengths), compared to the reverse counterclockwise (CCW) sense. The flickering sequence appears to contain different predominant colors depending on the direction of procession. CW stimuli typically appear orange and sky blue, while CCW stimuli typically appear lime and magenta. The effect was later described by Davidoff, Aspinall, & Hill (1978) and Hill, Rodger, & Smalridge (1980). These authors constructed their stimuli by CW or CCW rotation of tubes or discs colored with brightness-matched colored papers.

The rapid rates of presentation required for the Peake Effect (approximately seven to twelve cycles around the hue circle per second) far exceed the rates at which temporal order judgments can be made. So, the Peake Effect is counterintuitive in that we might expect stimuli that differ only in the order in which the hues are presented, and not in the loci of hues themselves, to be indistinguishable. During the transmission of the neural response associated with CW and CCW stimuli to the point at which the perceptual comparison can be made, the signals must become distorted, such that the locus of hues visited is different for the two senses of modulation. It is possible to imagine several classes of mechanism through which such distortions might be introduced – for example, asymmetries in temporal masking between opponent mechanisms that are unequally distributed around the hue-circle (Spillmann, 1990; Spillmann & Neumeier, 1984), or differential latencies between color mechanisms, possibly at a stage as early as the cones (Hill, et al., 1980). In support of the latter suggestion, Stromeyer et al. (1991) found that introducing a phase advance to the modulation of the S-cones was sufficient to null this latency, rendering the CW and CCW stimuli indiscriminable. Models of the Peake Effect that rely instead on asymmetric tuning of mechanisms around the hue-circle predict that no such null is possible. Furthermore, Stromeyer et al.'s result is

consistent with separate evidence that the signal originating in the S-cones is delayed at some stage during transmission to a central site (see above).

The ability to counteract the Peake Effect by adjusting only the phase of an S-cone isolating modulation implies that there is a neural delay, and that it is associated with a mechanism that is aligned with the tritan axis. This result is particularly interesting in the light of the growing evidence summarized above that suggests that the S-opponent mechanism is not characterized by a symmetric bipolar pathway and that there are post-receptoral channels not aligned with the cardinal axes of Krauskopf, Williams and Heeley (1982).

We start by replicating Stromeyer et al.'s (1991) experiment, with component modulations that are aligned to the cardinal axes, and check the null. Secondly we extend the experiment to use component modulations that are oriented intermediate to the cardinal axes. The lights that comprise a hue-circle locus can be constructed from component modulations along any pair of orthogonal axes. However, when the component modulations are intermediate to the cardinal axes, we predict that no null should be possible, since we now cannot introduce a phase advance that is confined to the short-wave pathway.

4.2.4 Predictions

To introduce formally the present experiments, and to make specific predictions, we first describe the geometric properties of the chromatic loci that comprise our stimuli. Temporal modulation around a circular locus in the equiluminant plane of DKL color space (Derrington, et al., 1984) can be thought of as the sum of two component modulations along two orthogonal directions, temporally offset by a phase difference of $\pi/2$ radians.

First consider component modulations that are aligned to the cardinal axes. When the tritan modulation *leads* the $\pm[L-M]$ modulation by $\pi/2$ (or lags by $3\pi/2$) the sense of procession is CW; whereas when the tritan modulation *lags* by $\pi/2$ the sense of procession is CCW. Figure 19 illustrates these relationships. The $\pm[L-M]$ modulation is represented by a thin black line. The solid blue line represents the phase-advanced tritan modulation (and CW procession around the hue circle) and the solid red line represents the phase-delayed tritan modulation (and the CCW procession around the hue circle). Inverting one of the component modulations (by introducing a phase difference of π or, equivalently, by multiplying the sinusoid by -1) reverses the direction of procession.

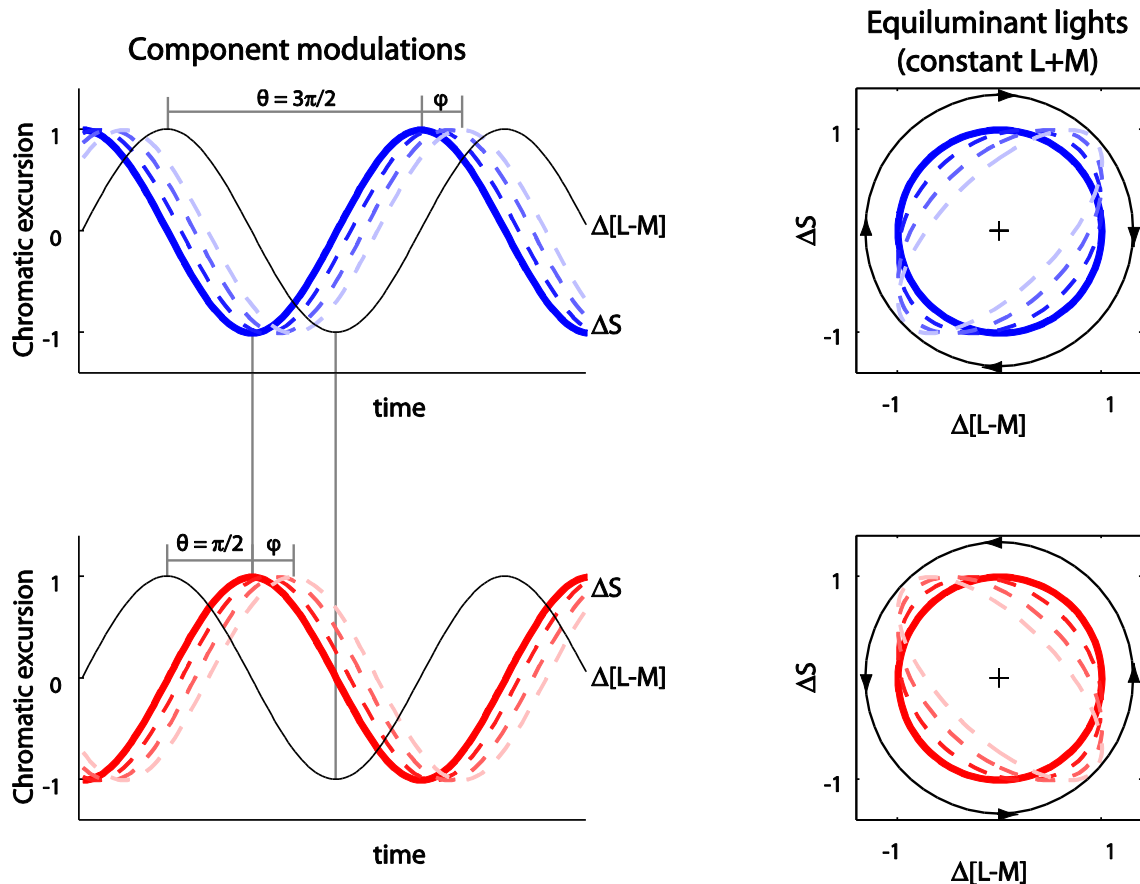


Figure 19. The sequences of chromaticities produced by sinusoidal $\Delta[L-M]$ and ΔS modulations in different phase relationships. The plots on the left show modulations as a function of time; the plots on the right show the loci of chromaticities that are visited. In the plots on the left the $\Delta[L-M]$ signal is used as a reference (thin, black, solid line). When this is combined with a ΔS modulation with a $3\pi/2$ phase shift (top-left panel, thick blue solid line), the resulting sequence of chromaticities follows a circular clockwise (CW) trajectory in the equiluminant plane of DKL color space (top-right plot, thick blue solid line). If the ΔS modulation has a phase of $\pi/2$ relative to the $\Delta[L-M]$ signal (lower-left panel, thick red solid line) the sequence of chromaticities plots out a circular counterclockwise (CCW) trajectory (bottom-right panel, thick red solid line). If the phase of the ΔS modulation is delayed from $3\pi/2$ (dashed curves in upper panels), then the chromaticities follow elliptical trajectories in color space. These ellipses are oriented with their major axis along the positive diagonal, becoming more eccentric with larger phase differences, up to a phase difference of π . For delays from $\pi/2$ (dashed curves in lower panels), the ellipses are oriented with their major axis along the negative diagonal. Importantly, inverting the S-cone modulation by introducing a phase difference of π reverses the sense of procession and the orientation of the ellipse (and defines the difference between upper and lower panels).

The dashed lines represent additional phase shifts (ϕ) of the S-cone signal, to simulate neural delays in the S-cone pathway. When combined with the $\pm[L-M]$ modulation the resultant chromatic loci are no longer circular. In the CW case, the phase difference between the component modulations becomes $\pi/2 + \phi$, and the chromatic locus is an ellipse oriented along the positive diagonal in DKL space. In the CCW case, the phase difference becomes $-\pi/2 + \phi$, and the elliptical locus is oriented along the negative diagonal. (The orientation of the ellipses is determined by the relative magnitudes of the component modulations; in a space in which the excursions are balanced, the ellipses are

oriented at 45°). The elliptical loci in Figure 19 provide a plausible explanation of the Peake Effect (Stromeyer, et al., 1991). The stimuli appear to flicker between the dominant hues at the extremes of the major axes of the elliptical loci: Between orange and sky blue for CW stimuli, and between lime and magenta for CCW stimuli. At the rates of presentation required for the effect, a neural delay of a few tens of milliseconds produces ellipses with the lengths of major and minor axes in the ratio 2:1.

If the response to tritan modulation is delayed, it should be possible to find a physical advance of the tritan modulation that nulls the neural delay, causing the resultant chromatic loci to become circular once more. Figure 20A shows the combined effects of putative neural delays of the S-cone response (φ) and physical phase differences between component modulations (θ). The blue and red lines in each plot represent the chromatic loci generated by combining chromatic modulations with phase differences of $\theta - \varphi$ and $\theta - \varphi - \pi$ respectively, which correspond to CW and CCW processions. The top row in Figure 20 represents zero neural delay of the S-cone response, and is a description of the physical stimuli. Subsequent rows represent neural phase delays (φ) of 0.1π , 0.2π and 0.3π . The successive plots from left to right in each row indicate loci that are generated from component modulations with physical phase differences (θ) of 0 to π , in steps of 0.1π . In each row, the plot outlined in orange corresponds to resultant phase differences of $\pm\pi/2$ (i.e. $(\theta - \varphi) = \pm\pi/2$), where CW and CCW loci overlap. In the extreme cases when the resultant signals are perfectly in phase or in counter-phase (i.e. $(\theta - \varphi) = 0$ or $(\theta - \varphi) = \pi$) the loci collapse to straight lines along the positive and negative diagonals. In these cases the 'CW' and 'CCW' loci are maximally discriminable.

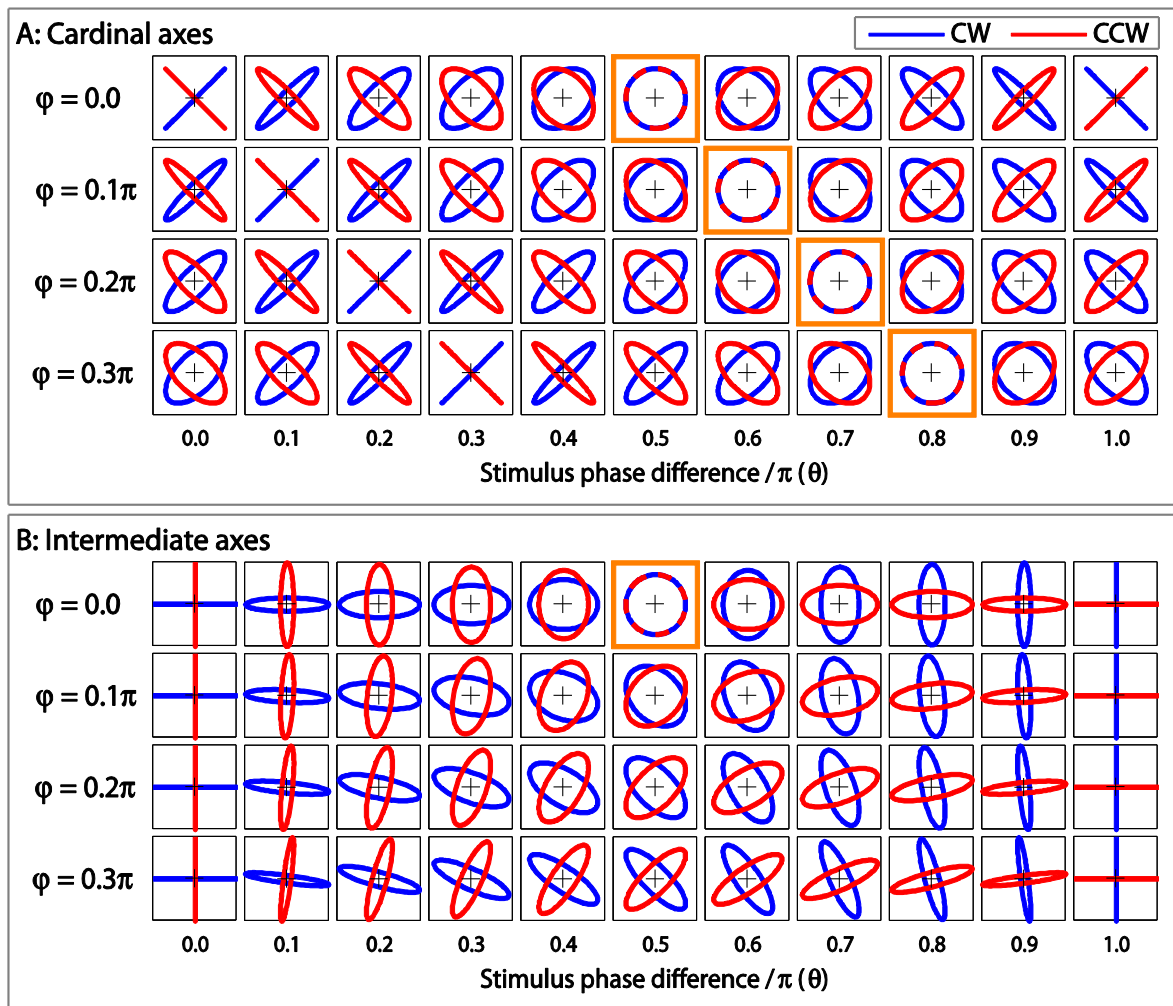


Figure 20. Simulated chromatic loci reaching a central site when the stimuli are composed of sinusoidal modulations along the cardinal directions of color space (upper panel, A) or along our intermediate axes (bottom panel, B). Each square represents the equiluminant plane of DKL color space. CW stimulus loci (where the component modulations have a phase difference of θ) are shown in blue, and CCW loci (where the component modulations have a phase difference of $\theta - \pi$) are shown in red. Each column of plots represents a different value of θ and each row represents a different simulated neural phase shift φ introduced by the S-cones or the putative $\pm[S-(L+M)]$ mechanism. Note how, in the cardinal axes condition, increasing φ translates the pattern of ellipses to the right but in the intermediate axes condition, increasing φ rotates the ellipses and the patterns remain symmetrical around $\theta = 0.5\pi$.

The panels of Figure 20A provide illustrations of the family of chromatic loci used by Stromeyer et al. (1991) to measure the latency of the S-cone response. By measuring discriminability of CW and CCW loci as a function of the phase delay between component modulations they were able to find a minimum of discriminability from which they inferred the neural phase delay.

When component modulations are along axes that are intermediate to the cardinal axes (our ‘intermediate axes’ condition), the stimulus loci are effectively rotated in DKL space. These loci are represented in the top row of Figure 20B

. A phase difference of $\theta = \pi/2$ or $\theta = 3\pi/2$ again results in a stimulus that changes chromaticity in time around a circular locus, and the corresponding plot is outlined in orange. For phase differences other than $3\pi/2$ and $\pi/2$, the stimulus loci are elliptical, but now the major axes align with the cardinal axes, and in the extreme cases when $\vartheta = 0$ and $\vartheta = \pi$, the locus collapses to a straight line along one or other cardinal axis. The subsequent rows of Figure 20B indicate the distorted loci that result from phase shifts of the S-cone response (φ) of 0.1π , 0.2π and 0.3π . In contrast to the cardinal axis plots shown in Figure 20A, there is now no stimulus phase difference (θ) at which the CW and CCW loci overlap. The maximally discriminable loci are in all cases at phase differences of $\theta = 0$ or $\theta = \pi$. Here the CW and CCW loci are aligned with the cardinal axes and stimuli would appear as flickering exchanges between cherry and teal or between violet and chartreuse. As the phase differences approach $\theta = \pm\pi/2$, CW and CCW loci become more similar, with the most similar occurring at $\theta = \pm\pi/2$. The similarity at $\theta = \pm\pi/2$ depends upon the simulated phase shift of the S-cone response (φ): Loci overlap when $\varphi = 0$, and are progressively dissimilar as φ increases.

In our experiment, in addition to measuring the discriminability of CW and CCW loci constructed from phase-shifted modulations along the cardinal axes, we test loci constructed from phase-shifted modulations along the intermediate axes.

4.3 Methods

Our methods follow those of Stromeyer et al. (1991). Throughout the investigation we sought to measure the threshold amplitude for discriminating between CW and CCW stimuli as a function of the phase difference (θ) between component modulations. We used a two-alternative temporal forced-choice (2ATFC) task. One interval contained stimuli processing in a CW sense, and consisting of component modulations with a phase difference of θ . The other interval contained stimuli processing in a CCW sense, with a phase difference of $\theta - \pi$. The order of the stimuli was chosen at random, with equal likelihood of CW followed by CCW or vice-versa.

During an experimental session, the amplitude of one component modulation was adjusted until the observer could no longer discriminate CW and CCW processions. Initial modulation amplitudes were made as large as possible within the color space gamut of our apparatus whilst maintaining perceptually equated amplitudes for both components (see below). At these amplitudes either component alone was well above detection threshold. Therefore, by adjusting only one modulation at

a time, we guaranteed that the CW and CCW modulations were always suprathreshold, even when their discriminability was at threshold.

There were several reasons why we adopted this approach, rather than requiring observers to manipulate θ in order to find a point at which discrimination was impossible: (i) the task becomes one of objective performance rather than phenomenal judgment, (ii) the variation of threshold with θ allows us to estimate the null point from the entire data set, (iii) we are able to test the full curve against a quantitative model, (iv) in the case of the intermediate axes, there proves to be no null point.

4.3.1 Apparatus

All stimuli were generated with a computer-controlled four-primary Maxwellian view optical system (Pokorny, et al., 2004). This apparatus was configured to present chromatic stimuli in a center-surround spatial configuration. The circular center field subtended 2° of visual angle, and the annular surround subtended 8° . Stimuli were viewed monocularly through an artificial pupil of diameter 3mm. The observer's head was held stationary with a bite bar.

The center and surround fields were generated independently by mixing two sets of four LEDs behind interference filters that provided a narrow-band output. The primaries had peak outputs at 459nm (blue), 516nm (cyan), 561nm (green) and 664nm (red). The relation between the intensities specified by the program and those produced by the diodes was established with a radiometer (UDT Instruments, Orlando, FL). A linearizing look-up table was then created to generate a mapping from the level requested in software to the luminance output of each LED.

Since four primaries were mixed, we were able to specify the relative stimulations of four photoreceptor types – the three classes of cone, and the rods. A transformation matrix was used to convert between desired photoreceptor excitations and outputs of the primaries.

The transformation matrix was derived from the cone sensitivity functions (Stockman & Sharpe, 2000; Stockman, et al., 1999), the sensitivity function of the rods ($V'(\lambda)$) and spectral measurements of the primaries, measured with a telescopic spectroradiometer (Gamma Scientific, San Diego, CA). The L- and M-cone fundamentals were scaled so that $L+M = V^*(\lambda)$, the photopic luminosity function appropriate for the Stockman & Sharpe standard observer fundamentals (Sharpe, et al., 2005), itself normalized to peak at unity. The S-cone fundamental was scaled so that $S/(L+M) = 1.0$ for the point on the spectrum locus that corresponds to the maximum S-cone stimulation (MacLeod & Boynton, 1979).

4.3.2 Calibration for individual observers

The relative scalings of the photoreceptor sensitivity functions were further adjusted to account for individual differences, using the color matching technique described by Pokorny et al. (2004). In this procedure, the observer makes chromatic matches between a mixture of the red and cyan primaries and a mixture of the green and blue primaries. The relative outputs of each of the center primaries after matching were compared to the matching values for the standard observer and used to scale the standard sensitivities to the primary lights. This procedure should correct for individual differences in photopigment sensitivities and macular pigment and lens density (Shapiro, et al., 1996). The scaling for each individual observer allows us to produce luminance-equated stimuli. To minimize the contrast at the border between the central and surround fields, each of the surround LEDs in turn was perceptually matched in brightness to the center LED having the same wavelength composition.

The intermediate axes modulations were intended to create balanced stimulation along S-opponent and L/M-opponent directions in color space. Since modulation in one channel does not have an intrinsically equivalent magnitude in the other channel, there is no accepted method of achieving this scaling. We chose to equate the perceived saturation of colors at different angles in color space around equal-energy-white (EEW). We presented a 1 Hz modulation around the hue circle and allowed the observer to adjust the relative amplitudes of the component modulations such that the number of perceived hues was maximized. The procedure was repeated for CW and CCW processions, and the optimal scaling was taken to be the average of four settings for each direction. A scaling factor of $L/(L+M) : S/(L+M) = 1.14:1$ was appropriate for both observers.

A mismatch in the scaling of stimuli in the S-opponent and L/M-opponent directions in color space is not predicted (according to the model outlined below) to change the critical phase delay in the cardinal axes condition, nor the phase delay at which discrimination is easiest in the intermediate axes condition.

4.3.3 Stimulus characteristics

The nominal luminance of the time-averaged stimulus (for both center and surround fields) was approximately 20cdm^{-2} , and the time-averaged cone-signals were those that would be elicited in the Stockman and Sharpe observer by an EEW spectrum. In a version of the MacLeod-Boynton chromaticity diagram that was constructed from these cone-sensitivities, the mean chromaticity coordinates were $[L/(L+M), S/(L+M)] = [0.6652, 0.0194]$.

Using the four-primary colorimeter we were able to hold constant the summed L- and M-cone stimulation to maintain the modulations in the equiluminant plane, and simultaneously to hold rod stimulation constant, whilst generating chromatic modulations.

The surround field was held steady in order to maintain uniform adaptation over an extended area of retina so that even small lapses of fixation would not have caused the test stimulus to fall on a non-adapted region. Test stimuli comprised temporal modulations of the center field, generated by combining two sinusoidal temporal modulations along orthogonal directions in the equiluminant plane. The extreme points visited had chromaticity coordinates $L/(L+M)_{\max}=0.6777$; $L/(L+M)_{\min}=0.6527$ and $S/(L+M)_{\max}=0.0304$; $S/(L+M)_{\min}=0.0084$.

To simplify description, we use the convention introduced by Derrington, Krauskopf and Lennie (DKL, 1984) in which the origin of the space corresponds to the white point. Our space is modified from that of DKL only in that it is constructed from the cone sensitivities of Stockman and Sharpe. Thus, in this space, sinusoidal modulations in $\Delta[L-M]$ and ΔS had amplitudes of $1.14 \times 2a$ and a respectively, where a set the amplitude of the modulations and was 0.0110 in our experiment.

The temporal frequency of the sinusoidal modulations was 10Hz. Note that one cycle of the component modulations corresponds to one complete procession around the elliptical locus. Each stimulus lasted 1356 milliseconds (13.56 cycles), including 339ms at the beginning and end during which the modulation depth of the flickering stimulus was ramped on or off with a raised cosinusoidal envelope. The temporal phase of the stimulus was randomized relative to the envelope on every presentation.

We use (i) component chromatic modulations that are aligned with the cardinal axes of the equiluminant plane (i.e. $\Delta[S-(L+M)]$ which simplifies to ΔS since $L+M$ is held constant, and $\Delta[L-M]$), and (ii) component modulations that are oriented intermediate to the cardinal axes (i.e. $\Delta[S+L-M]$ and $\Delta[S-L+M]$). The component modulations had a relative phase difference of θ ("CW") or $\theta-\pi$ ("CCW").

4.3.4 Experimental procedure

The observers' task was a 2AFC in which they had to select the interval that contained the stimulus with a θ phase difference, responding by pressing one of two buttons – each button corresponding to one of the intervals. Since the stimuli with a θ phase difference were not associated with a consistent appearance across stimulus conditions, it was necessary to give observers the opportunity to learn

which of the combinations they should be reporting. Feedback was given immediately in the form of a tone of high or low pitch.

After two practice trials, a staircase reduced the amplitude of one of the component modulations. We estimated the 71%-correct point on the psychometric function using a two-down-one-up staircase procedure, which reduced the amplitude of one component after two correct responses and increased it after one incorrect response, in logarithmic steps. The staircase progressed according to observers' performance in reliably identifying the correct stimulus, and we did not restrict the cues on which the observer was allowed to base his decision, although in most cases both observers reported that they were using the difference in predominant hues. Each staircase terminated after eight reversals, and we geometrically averaged the final six of these to estimate the threshold amplitude of the adjusted component modulation for this value of θ . Two independent staircases, with identical stimulus starting parameters, were interleaved in each run.

4.3.5 Experimental conditions

In each run we measured amplitude thresholds at a particular value of θ , to obtain threshold as a function of θ , from 0 to 1.2π . Measurements for different values of θ were run in a counterbalanced order. For each value of θ in each condition, we obtained at least two pairs of staircase estimates of threshold. In separate conditions we adjusted one or other of the component modulations. We ran counterbalanced sets of measurements for component modulations that were along cardinal axes and for component modulations that were along intermediate axes. So in total we obtained threshold as a function of ϑ for four combinations of stimulus conditions: (i) $\Delta[L-M]$ and ΔS , adjust ΔS ; (ii) $\Delta[L-M]$ and ΔS , adjust $\Delta[L-M]$; (iii) $\Delta[S+L-M]$ and $\Delta[S-L+M]$, adjust $\Delta[S-L+M]$; (iv) $\Delta[S+L-M]$ and $\Delta[S-L+M]$, adjust $\Delta[S+L-M]$.

Two observers, both authors, completed all conditions. They both had normal color vision and corrected-to-normal acuity.

4.4 Results

Figure 21 shows the complete dataset for our study. Thresholds for discriminating between the two processions are plotted against the phase difference between component modulations, for all four stimulus conditions for both observers. Panels A, B, E and F show data for the cardinal axes condition, panels C, D, G and G show data for the intermediate axes condition. Panels A and E show data obtained by reducing the amplitude of the ΔS modulation. Panels B and F show data obtained by reducing the amplitude of the $\Delta[L-M]$ modulation. Panels C and G show data obtained by reducing the amplitude of the $\Delta[S-L+M]$ modulation. Panels D and H show data obtained by reducing the

amplitude of the $\Delta[S+L-M]$ modulation. In each panel, the ordinate represents the amplitude of the staircase-adjusted modulation, relative to the fixed-amplitude modulation, at threshold. So values of 1.0 indicate equal modulation amplitudes, in DKL color space, for the two component modulations. Symbols show the geometric mean of at least two pairs of staircase endpoints. Error bars show one geometric standard deviation above and below the geometric mean. Smooth curves through the data points show predictions of the model described below.

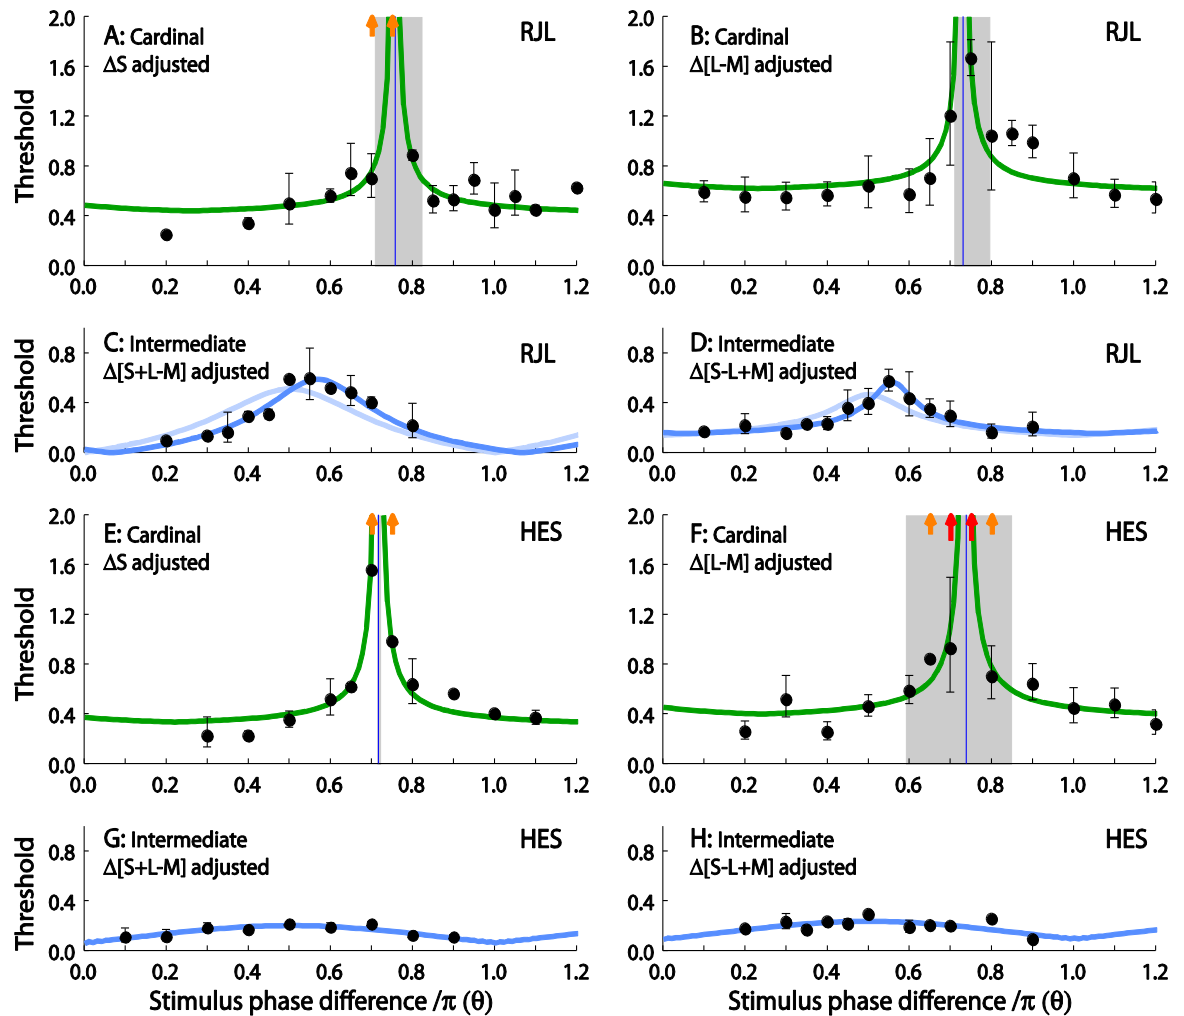


Figure 21. Measured thresholds for discriminating CW and CCW processes, as a function of θ , the phase difference between component modulations. Panels A-D represent results from observer RHL and panels E-H are for observer HES. Panels A, B, E and F show data from the experimental conditions when component modulations were along the cardinal axes. Panels C, D, G and H show data from conditions in which component modulations were along intermediate axes. The left-hand column of panels represents conditions in which the ΔS modulation or the $\Delta[S+L-M]$ modulation was adjusted by the staircase. The right-hand column of panels represents conditions in which the $\Delta[L-M]$ modulation or the $\Delta[S-L+M]$ modulation was adjusted by the staircase. Data points are the geometric mean of at least four staircases. Error bars indicate one geometric standard deviation above and below the geometric mean of several staircase results. Arrows along the top edges of each panel indicate that we failed to measure thresholds at the corresponding phase difference on some attempts (small orange arrows) or on all attempts (large red arrows). The green and blue lines represent predictions of the model described in the text and shown in Figure 22.

4.4.1 'Cardinal' axes

For both observers in the cardinal axes conditions, thresholds were strongly dependent on θ , the physical phase difference between component modulations. There were some phase differences at which it was not possible to make a threshold measurement, since the task was impossible at the maximum modulation amplitudes in gamut. Thresholds that we were unable to measure are marked as red arrows along the top edge of the plots, at the appropriate value of θ . Thresholds that we were able to measure on some occasions but not others are marked with orange arrows. In both cases these should be taken as underestimates of threshold. For both observers, the peak

values occurred at values of θ around 0.75π . In comparing panels A vs. B and E vs. F, it is also clear that changing which of the component modulations was adjusted by the staircase does not affect the pattern of thresholds.

4.4.2 'Intermediate' axes

Data for the intermediate axes conditions are shown in Figure 21 panels C, D G and H. The format of these plots is the same as for the cardinal axes conditions, but now the physical phase difference θ is the difference between the component modulations along our intermediate axes. Thresholds show some dependence on θ , but in the intermediate axes condition there was no value of θ for which the task became impossible. The highest thresholds were always close to $\theta = 0.5\pi$, and not as far from $\theta = 0.5\pi$ as the locations of the peaks in the cardinal axis case.

4.5 Models

4.5.1 Model description

The Lissajous figures presented in Figure 20B are constructed by considering the way in which CW and CCW loci would be distorted by a phase shift of the S-cone response. Here we derive a model to predict our measured discrimination thresholds as a function of the physical phase delay between component modulations. Our predictions are based on first estimating the signals presented to a central site for CW and CCW processions (represented schematically by the red and blue lines in Figure 20, assuming different delays of the S-cone response in successive rows) and then deriving a decision variable from the difference between signals for CW and CCW modulations.

Here, we make an explicit link between discrimination threshold and the similarity of predicted CW and CCW loci. A simple metric is to calculate the Weber contrast between the maximum excursion of the CW stimulus and the minimum excursion of the CCW stimulus, both measured as Euclidian distances from the time-averaged chromaticity in the equiluminant plane of DKL color space, and to use the inverse of this contrast to predict thresholds. Similar templates are obtained by calculating the root-mean-square radial difference between CW and CCW loci. In their gross features, these templates resemble that used by Stromeyer et al. (1991) based on the form $|\cos(\theta-\varphi)|-I$.

The derivation of the model is summarised in the Appendix. The most important parameter in the model is the phase shift of the S-cone response and varying this produces the family of templates presented in Figure 22. To fit these templates to our data, we also include vertical offset and scaling.

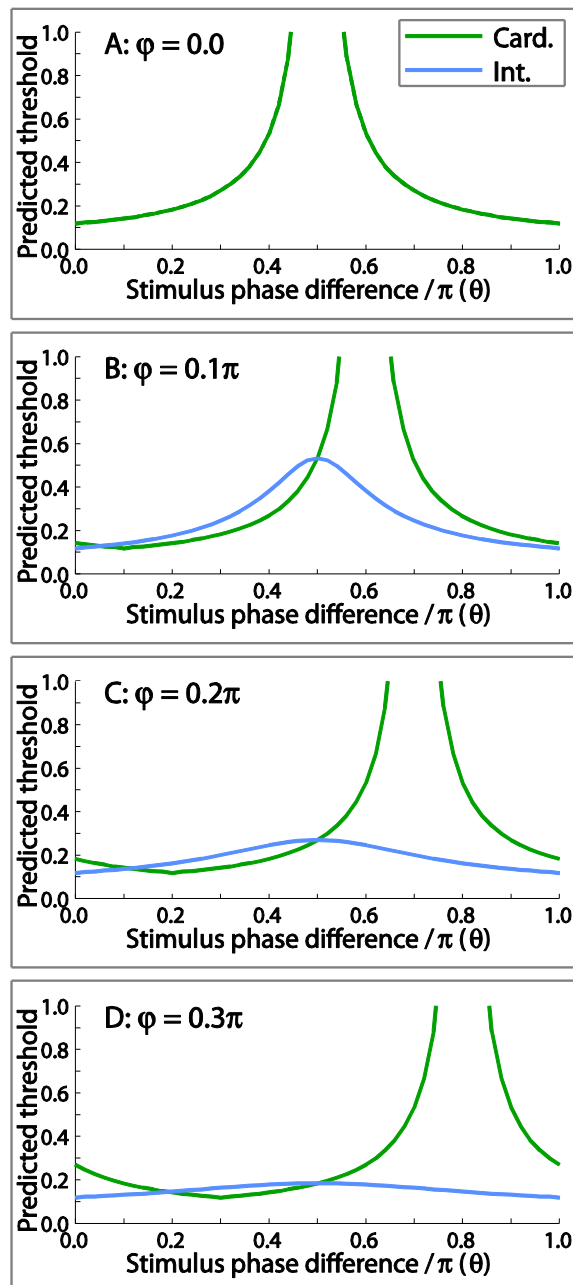


Figure 22. Predicted thresholds for discriminating between a stimulus in which component sinusoidal modulations have a phase difference of θ and a stimulus in which the phase difference is $\theta - \pi$, as a function of θ . Green lines show predictions when component modulations aligned with the cardinal axes, and blue lines show predictions for the intermediate axes. The uppermost panel assumes no neural delay between the response to ΔS and $\Delta[L-M]$, and each subsequent panel shows predictions for successive small increases in the phase delay of the ΔS response, denoted by φ .

If there were no phase delays, discrimination thresholds would tend to infinity at $\theta = \pi/2$, and would show the same pattern for the cardinal axes and intermediate axes conditions. If the phase delay (φ) between mechanisms aligned to the cardinal axes were nonzero, thresholds in the cardinal axis case would tend to infinity when $(\theta - \varphi) = \pi/2$. The location of the peak would therefore identify the phase delay. In the intermediate axes condition the peak threshold would be at $\theta = \pi/2$, irrespective of the neural phase delay. For a nonzero phase delay, discrimination would be possible at all values of θ , and thresholds would not tend to infinity.

4.5.2 Model fits

The smooth curves in Figure 21 show the results of a least-squares fit to the data obtained in each condition (see Appendix for best fitting values of the three free parameters: neural phase lag (φ) between the putative underlying mechanisms; vertical scale factor; and vertical offset). For both observers, the fit parameters are similar after adjusting either axis in the staircase. We tested this statistically using an F test to compare the difference between the total sum-squared residuals of fits to the data from each condition and the sum-squared residuals of a fit to the pooled data from both conditions combined, and no found no evidence to suggest a difference

4.5.3 Cardinal axes conditions

Smooth green lines in panels A, B, E & F show fits for the cardinal axes conditions. In terms of neural mechanisms, the phase lag parameter is the most interesting. It is determined by the position on the abscissa at which thresholds peak. For a neural phase lag $\varphi = 0$ the peak occurs at $\theta = 0.5\pi$, and in general the peak occurs at $\theta = 0.5\pi + \varphi$. To obtain 95% confidence limits on this parameter we stepped through a range of φ -values and found the limits within which the other two parameters could be adjusted to obtain fits not statistically different from the original, optimal one.

Importantly, the confidence limits on the neural phase lag parameter indicate that it was significantly different from zero, and therefore that the S cone pathway is significantly delayed relative to the L vs. M pathway. Averaging the estimates obtained from the condition in which $\Delta[L-M]$ was adjusted and those obtained when ΔS was adjusted, gives values of $\varphi = 0.25\pi$ and $\varphi = 0.23\pi$ for observers RJL and HES respectively. At the modulation frequency of 10 Hz, these phase differences correspond to a neural delay, of the S cone pathway relative to the L vs. M pathway, of 12.2ms and 11.4ms for observers RJL and HES respectively. Based on the highest and lowest confidence limits as described above, we estimate the delay to be between 10.1ms and 16.1ms for the two observers.

There are some subtle features of the data that are not well captured by the simple model. For example, in the cardinal axes condition, thresholds on the left-hand-side of the peak are lower than thresholds on the right-hand-side of the peak. This occurs for both observers, irrespective of whether we adjust the S or L/M modulation amplitude. The model predicts that thresholds are symmetric about the maximum threshold, as can be seen from the Lissajous figures (Figure 20) and from the template curves (Figure 22). Systematic differences between the data and the model suggest that there are measurable asymmetries in the thresholds. We consider a possible source of these asymmetries in the Discussion.

4.5.4 Intermediate axes conditions

Smooth faint blue lines in panels C, D, G & H show fits for the intermediate axes conditions. The neural delay is assumed to be between ΔS and [L–M], but the model now considers component modulations that are angled 45° to the cardinal axes in our color space, so the physical delay cannot directly null the neural delay. This model can account for the main properties of our data and has the same three free parameters as before: neural phase lag (φ); vertical scale factor; and vertical offset. Best fitting parameters are provided in the Appendix. For the intermediate axes condition the neural phase lag parameter is not well constrained by the data, since once the delay is greater than zero its effects are similar to the effect of the vertical scale factor. Upper confidence limits cannot be obtained for this parameter but lower confidence limits confirm that the delay is not zero.

In this form, the model for the intermediate axis condition is rigid in that the location of the peak is constrained to be at $\theta = 0.5\pi$. The data, particularly for observer RJL, however, show systematic deviations from this, with the maximum thresholds obtained slightly above $\theta = 0.5\pi$. A simple explanation of this discrepancy is that, in the intermediate axes condition, the two component modulations may not be perfectly matched in their effect on the opponent mechanisms. For example, one component might produce relatively more excitation in the S-opponent mechanism than the other. By allowing an additional parameter to vary in the model, namely the rotation of the intermediate axes relative to the cardinal axes (previously fixed at 45°), the model can account for the slight adjustment of the peak away from $\theta = 0.5\pi$. The rotations required to give the best fit for the data, for observer RJL and for each of the adjusted modulations, are 32° (ΔS adjusted), 21° ($\Delta[L-M]$ adjusted). These fits are shown with heavy lines in Figure 21 panels C and D. The extra variation accounted for by the additional parameter significantly reduces the overall variation, as determined by comparing the sum-of-squared-deviations of the data in the two models with an F test. Systematic deviations (i.e. an over-estimation of thresholds on the left shoulder of the peak, and an under-estimation of thresholds on the right of the peak) are improved by the extra parameter. For observer HES the additional free parameter did not significantly improve the fits. Again, best fitting parameters are included in the Appendix.

4.6 Discussion

4.6.1 Do S-cone signal delays explain the Peake Effect?

A simple model that assumes a phase shift of the S-cone response provides a good fit to the datasets, for both the cardinal and intermediate axes conditions. Thresholds in cardinal axes conditions tend to infinity at $\theta = 0.5\pi + \varphi$, whereas CW and CCW stimuli are discriminable at all phase delays between the intermediate component modulations.

In the cardinal axes condition, the equiluminant S-opponent modulation is a pure tritan stimulus – it modulates only the S-cones. If there were a delay only in the response of the S-cones themselves, the physical phase delay we introduce to the modulated lights would cancel the delay in the cones, and hence the input to all subsequent pathways. For the peak thresholds in the cardinal axis condition to be shifted from $\pi/2$, we assume that the delay must be apparent at the point at which the responses to the component modulations are combined to solve the perceptual task. Therefore, if the neural delay arises in the S-cones, we additionally infer that it must be propagated through the system, implying that the S-cone signal must remain confined to pathways that maintain a delay relative to the L/M-opponent signal.

Mechanisms later in the S-cone signal pathway may also contribute to the apparent delay of the S-cone signals revealed in our perceptual task. If these mechanisms are aligned with the tritan axis, as $\pm[S-(L+M)]$, delays in receptor and post-receptor stages simply add. The locations of the peaks in the cardinal axes data identify the combined delay, and there is no way in this experiment to tease apart the two sources.

However, if the chromatic tuning of these later mechanisms is not aligned with the tritan axis, delays injected here would be predicted to distort the CW and CCW loci, generating signals that might support discrimination of the CW and CCW stimuli at all values of θ . Since CW and CCW loci in the cardinal axes conditions became indiscriminable at critical values of θ , we found no evidence that delayed responses from non-tritan-aligned mechanisms were reliably available to observers in our psychophysical task. Indeed, if opponent mechanisms show a diversity of tuning around the tritan axis, and if these mechanisms contribute to performance in our task, then the neural delay that we null must be at a pre-opponent stage.

The purpose of our intermediate axes condition was to use component modulations that are unlikely to be aligned with the neural mechanisms that exhibit relative delays. As predicted, loci for the intermediate axis condition remained sufficiently distinct at all physical phase delays to support good discrimination.

4.6.2 Comparison to Stromeyer et al. (1991)

For the case of modulations on the cardinal axes, our replication yields the same pattern of thresholds as was found by Stromeyer et al. (1991). There were some differences between our experimental conditions and those of Stromeyer et al. However, our use of cone-isolating modulations and maintenance of rod excitation at a constant level did not change the results, supporting the conclusion that cone mechanisms were successfully isolated in the original study. A potentially more interesting difference is the surround adaptation in the two experiments: Our test stimulus had a time-averaged chromaticity of white and was surrounded by a white annular

field; Stromeyer et al.'s test field time-averaged to white, but their surround had a yellow-green appearance. However, we do not see any gross differences between the results obtained in the two cases. Our replication therefore suggests that the original results were not dependent on the biased surround adaptation.

Stromeyer et al. (1991) proposed that differences in the appearance of CW and CCW processions around a hue circle stem from a phase-delay of the S-cone signal. From their discrimination data they estimated that the S signal lagged the L/M opponent signal by about 0.41π - 0.50π at 10 Hz (a latency of about 21-25 ms). Our estimates of the phase delay of the S-cone signal from the cardinal axes condition are lower, with mean best-fitting values of 11 and 12 ms for the two observers. Adaptation state critically determines the temporal response of visual mechanisms (e.g. Stockman, et al., 2006), with faster responses obtained at higher adaptation levels. The S-cone quantal-catch for the stimuli used by Stromeyer et al. was higher than ours (by approximately one log unit), but the adapting chromaticity was similar, such that the relative adaptation states of color mechanisms were approximately equivalent in the two studies.

Critically, the interpretation of the Peake Effect in terms of a delay of the S-cone signal is further supported by the results in our intermediate axes condition. If there were no significant relative delays, discrimination of CW and CCW stimuli would be predicted to be impossible when loci were physically overlapping. This is not the case, and CW and CCW stimuli are discriminable at all phase delays between the intermediate component modulations. Model fits to our results confirm that the required value of the phase-delay parameter is significantly greater than zero.

4.6.3 Comparison to other estimates of the magnitude of the S-cone signal delay

There has long been interest in determining the relative latencies of the cone signals.

Physiological studies remain inconclusive: Some authors (e.g. Schnapf, et al., 1990; Yeh, et al., 1995) argue for similar temporal resolutions at early stages while other authors (e.g. Tailby, et al., 2008) suggest that the poor temporal resolution of the S-cone pathway arises early in the retina. However, few studies have reported phase delays explicitly and Tailby et al. at least show a dissociation between phase delay and temporal resolution.

In psychophysical studies the question is complicated by the potential involvement of multiple post-receptoral pathways. At the simplest level, residual luminance transients in the test stimuli can support rapid responses (Mollon, 1980), effectively hiding differences between stimuli intended to isolate chromatic pathways (Ueno, Pokorny, & Smith, 1985). Even equiluminant exchanges in L- and M-cone excitations can elicit responses in parasol ganglion cells (Lee, Martin, & Valberg, 1989), which might support performance in psychophysical tasks. Interestingly, this parasol cell response is not elicited by stimuli that isolate the S-cones, so S-cone stimuli might be

at a disadvantage in a behavioural task that compares equiluminant S-cone and L/M cone exchanges. Chatterjee and Callaway (2002) claim a small (10%) but consistent S-cone input to magnocellular neurons in Macaque LGN. However, Sun, Smithson, Zaidi & Lee (2006a, 2006b) find no significant S-cone input to magnocellular and parvocellular ganglion cells, at least under neutral adaptation. Using dynamic luminance noise to isolate chromatic channels, and using similar chromatic adaptation conditions to ours, Smithson and Mollon (2004) found that reaction times to liminal S-opponent and L/M-opponent stimuli showed mean differences of 13, 6 and 12 ms for their three observers. Other groups have found larger differences, of 40ms or more (McKeefry, et al., 2003), with similarly equated stimuli but without luminance noise.

Interestingly, Stromeyer et al. (1991) compared the results for their main discrimination task (in which one component modulation was S-isolating and the other was an equiluminant L–M-exchange) with results obtained when an S-isolating modulation was paired with a luminance (L+M) modulation. This revealed a further delay of S-cone signals in the inferred luminance pathway of 28 ms, consistent with other reports that delayed S-cone signals feed into luminance pathways (Lee & Stromeyer, 1989; Stockman, et al., 1987). They additionally used conditions in which observers were able to switch their criterion from a chromatic judgment to an achromatic judgment and again they found a dissociation between phase delays estimated for inferred chromatic and luminance pathways.

When comparing across photoreceptors with different spectral sensitivities it is essential to specify their relative adaptation states. Furthermore, perceptual latency depends on signal strength, with stronger stimuli eliciting more rapid responses (Mollon & Krauskopf, 1973; Piéron, 1931). Paying particular attention to these two factors, Blake, Land and Mollon (2008) have demonstrated that when the S-cones are given a “fair” chance (i.e. the S-cone adaptation state is equated to the L- and M-cone adaptation states such that the thresholds for both are equivalently elevated above absolute threshold, and excursions of equal discriminability are used), there is negligible delay in the phase of a moving S-cone stimulus relative to a long-wave stimulus. The contrasts of our component modulations were matched in units of discrimination threshold, but the equal-energy time-averaged adaptation placed the S-cones in a relatively more adapted state than the adaptation used by Blake et al. (2008). More adapted receptors are predicted to have shorter latencies so this difference in relative adaptation is unable to account for the relative delay we measured.

Given that the absolute response latency depends on adaptation state, it is interesting to determine the way in which relative latencies vary with adaptation level for different latency estimates. Irrespective of whether the S-cone delay was assessed against a chromatic modulation or against a luminance modulation, Stromeyer et al. (1991) found that a one log unit increase in

the adapting level of the S-cones decreased the S-cone delay by 17ms. The relative adaptation state of the S-cones did not change the *difference* between the S-cone delay relative to the inferred luminance pathway and the S-cone delay relative to the inferred chromatic pathway. A parsimonious explanation of this result is that the adaptation introduces latency differences early in the visual pathway, probably in the S-cones themselves.

4.6.4 Alternative models

The gross characteristics of CW and CCW discriminations in the cardinal axes conditions – namely that we reliably find that a simple advance of the tritan modulation is sufficient to render CW and CCW loci indiscriminable – are well captured by assuming a delay in a tritan aligned mechanism. However, as noted above, there are small systematic differences between our data and the model. In particular, the model cannot account for asymmetries in thresholds around the peak.

Our model assumes bipolar opponent mechanisms, where the tuning of ON and OFF sub-mechanisms is co-linear in DKL space. To account for the asymmetries in our data, we might appeal to additional cone opponent mechanisms that may not be aligned to the cardinal axes. Contributions from such mechanisms, perhaps with their own phase delays, have the potential to distort CW and CCW loci differently for phase advances and phase delays relative to the null, thus causing asymmetric thresholds around the peak.

Several authors have suggested that color mechanisms might be more accurately described as unipolar mechanisms, with rectified outputs, rather than bipolar mechanisms (e.g. Chen, Foley, & Brainard, 2000; Krauskopf, et al., 1982). For the S-opponent system the mechanisms that respond to increments and decrements have been shown to be separable (Shinomori, et al., 1999), and to respond with different latencies (Shinomori & Werner, 2008). Given that different classes of cell respond to opposite directions along the same axis of color space, we must also acknowledge that responses of any mechanism may be non-linear, and that these non-linearities may differ between cells responding to chromatic changes in different directions (e.g. Giulianini & Eskew, 2007; Zaidi & Shapiro, 1993).

It is notoriously difficult to decide how best to model the combination of signals from different mechanisms. Indeed, there has been a long debate about this issue for the combination of information from the cardinal mechanisms (see Eskew (2009) for review). In the model that we fit to our data, the signal reaching central sites is the vector sum of modulations along the $\Delta[L-M]$ and $\Delta[S-(L+M)]$ axes. One way to extend the model to incorporate a chromatic mechanism with a preferred direction intermediate to the cardinal axes (e.g. the S-OFF LGN cells characterized by Tailby et al. (2008)) is to assume inputs from rectified independent mechanisms (+L-M; -L+M; +S-L-M; and -S+L-M). Another scheme would be to first combine multiply-tuned opponent cell

signals to generate only two color opponent dimensions, and then treat these as orthogonal axes of a 2D chromaticity space, as in the classical model.

These complications present a vast array of possible models, each with many unknown parameters. Detailed prediction of psychophysical thresholds is impossible from such modeling. Instead, we concentrate on correspondence between qualitative features of the data and models, namely the existence, location and symmetry of peaks in threshold as a function of phase difference between component modulations.

On the one hand, the slight asymmetries in our data indicate the involvement of mechanisms more complex than those captured by our simple model. However, given that we find substantial failures of discrimination in the cardinal axes condition, it is likely that the neural delay arises in mechanisms tuned to the tritan axes, and not in a family of mechanisms with diverse tuning. The simplest explanation is that delays are injected at the receptor level, but injection at a unitary tritan-aligned, post-receptoral mechanism cannot be ruled out. The results of the discrimination experiment imply that the delay of the S-cone signal propagates through the visual system to at least the neural locus at which responses to the component modulations are combined. It is an open question whether later stages are perceptually calibrated to represent physical simultaneity of events in the world.

4.6.5 Summary

Every model we have considered predicts that, if there are no relative delays in the pathways supporting the discrimination of CW and CCW stimuli, we should see large peaks at $\pi/2$ and $3\pi/2$, in both the cardinal and intermediate axis conditions. We confirm Stromeyer et al.'s (1991) finding that the peaks are displaced from these locations for the cardinal axes condition, and further show that peaks are much reduced in the intermediate axes condition, a result that similarly is consistent with there being significant neural delays.

Given that post-receptoral mechanisms show a diversity of tuning around the tritan axis, and assuming these mechanisms contribute to performance in this task, we suggest that strong peaks in the cardinal axis data are not consistent with delays injected at a post-receptoral site. The source of the delay is therefore likely to be before the S-opponent mechanisms are constructed, possibly in the S-cones themselves. Alternatively, if the delay does arise late in the pathway, then our results suggest that the only channels contributing to the task are ones aligned with the cardinal axes of color space.

4.7 Appendix

4.7.1 Model equations

The equation predicting discrimination threshold in our model takes the general form

$$T = \left(\frac{\max(r_{CW})}{\min(r_{CCW})} - 1 \right)^{-1} \quad (43)$$

where T is proportional to threshold and r_{CW} and r_{CCW} are functions describing the radii of the CW and CCW ellipses, respectively:

$$r_{CW} = ((\sin(t) \cos(\alpha) - \sin(t - \theta) \sin(\alpha))^2 + (\sin(t + \varphi) \sin(\alpha) + \sin(t - \theta + \varphi) \cos(\alpha))^2)^{\frac{1}{2}} \quad (44)$$

$$r_{CCW} = ((\sin(t) \cos(\alpha) - \sin(t - \theta - \pi) \sin(\alpha))^2 + (\sin(t + \varphi) \sin(\alpha) + \sin(t - \theta + \varphi - \pi) \cos(\alpha))^2)^{\frac{1}{2}} \quad (45)$$

where α is the rotation of the stimulus component axes relative to the cardinal ones, θ is the phase difference between the two components, φ is a phase delay introduced by the S-cones or any tritan-aligned mechanism and t is a parameter such that $0 < t < 2\pi$.

In the cardinal axes condition, these equations simplify to

$$r_{CW} = (\sin^2 t + \sin^2(t - \theta + \varphi))^{\frac{1}{2}} \quad (46)$$

$$r_{CCW} = (\sin^2 t + \sin^2(t - \theta + \varphi - \pi))^{\frac{1}{2}} \quad (47)$$

and the minima and maxima can be found by taking the derivatives:

$$\frac{dr_{CW}}{dt} = (\sin^2 t + \sin^2(t - \theta + \varphi))^{-\frac{1}{2}} (\sin t \cos t + \sin(t - \theta + \varphi) \cos(t - \theta + \varphi)) \quad (48)$$

$$\frac{dr_{CCW}}{dt} = (\sin^2 t + \sin^2(t - \theta + \varphi - \pi))^{-\frac{1}{2}} (\sin t \cos t + \sin(t - \theta + \varphi - \pi) \cos(t - \theta + \varphi - \pi)) \quad (49)$$

The minima and maxima of the radius functions occur when

$$t = \frac{\theta - \varphi}{2} + \left(k + \frac{1}{2}\right)\pi \quad (50)$$

or

$$t = \frac{\theta - \varphi}{2} + k\pi \quad (51)$$

where k is an integer $k \in \mathbb{Z}$.

Substituting these into above gives

$$T = \left(\frac{\left(\sin^2 \left(\frac{\theta - \varphi + (2k+1)\pi}{2} \right) + \sin^2 \left(\frac{\varphi - \theta + (2k+1)\pi}{2} \right) \right)^{\frac{1}{2}}}{\left(\sin^2 \left(\frac{\theta - \varphi + (2k+2)\pi}{2} \right) + \sin^2 \left(\frac{\varphi - \theta + 2k\pi}{2} \right) \right)^{\frac{1}{2}}} - 1 \right)^{-1} \quad (52)$$

for $0 < (\theta - \varphi) < \frac{\pi}{2}$ and $\frac{3\pi}{2} < (\theta - \varphi) < \pi$ and

$$T = \left(\frac{\left(\sin^2 \left(\frac{\theta - \varphi + 2k\pi}{2} \right) + \sin^2 \left(\frac{\varphi - \theta + 2k\pi}{2} \right) \right)^{\frac{1}{2}}}{\left(\sin^2 \left(\frac{\theta - \varphi + 3k\pi}{2} \right) + \sin^2 \left(\frac{\varphi - \theta + (2k-1)\pi}{2} \right) \right)^{\frac{1}{2}}} - 1 \right)^{-1} \quad (53)$$

for $\frac{\pi}{2} < (\theta - \varphi) < \frac{3\pi}{2}$.

In the intermediate axes condition, the derivatives of the radius functions do not simplify, leaving equations that are difficult to solve. Our templates for the intermediate axes were generated numerically.

4.7.2 Model fit parameters

Table 2. A summary of the parameters of the models fitted to data across all experimental conditions

Observer	Adjusted modulation:	Cardinal axes		Intermediate axes, fixed rotation		Intermediate axes, free rotation	
		ΔS	$\Delta L-M$	$\Delta S+L-M$	$\Delta S-L+M$	$\Delta S+L-M$	$\Delta S-L+M$
RJL	S-(L+M) phase delay / π	0.258	0.231	0.169	0.085	0.140	0.075
		(0.208 < φ < 0.322)	(0.208 < φ < 0.294)	(0.033 < φ)	(0.031 < φ)	(0.075 < φ)	(0.049 < φ)
	S-(L+M) time delay (ms)	12.9	11.6	8.5	4.3	7.0	3.7
		(10.4 < t < 16.1)	(10.4 < t < 14.7)	(1.7 < t)	(1.5 < t)	(3.7 < t)	(2.4 < t)
	Vertical scale factor	0.062	0.064	0.383	0.102	0.290	0.069
	Vertical offset	0.445	0.624	0.005	0.144	0.001	0.158
	axis angle (°)	0	0	45	45	31.4	20.6
	RMS error	0.095	0.113	0.100	0.061	0.041	0.022
HES	S-(L+M) phase delay / π	0.218	0.239	0.419	0.385	0.416	0.384
		(0.214 < φ < 0.221)	(0.201 < φ < 0.279)	(0.043 < φ)	(0.013 < φ)	(0.047 < φ)	(0.020 < φ)
	S-(L+M) time delay (ms)	10.9	11.9	20.9	19.235	20.8	19.2
		(10.7 < t < 11.0)	(10.1 < t < 13.9)	(2.2 < t)	(0.7 < t)	(2.3 < t)	(1.0 < t)
	Vertical scale factor	0.070	0.085	0.955	0.644	0.949	0.669
	Vertical offset	0.339	0.403	0.065	0.096	0.061	0.091
	axis angle (°)	0	0	45	45	44.5	44.5
	RMS error	0.065	0.089	0.018	0.038	0.017	0.037

4.7.3 Further comments

The following discussion is a result of comments made by others since the material in this chapter was published. As such it is appropriate to place it in this appendix, rather than alter the text above.

We base our estimates of the phase difference between mechanisms carrying ΔS and $\Delta[L-M]$ signals on the φ parameter obtained by fitting the model to data from the cardinal axes condition (see section 4.5.3). We do not comment on the magnitude of the phase difference or temporal delay that might be estimated from values of φ obtained from the intermediate axes condition data. As can be seen from the values given in the appendix, the estimates of temporal delay obtained from the intermediate axes condition are quite different from those obtained from the cardinal axes condition and have more variability. As mentioned in section 4.5.4, adjusting the φ and vertical scale factor parameters has a similar effect on the predicted threshold curve. Because of this, we do not believe that the φ or temporal delay magnitudes can be reliably estimated from the intermediate axes data. The primary purpose of the study was not to generate another estimate of the temporal lag of the ΔS mechanism, but to show that it was not possible to null this delay unless our stimuli were constructed from modulations along cardinal chromatic axes. This leads to our conclusion that the delay occurs early in the pathways and in a tritan-aligned mechanism.

Stromeyer et al. (1991) performed their experiment (on which our cardinal axes condition is modelled) with stimuli modulated at several temporal frequencies. We confine our stimuli to 10Hz modulations. Since this was sufficient to demonstrate that nulling a phase difference between modulations along intermediate chromatic directions, we do not believe that repeating the experiment with stimuli of other temporal frequencies would offer any new information, other than estimates of the phase difference and temporal delay at these frequencies. This was, as discussed above, not the goal of our study.

5 Experiment 2: adaptation of central chromatic mechanisms

5.1 Introduction

The previous study provides further evidence that signals in the pathways carrying S-cone signals are delayed by some milliseconds relative to signals in the L-M pathway. In this experiment we attempt to exploit this temporal difference and, with a similar stimulus to that used in the previous experiment, investigate the stage at which chromatic signals from all three cone classes are opponently combined. The prevailing model of chromatic processing, based on physiological evidence from the last 30 years, suggests that this does not occur until some stage in the cortex (see section 1.1.1.6).

This study attempts to use psychophysical methods to provide evidence for the existence of a stage, which we refer to as the central site, where signals from the early chromatic pathways are combined into a distinct set of chromaticity encoding mechanisms. This would most likely occur in the visual cortex, and the possible neural substrates of this stage are discussed in section 1.1.1.6. Traditionally these mechanisms have been contrasted with the cardinal chromatic mechanisms that seem prevalent in retina and LGN. They have become known as “higher order” colour mechanisms (Krauskopf, et al., 1986). This term has come to carry two meanings, both that the mechanisms exist at central stages of the visual system (as opposed to early stages) and that these mechanisms are more numerous and have narrower tuning than the cardinal mechanisms. The evidence for and against such mechanisms is comprehensively reviewed by Eskew (2009). It is also possible that the off-axis chromatic tuning of cells in LGN (e.g. Tailby, et al., 2008) could contribute to such “high order” mechanisms, though on the other hand these mechanisms may not contribute to “colour” processing explicitly. Importantly, for the present experiment, we are interested in colour mechanisms that are tuned differently from the cardinal mechanisms.

We aim to demonstrate adaptation in these mechanisms in which signals from S-cone have been combined with L/M opponent signals. Adaptation occurs in many sensory systems and is discussed in more detail in section 1.3. We attempt to show two different patterns of discrimination thresholds after viewing two different stimuli, when these stimuli are designed such that they are equivalent at stages of chromatic processing before mechanisms that respond to non-cardinal chromaticities are constructed. Krauskopf et al. (1982) demonstrated selective adaptation in the cardinal mechanisms with sinusoidal modulations in their corresponding chromatic directions. We attempt to adapt the hypothetical higher order mechanisms with a stimulus, similar to the one used in the previous experiment and described in section 2, which we can manipulate to create two conditions. Each of these conditions will present different relative stimulations to the central mechanisms while presenting equivalent modulations to the early

cardinal mechanisms. Both conditions consist of an equiluminant L/M exchange and an S-cone isolating modulation that isolates the putative bipolar $\pm[S/(L+M)]$ mechanism.

We use component modulations along the cardinal directions of equiluminant colour space, with relative phase differences of $\pi/2$ and $-\pi/2$ ($3\pi/2$). These are the stimuli that trace a circular locus in the equiluminant plane, one in a clockwise (CW) direction and the other in a counterclockwise (CCW) direction. The sets of lights visited by the stimulus in both directions of modulation are the same, but the order is simply reversed between the two conditions. If this stimulus is presented at a high enough modulation frequency ($>5\text{Hz}$) it is not possible to distinguish the order in which hues are presented. Both stimuli appear as a coloured flicker. However, the predominant colours in the flicker appear different between the two stimuli.

If we take into account the delay in the S-opponent pathway described above, and model this as an additional phase difference (which should be constant at any given frequency), identical in both CW and CCW modulations, we find that the stimulus loci reaching the central site at which the two component signals are combined are no longer circular. This combination of sinusoidal modulations, and the shapes of the resulting elliptical loci were discussed in the previous experiment, however the smaller subset of stimuli should be summarised again. The stimuli in both the CW and CCW conditions maintain their general direction of procession around the loci, but the loci are now elliptical and oriented differently. In the CW direction, the elliptical locus is aligned with its major axis along the positive diagonal of colour space (i.e. $+L-M+S$ to $-L+M-S$) and in the CCW direction, the elliptical locus is aligned with its major axis along the negative diagonal of colour space (i.e. $-L+M+S$ to $+L-M-S$). Figure 23 shows the relationships between the physical modulations we present and the theoretical modulations reaching the central site.




Figure 23 The phase relationships between the physical stimuli designed to modulate the L-M opponent mechanism (black lines in left-hand panels) and the S-(L-M) opponent mechanism (solid coloured lines in left-hand panels), and the modulations reaching a central site (dashed coloured lines in left-hand panels). The locus of chromaticities resulting from the modulations are shown in the right hand panels, and the direction of procession is CW in the top panels and CCW in the bottom panels. Note the differently oriented ellipses in the CW and CCW conditions.

The stimulation provided to the three cone types and both cardinal mechanisms, when each mechanism is considered independently, is the same in both of these experimental conditions. Furthermore, the time-averaged excitation in each of these mechanisms is the same. Therefore, prolonged exposure to either condition should not differently adapt any of the cone types or cardinal mechanisms. However, since the resultant loci imply that different sets of chromaticities are visited in the different conditions, mechanisms not tuned to the cardinal directions will be adapted in different ways.

It is readily apparent that the two directions of procession lead to different sets of perceived colours. Within a range of frequencies ($\sim 7\text{Hz}$ to $\sim 12\text{Hz}$), the dominant colours seen when viewing the two stimuli are different. This effect was discovered some time ago (see section 1.4.1.5).

After adaptation to either the CW or CCW stimulus, we probe discrimination thresholds for various hues. If there are many mechanisms at the central site, each selective to a small range of chromaticities, we expect to see patterns of adaptation in which thresholds are raised most to those colours along the major axis of the adapting stimulus loci – opposite in the two adapting conditions. Even if there are only two mechanisms at the central site, but their preferred colour

directions are rotated relative to the cardinal axes defined by the opponent mechanisms, we would expect to see different patterns of adaptation after adaptation to CW or CCW stimuli. As mentioned above, this sort of mechanism, which combines signals from earlier mechanisms is usually referred to as a “higher-order mechanism”. Eskew (2009) points out that the existing evidence for such mechanisms is weak, so our experiment is potentially important in revealing non-cardinal adaptable mechanisms, if they exist.

5.2 Methods

All stimuli were displayed on a Sony Trinitron monitor, controlled by a VSG 2/5 graphics board (Cambridge Research Systems, Rochester, UK). Emission spectra of the red, green and blue phosphors were measured with a spectroradiometer (RadOMA GS-1271, Gamma Scientific, San Diego, CA), and then relative outputs of these were specified to generate desired cone excitations, calculated using the Stockman & Sharpe cone fundamentals (Stockman & Sharpe, 2000; Stockman, et al., 1999). The calibration procedures, and calculations involved in specifying desired cone excitation and opponent mechanism modulations are described in detail in section 3. The calibration was adjusted for each observer as described in section 3.1.3. The monitor was viewed from a distance of one meter, had dimensions of 0.36×0.30m and subtended a horizontal visual angle of 22°, and was the only source of light in the room. The image had dimensions of 800×600 pixels (each pixel measured 0.5mm) and the frame rate was 120Hz.

5.2.1.1 Stimuli

We measured chromatic discrimination thresholds with a probe stimulus, which consisted of a circle subtending 4° visual angle, divided into quadrants separated by 10 pixels ($\approx 0.2^\circ$), see Figure 24. Two of the adjacent quadrants were given the same hue, while the remaining two were given the complementary hue. The two hues were assigned so that either the top two quadrants had the same hue and the bottom two quadrants had the opposite hue or the left two and right two had opposite hues. The observer’s task was to differentiate between a circle chromatically divided horizontally (i.e. the top quadrants had opposite hues to the bottom) or vertically (i.e. the left quadrants had opposite hues to the right). Luminance noise was introduced so that any two of the quadrants were more luminous than the background, and the other two quadrants were less so, by an amount proportional to the saturation. This was done to render any remaining luminance differences between the two chromaticities unreliable for performing the task.




Figure 24. Representation of two examples of the probe stimulus. In the left example, the disc is chromatically divided horizontally, so that the top and bottom halves differ in chromaticity. In the right example, the disc is chromatically divided vertically, so that the left and right halves differ in chromaticity. Each quadrant has a luminance offset above or below the background.

The adapting stimulus was a disc, subtending 11° whose chromaticity was varied over time. The size of this stimulus influenced our choice of the size of the minimum-motion stimulus in the individual-observer calibration procedure (see section 3.1.3) – we wanted to compensate for individual equiluminance differences in the same region of retinal eccentricity. The chromaticity of the adaptation stimulus was constructed from two sinusoidal modulations, each designed to cause modulation in one of the opponent mechanisms. One component modulated the L and M cones in counterphase while maintaining the total excitation, so as to modulate the L-M pathway. The other component modulated the S cones while maintaining L and M cone stimulation, so as to modulate only the S-(L+M) pathway. The amplitude of both modulations, and therefore the radius of the resulting circular locus, was 0.01 units in the MacLeod-Boynton chromaticity diagram scaled as discussed in section 1.2.2. The time-averaged stimulation of all three cones was metameric to equal-energy-white (EEW). All modulations had a frequency of 10Hz. In one experimental condition, the $\pm[S-(L+M)]$ modulation led the L-M modulation by a phase difference of $\pi/2$, resulting in a CW procession around a circular locus in colour space. In another condition this phase difference was reversed so that the S-L+M modulation lagged the $\pm[L-M]$ modulation by $\pi/2$, resulting in a CCW procession. These phase relationships were discussed earlier and illustrated in Figure 23. The Hue Circle stimuli were described in further detail in section 2.

5.2.1.2 Experimental procedure

Three adaptation conditions were used: (a) in which the adaptation stimulus was modulated in a CW sense, (b) in which the adaptation stimulus was modulated in a CCW sense, and (c) the baseline condition in which the adaptation stimulus had the same chromaticity as the background and was not modulated.

In all conditions, adaptation and probe stimuli were presented in the centre of the monitor, against an EEW background identical to the average of the adaptation modulations. The observer

was instructed to fixate a dot in the centre of the screen throughout the experiment. For each trial, in the two adaptation conditions, the adapting stimulus was presented for 4.0 seconds. The stimulus amplitude was ramped on during the leading 500 milliseconds and off during the trailing 500ms. A delay of 167ms followed the adaptation stimulus, then the probe stimulus was displayed for 167ms. The next trial began after 167ms, during which time the observer made a response using a button box. This temporal sequence is illustrated in Figure 25. Before the first trial, the adaptation stimulus was repeated for 2 minutes. In the baseline condition, the timing parameters were the same, but the adaptation stimulus had a chromatic contrast of zero and so was identical to the uniform background.

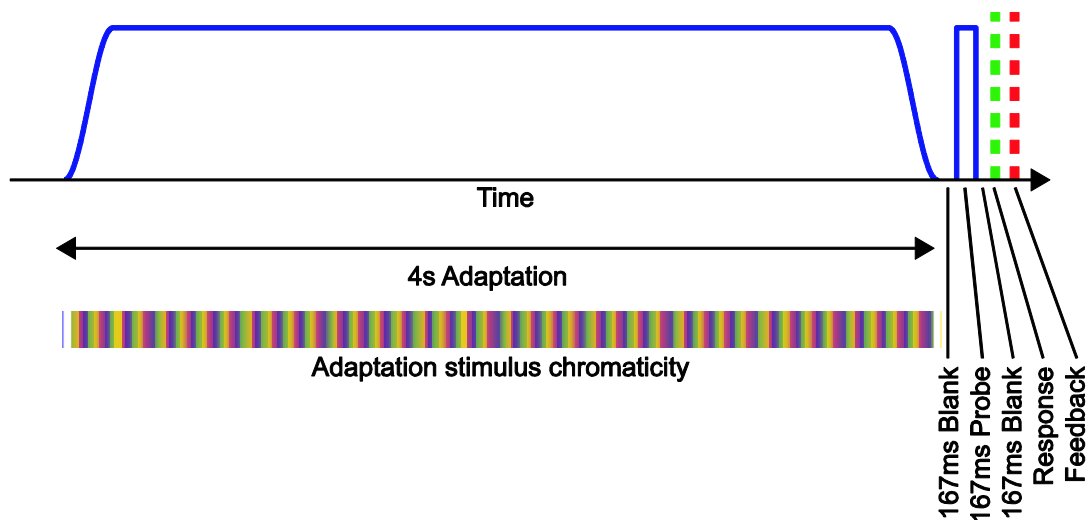


Figure 25. Temporal sequence of one trial.

The chromatic contrast of the probe stimulus was adjusted for each trial based on the responses to the previous trials in a staircase procedure. The contrast was decreased by 0.1 log unit after two consecutive correct answers and increased by the same factor after one incorrect response. The staircase terminated after eight reversals and the contrast of the final six were averaged to obtain a measure of discrimination contrast threshold. Eight staircases were randomly interleaved in each run, each with a different pair of hues so that sixteen hues, evenly distributed around EEW, were used in all. An additional baseline condition was run before the adaptation conditions and each observer's ratio of threshold measurements in the hue directions along the cardinal axes were used to scale the ratio of their adaptation stimuli modulation amplitudes. For all observers, the discrimination thresholds were higher along the $\pm[L-M]$ axis than along the $\pm[S+(L+M)]$ axis, so the L-M modulation was set to the maximum amplitude allowed by the monitor gamut and the $\pm[S-(L+M)]$ modulation scaled according to the observer's relative discrimination thresholds in this direction. On each trial, the initial phases of the adaptation stimulus modulations were offset by a random amount. The same random offset was applied to both component modulations, in addition to the $\pi/2$ or $-\pi/2$ phase difference. This is equivalent to starting the procession at a

random point around the hue circle, and was done to prevent preferential adaptation by the colour on which the sequence ended.

The three conditions were run in a counterbalanced order. Three observers participated, RJJ is the author, and KLB was naive to the purpose of the study. RJJ and HES are experienced colour psychophysical observers. Each observer repeated the three conditions at least eight times.

5.3 Results

We show thresholds in as a function of hue angle, on radial plots, in Figure 26 (observer HES), Figure 27 (observer KLB), and Figure 28 (observer RJJ). Panel A in each figure show thresholds in the baseline condition when no adaptation modulation was shown, panel B shows thresholds after adaptation to the CW stimulus, and panel C shows thresholds after adaptation to the CCW stimulus. The direction of each point from the centre represents the direction of the test chromaticity from our EEW background, when represented in the MacLeod-Boynton chromaticity diagram, and the points are coloured in approximations to the appropriate hues. The distance of a point from the centre represents the threshold for distinguishing this hue from the hue on the opposite side of the centre. Since the task required the observer to distinguish a hue from its opposite, thresholds for opposite pairs of hues are the same and the plots are rotationally symmetric. Each point is the geometric mean of all repeated threshold estimations, and the shaded areas show one geometric standard deviation above and below the mean. We define an ellipse and allow its size, eccentricity and rotation about the centre to vary in order to obtain a least-squares best-fit to the mean data.

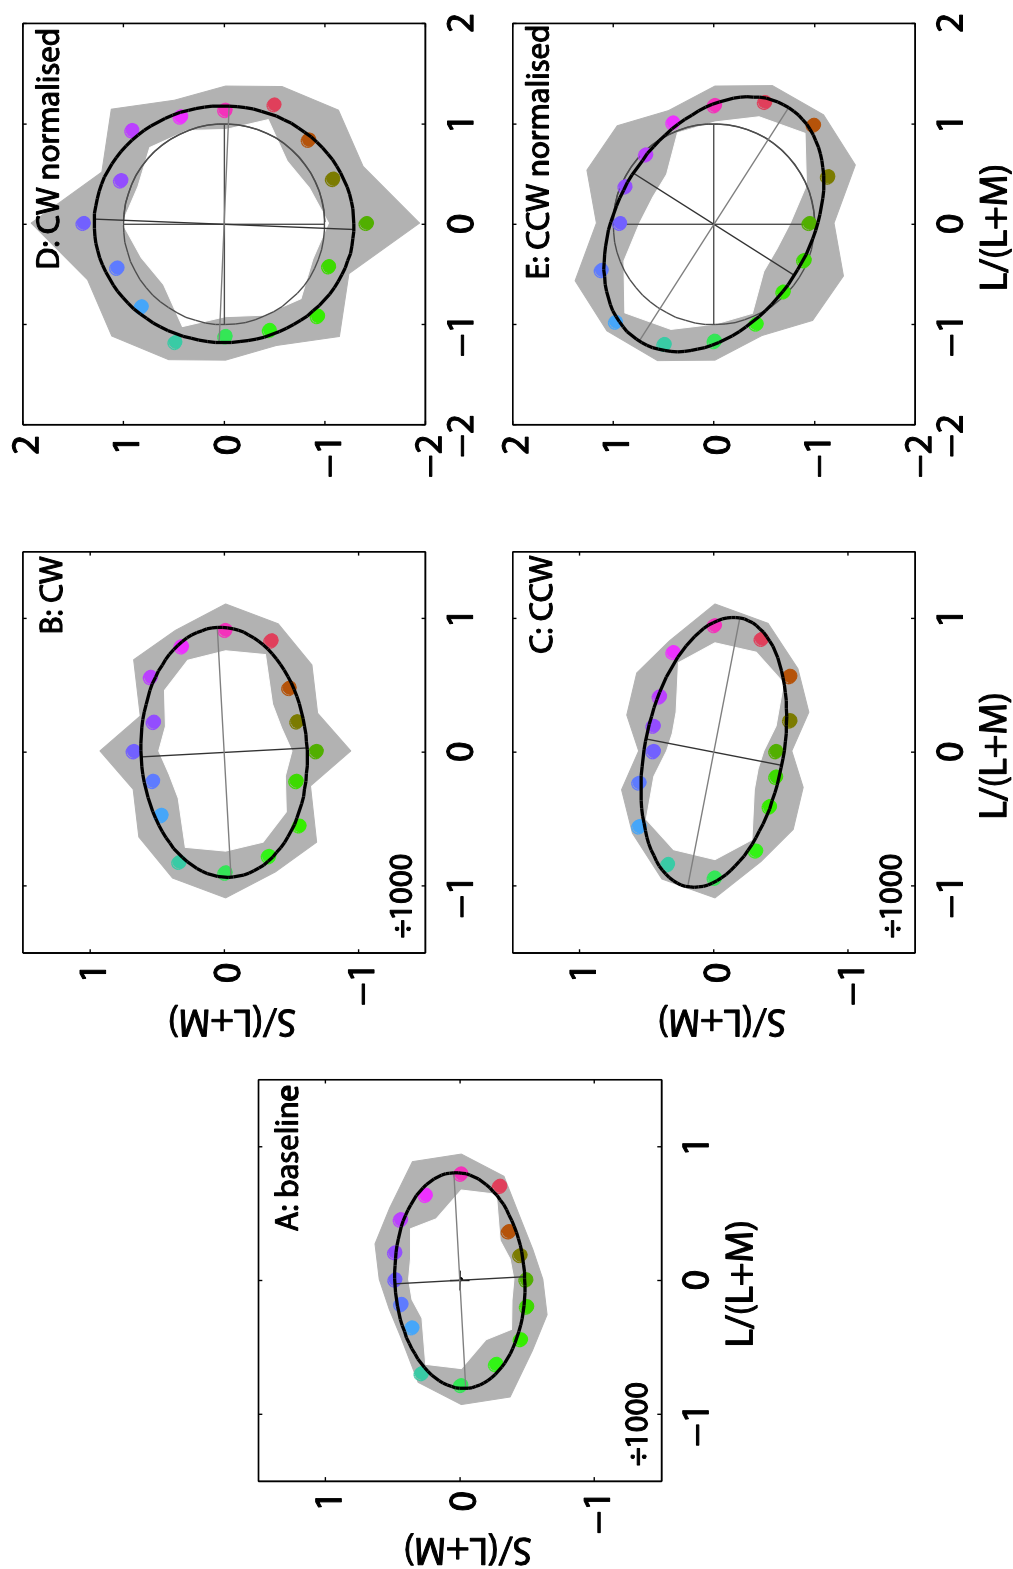


Figure 26. Thresholds for observer HES when discriminating between chromaticities opposite EEW in the equiluminant plane. The distance of each point from the centre of each plot represents the geometric mean of at least eight repeated measurements of threshold. The angular direction of each point from the centre of each plot represents the chromatic direction of the test chromaticities, illustrated by the colour of the point. The heavy black ellipses are fitted to the points with a least-squares procedure, and the filled grey regions illustrate one geometric standard deviation above and below the mean of each threshold. Panel A shows thresholds in the baseline condition after adaptation to unmodulated EEW. Panels B and C show thresholds after adaptation to the CW and CCW stimulus, respectively. Panels D and E show the CW and CCW thresholds from panels B and C normalised by the radii at the appropriate chromatic angle of the ellipse fitted to the baseline data in panel A.

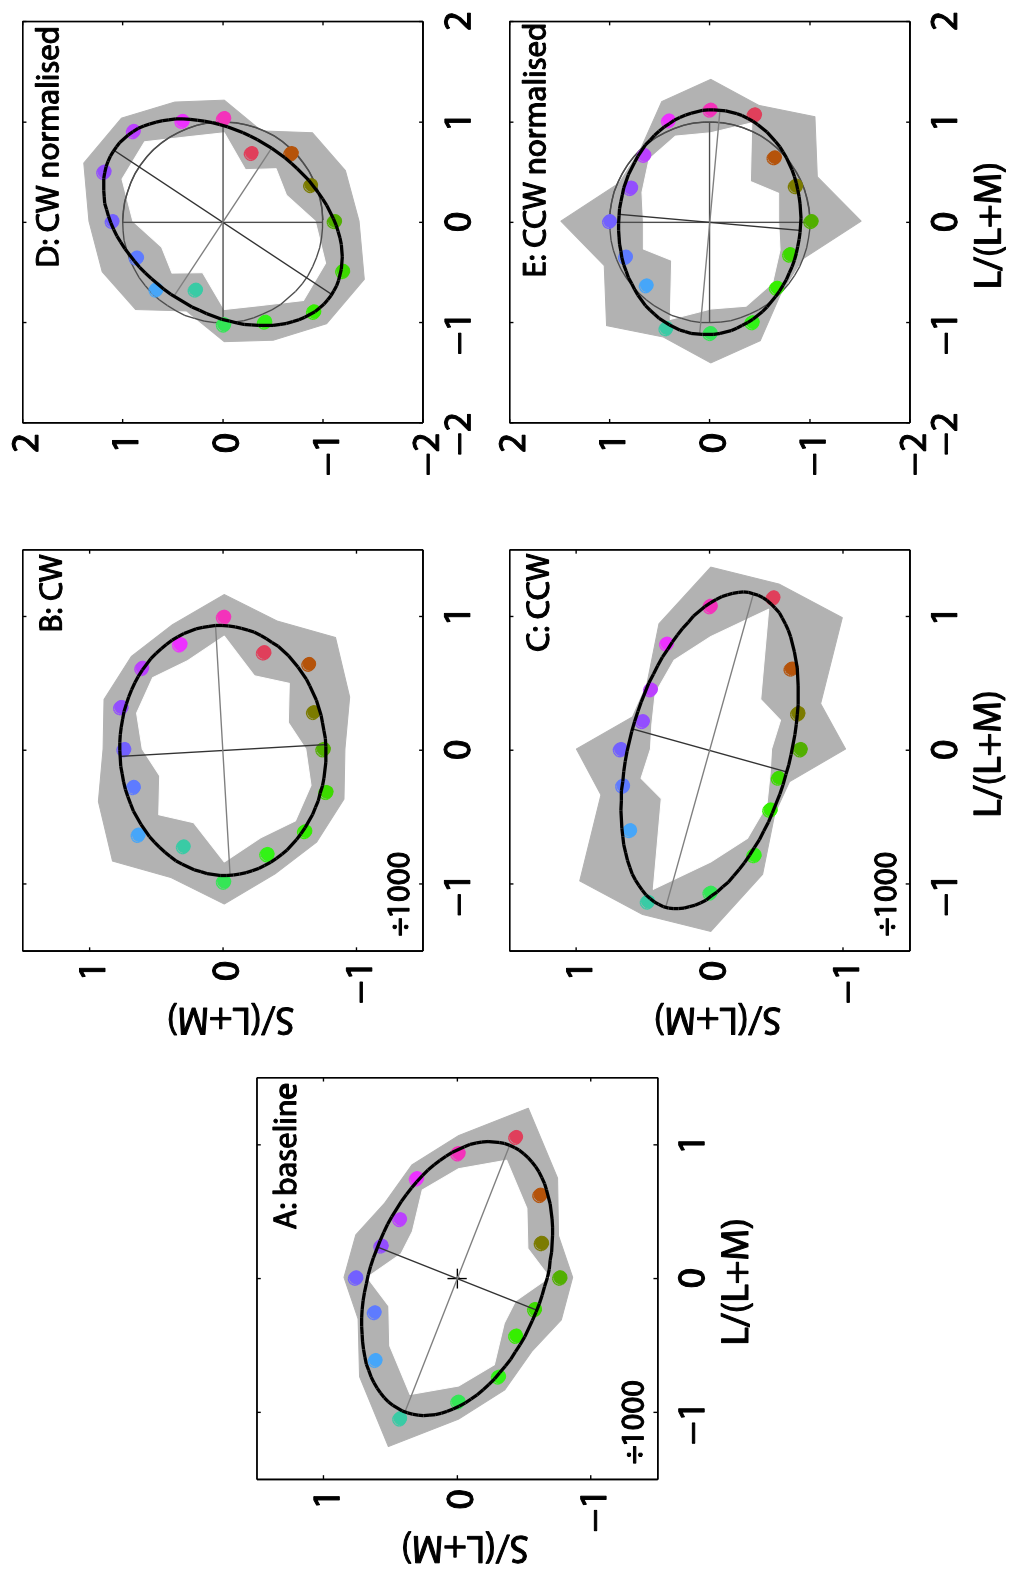


Figure 27. As Figure 26, for observer KLB.

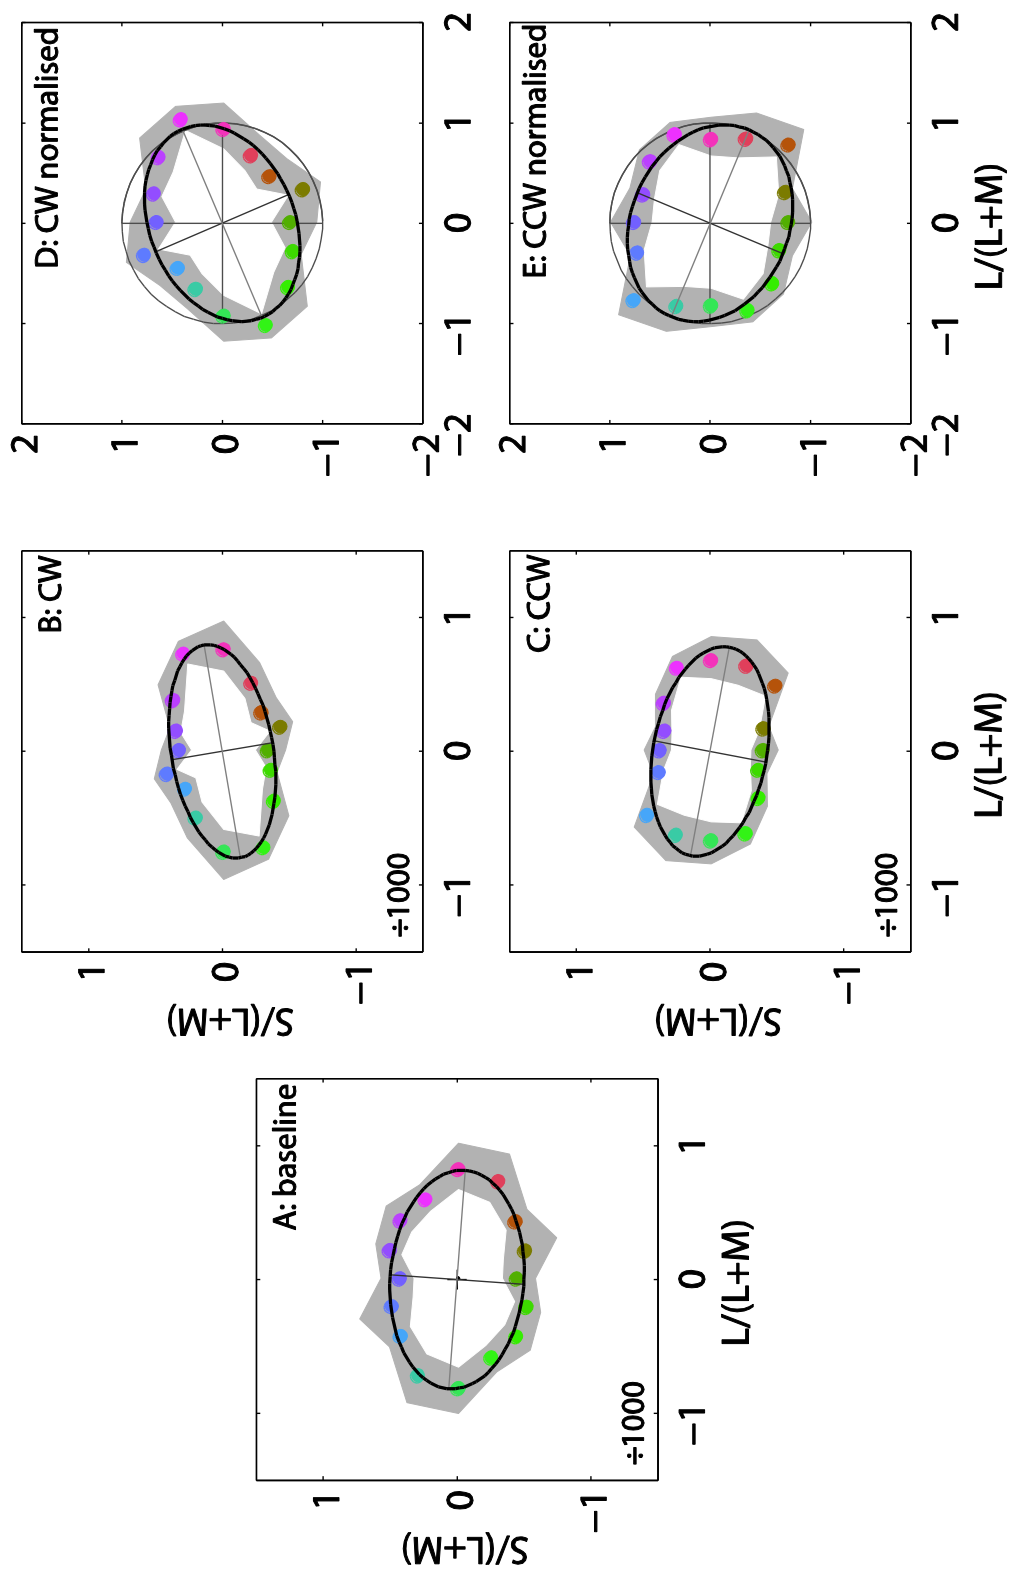


Figure 28. As Figure 26, for observer RJL.

To represent the magnitude of adaptation at different hue angles, we normalise the data in the CW and CCW adaptation condition by the curve fitted to the baseline data, in a similar way to how Krauskopf et al. (1982) scaled their adaptation stimuli. Each average adapted threshold is divided by the baseline threshold predicted by the distance from the centre of the curve fitted to the baseline data at the appropriate hue angle. Panels D and E in each figure show these normalised data. The distance of each point from the centre now represents the ratio of adapted to baseline threshold for discriminating the appropriate hue-pair.

The most important feature of these plots to note is that the orientation of the major axis of the ellipses in the adaptation plots is different after adaptation to the CW and CCW stimuli. As discussed above, both adaptation stimuli are designed to cause equivalent adaptation in the receptors and cardinal opponent mechanisms. The different patterns of threshold changes we see after adaptation to CW and CCW stimuli, must therefore be a result of adaptation of mechanisms with non-cardinal tuning.

5.4 Discussion

Our two adaptation stimuli are designed to cause different patterns chromatic discrimination thresholds. The plots of results in Figure 26, Figure 27 and Figure 28 would indicate that this is the case. The ellipses fitted to the discrimination thresholds of all three observers are oriented differently after adaptation to the CW and CCW stimulus. In general, the ellipses are oriented towards the direction predicted by modelling latency in the S-cone signal pathway and illustrated in Figure 23. The model predicts that the ellipses should be oriented with their major axis along a 45° diagonal line in MB-DKL colour space. This is not always the case in these data, but for example in the case of observer RJL, a simple linear scaling of the plots in the S-(L+M) direction will result in the major axes of the two ellipses being oriented at roughly $\pm 45^\circ$. Deviations from orientations of 45° are unsurprising, since there is no simple way of equating excursions in different chromatic directions which would create equivalent responses in the two putative cardinal mechanisms. The adaptation stimulus amplitudes were scaled based on discrimination thresholds, but this may not have the same effect for suprathreshold stimuli such as those used as the adaptation stimulus.

There are some features of the plots that remain unexplained. Firstly, in some plots, the ellipses fitted to normalised thresholds after adaptation are not oriented along mirror-symmetric directions after CW and CCW adaptation, contrary to our predictions. This is most noticeable for observer HES after adaptation the CW stimulus and for observer KLB after adaptation to the CCW stimulus. In these plots, the major axes of the ellipses do not appear significantly rotated from the cardinal directions, although the small amount of rotation is in the predicted direction. This could be due to noise in the measurements in combination with the fact that the increases in threshold

that we measure are very small (approximately 25%), meaning that the magnitude of threshold change after adaptation is, at certain critical points on around the ellipse, no greater than we might expect by chance. Also, any error in estimating the orientation of the ellipse fitted to the baseline thresholds will bias the rotation of the ellipses fitted to normalised thresholds. Again, the most important feature of the normalised threshold data is that the orientations are different after adaptation to CW and CCW stimuli.

Secondly, we would predict that the adaptation stimulus would cause an increase in thresholds which is maximal along the direction of the major axis of the resultant chromatic locus reaching the central site, but still nonzero in other directions. Even in chromatic directions in the direction of the minor axis of the resultant ellipse, there is some contrast from EEW. However, most notably in the plots of data from observer RJL, nearly all the ratios of adapted and baseline thresholds are below unity (indicated by the grey circle in panels D and E in Figure 28) which in fact suggests that thresholds were reduced after adaptation, maximally in chromatic directions aligned with the minor axis of the fitted ellipse. This could have a simple explanation: baseline thresholds may be overestimated, relative to adapted thresholds, due to temporal uncertainty of the onset of the probe stimulus in the baseline condition. In the adaptation conditions, the offset of the adaptation stimulus gave a cue to the onset of the probe stimulus, but in the baseline condition this cue was absent. The probe stimuli in the baseline conditions were preceded by unmodulated EEW for more than four seconds so observers could not anticipate the probe stimulus and may have been unprepared to make the discrimination.

Despite these unexplained features, the results do indicate that we have succeeded in producing different patterns of discrimination thresholds after adaptation to two stimuli which are, not only comprised of the same set of lights, but which are designed to provide equivalent stimulation to the cones and cardinal mechanisms, and so to isolate the stage at which signals in these mechanisms are combined (with a relative delay).

Since the elliptical patterns of threshold are not aligned with the cardinal axes of the MacLeod-Boynton chromaticity diagram, we conclude that the adaptation we see here is not solely due to the preferential adaptation of the early opponent mechanisms that define this space, and demonstrated by Krauskopf et al. (1982). The simplest interpretation of these data might be that we have differently adapted two chromatic mechanisms that are tuned to chromatic directions other than the cardinal axes. One of these mechanisms would be tuned to chromaticities closer to the major axis of the ellipses fitted to our data showing thresholds after, say CW adaptation, while the other would be tuned closer to the minor axis. These relationships would be reversed after CCW modulation. From these data, it is not possible to infer much about the tuning direction or tuning selectivity of these mechanisms, only that they are tuned to directions other than the

cardinal directions. It is possible that they are not tuned to orthogonal chromatic directions. An alternative interpretation is that the stage at which this adaptation occurs contains multiple chromatically tuned mechanisms, each with a different tuning direction and each more narrowly tuned than the early opponent mechanisms (e.g. Lennie, et al., 1990). Those mechanisms tuned to directions closer to the long axes of our fitted ellipses, in either the CW or CCW condition, are adapted more than those tuned to other directions because the resulting contrast in their chromatic direction (to which they are maximally sensitive) is greater. By the same argument, those mechanisms tuned closer to the minor axes of the ellipse are adapted the least. In either case, assuming that chromatic information is carried at some stage between the receptors and this central site in only two independent orthogonal mechanisms, signals in these mechanisms must be combined in some fashion in order to generate the patterns of selective adaptation that we see here.

6 Experiment 3: Adaptation of central and peripheral mechanisms measured with reaction times

6.1 Introduction

This experiment follows on from the previous one, in which CW and CCW modulations at 10Hz around a hue circle were used to adapt central colour mechanisms, reducing sensitivity to chromatic directions aligned intermediate to the cardinal axes of the MacLeod-Boynton chromaticity diagram. In this experiment, we aim to show that the same adaptation can also be measured, and perhaps more quickly, by measuring reaction times (RTs). A number of studies have measured the effects of chromatic stimuli on RTs (McKeefry, et al., 2003; Nissen & Pokorny, 1977; Smithson & Mollon, 2004) and these include studies in which RTs are affected by adaptation stimuli (Murray, Parry, & McKeefry, 2002; Parry, Murray, & McKeefry, 2008). Also, RTs have been shown to be a reliable way of measuring sensitivity (Murray & Plainis, 2003; Plainis & Murray, 2000). Reaction times increase with less-detectable stimuli, so are monotonically related to threshold. We decided to use the reaction time method, since we anticipated that fewer trials would be needed compared to a procedure measuring threshold with a staircase, as in the previous experiment. Also, as suggested by Parry et al. (2008), reaction time measurements provide more 'real-time' information since many trials are not required to make a single estimate, as in a threshold measuring procedure, and can therefore be accurate even if the time-course of the adaptation effect is short.

The results of the previous adaptation experiments show that the magnitude of adaptation is greatest in the chromatic direction along the major axis of the elliptical locus reaching the central site, where it is assumed that signals in the L/M- and S-opponent cardinal mechanisms are combined. This direction lies along the positive or negative diagonal of colour space, intermediate to the cardinal directions, depending on the direction of modulation. CW modulations result in higher thresholds along the positive diagonal ($\pm[+S+L-M]$), while CCW thresholds result in higher thresholds along the negative diagonal ($\pm[-S+L-M]$). The magnitude of adaptation is minimal along the minor axis of the ellipse, which lies along the opposite diagonal to the major axis. We measure reaction times to discriminating colours lying in these two directions only, predicting that RTs will be higher when discriminating chromaticities along the positive diagonal than the negative diagonal after adaptation to the CW stimulus, and vice-versa after CCW adaptation. We measure reaction times for chromaticities differing in the cardinal directions in a separate set of sessions and predict that the magnitude of adaptation in directions along the cardinal directions should not differ. The predicted loci of chromaticities reaching the central site cross at positions along the cardinal directions from the centre of the chromatic modulations and the thresholds were similar at these chromatic directions (see Figure 23).

However, adaptation to stimuli consisting of just one of the component modulations used to create the hue-circle, i.e. a sinusoidal modulation in time along one of the cardinal directions of colour space, should increase thresholds in that direction only (Krauskopf, et al., 1982). A modulation along the L/(L+M) direction should increase thresholds for discriminating chromaticities lying along this direction while leaving discrimination for chromaticities differing only in their S/(L+M) coordinate unaffected. Similarly, a S/(L+M) modulation should increase S/(L+M) thresholds while not affecting L/(L+M) thresholds. As before, increases in discrimination threshold are predicted to lead to a corresponding increase in RT. Krauskopf et al. (1982) used 1Hz modulations to demonstrate adaptation by measuring thresholds, and they concluded the adaptation was taking place in post-receptoral mechanisms. Similar results were obtained by Parry et al. (2008) who measured RTs. In this experiment, we attempt to replicate the results of these studies with 3Hz adaptation modulations, measuring RTs. We used 3Hz modulations, rather than 1Hz as Krauskopf et al. and Parry et al. because we wanted to keep the same time-course as the previous experiment, specifically with 4s top-up adaptation stimuli with 0.5s ramped offsets. With 1Hz modulations, this would mean that only 4 cycles of the stimulus would be shown and the adaptive effect could be stronger to one pole of the chromatic direction than the other if the adaptation was ramped off at a peak or trough of oscillation. We also investigate whether the same affect can be observed after adaptation to 10Hz stimuli.

We predict that hue circle adaptation stimuli constructed from 3Hz modulations are unlikely to result in differently increased thresholds along the diagonal directions in colour space, as we expect them to with the 10Hz stimuli. The differently oriented elliptical loci predicted to reach the central site result because of an additional phase difference between the responses to component modulations, and we assume that this is due to different latencies in the signals in the two cardinal opponent mechanisms. We assume that these latencies are fixed temporal delays, not phase differences – the S-opponent signal is offset from the L/M opponent one by a fixed amount of time, not a fixed phase. At 10Hz, this temporal delay is enough to result in a phase difference that can cause appreciably different chromatic loci when the two modulations are combined at the central site. The magnitude of this delay, however, is not great enough to cause a phase shift in 3Hz modulations that would result in elliptical loci with significantly different eccentricities. The phase shift (φ) is related to the temporal latency (l) and modulation frequency (f) by the following formula:

$$\varphi = 2\pi fl$$

Differences in eccentricities resulting from different phase differences are illustrated in Figure 5. Assuming a relative delay l of 10ms (the delay measured in the previous experiment is of this

order), then 10Hz modulations produce resultant ellipses a ratio of maximum and minimum excursion of 1.96. For 3Hz modulations, this ratio is 1.21.

6.2 Methods

The stimuli used in this experiment were very similar to those used in the previous experiment (see section 5). The apparatus were identical, the calibration was the same, pixel dimensions were 800×600 and the frame rate was still 120Hz. The observer viewed the monitor from a distance of one meter as before giving a horizontal viewing angle of 22°. As with the previous experiment, on each trial an adaptation stimulus preceded a probe stimulus. There were five adaptation stimulus conditions, which will be described below, but each consisted of either a temporally modulated chromatic modulation or adaptation to EEW in a baseline condition. Where before all the adaptation stimuli were modulated at 10Hz, in this experiment we additionally used 3Hz modulations in separate conditions.

The surround was metameric to an equal-energy-white (EEW) spectrum as before, and had the same luminance. On each trial, the adaptation stimulus appeared for 4000 milliseconds. The first and last 500ms were ramped on and off with a raised cosinusoidal envelope. An interval of 20 frames (167ms, at the framework of 120Hz) followed in which the adaptation modulation amplitude was zero, then this was followed by the probe stimulus for 12 frames (100ms). The next trial began 1000ms after the onset of the probe stimulus, and the observer has that period of one second to make a response (see below). The software used to control the stimulus presentation in this experiment was adapted from that written for the previous experiment, and stimulus durations were specified in frames. Conversion to nominal stimulus durations is achieved using the framerate. The adaptation sequence was repeated 30 times before the first trial began.

Where before the contrast of the probe was adjusted each trial by the staircase, in this experiment it was fixed. The chromaticities in the probe stimuli had a radial distance from EEW of 0.002 units in the MacLeod-Boynton chromaticity diagram, scaled as discussed in section 1.2.2. Luminance offsets were added to each quadrant as before. We did not use the whole set of eight bipolar directions in colour space that were used previously. In separate sessions we used test directions oriented either along the cardinal directions of colour space or along directions oriented 45° to the cardinal ones, leading to two bipolar directions being tested in each session. Each of the two chromatic directions was tested on forty trials in each session, giving eighty trials in total. In half of the trials for each direction the probe stimuli split horizontally and in the other half it was split vertically. The presentation order of the trials was randomised.

As before, the observer's task was to indicate whether the probe stimulus was split horizontally or vertically, and to respond with a button press. However, this time they were instructed to make

their response as quickly as possible. They were instructed not to sacrifice accuracy for speed, but to make the response as soon as they were able to make the discrimination. Audio feedback was given as before. The response made, and the time between the onset of the probe stimulus and the time at which the observer pressed a response button was measured and recorded. The VSG system allows this reaction time to be measured with sub-millisecond accuracy. If the observer did not press a button in the one second before the next stimuli began, this was classified as an incorrect response.

Since reaction times were being measured in this experiment, it was important to maintain the characteristics of the stimuli across the adaptation conditions. Also, while in the adaptation conditions the offset of the adaptation stimulus allowed the observer to anticipate the probe stimulus, this was not the case in the baseline condition. For this reason, we added a thin (5 pixels) black ring around the adaptation stimulus whose contrast with the surround was ramped simultaneously with the adaptation stimulus modulation amplitude. In the baseline condition, the ring had the same temporal characteristics, but simply appeared on an un-modulated adaptation field in the same temporal position as in the adaptation conditions. The offset of this ring allowed the observer to predict the onset of the probe stimulus in the same way as in the adaptation conditions.

The four adaptation conditions differed in the amplitude and phase of the sinusoidal modulations along the cardinal directions of the MacLeod-Boynton chromaticity diagram. In the first condition, the stimulus was a clockwise CW procession around a circle centred on EEW, as in the previous experiment. The second condition was a counterclockwise CCW procession, again, as before. In the third condition, the amplitude of the modulation in the S-opponent direction was set to zero, resulting in a sinusoidal modulation along the L/M-opponent direction only. In the fourth condition, the amplitude of the modulation in the L/M-opponent direction was set to zero, resulting in a sinusoidal modulation along the S-opponent direction only.

This results in twenty unique conditions: four adaptation conditions and the baseline condition x two test direction conditions x two adaptation frequency conditions. Both observers completed each session twice, so completed forty sessions in total, with a possible 80 repeated trials contributing to each RT estimate (assuming all responses were correct).

6.3 Results

6.3.1 Analysis of response times

Before testing hypotheses about the data, reaction times that were to outside a given range were removed. Reaction times greater than 1.5 times the interquartile range away from the central quartiles were considered outliers (Cardinal & Aitken, 2006). More accurately, a reaction time x

was considered an outlier if $x > Q_3 + 1.5I$ or $x < Q_1 - 1.5I$, where Q_1 and Q_3 are the first and third quartiles, respectively, and I is the interquartile range, $Q_3 - Q_1$. This is a standard method of classifying outliers.

This was intended to disregard fast responses likely to be anticipations and long reaction times likely due to occasional cases of loss of attention by the observer, or button-press errors. This is standard practice in dealing with reaction time data. This procedure was performed individually on the measured reaction times from each adaptation and test direction condition combination. This procedure removes only few reaction times, as in some cases the range of reaction times derived from the distributions (as above) included those with negative value, and reaction times longer than the 1000ms second response period.

The RTs of incorrect responses were also removed from the analysis. This, as well as the removal of outlying RTs, means that the results of each of the different conditions may be based on a different number of reaction times. There were differences in the number of errors made in different conditions, as can be seen below.

The means and standard deviations of the response times measured in each adaptation condition, as well as the baseline condition in which the adaptation stimulus was unmodulated EEW, are shown in Table 3.

Table 3. Mean and standard deviations of response times, in milliseconds, for each observer in each adaptation condition.

	Test direction	Adaptation direction	baseline	10Hz Adaptation				3Hz Adaptation			
				L/M	S	CW	CCW	L/M	S	CW	CCW
RJL	±[+L-M]	mean	401	398	370	440	403	432	418	401	406
		std. dev.	79	70	69	100	68	94	83	71	91
	±[S-(L+M)]	mean	391	388	387	397	413	400	420	413	389
		std. dev.	71	72	84	69	90	83	93	86	84
	±[+S+L-M]	mean	407	399	397	375	456	408	413	393	441
		std. dev.	94	87	97	73	123	78	106	103	112
	±[-S+L-M]	mean	425	418	398	417	418	428	452	404	456
		std. dev.	109	88	74	91	63	118	138	87	104
HES	±[+L-M]	mean	515	580	468	552	530	479	451	458	443
		std. dev.	118	147	90	118	119	100	87	80	84
	±[S-(L+M)]	mean	512	525	565	549	526	455	599	501	492
		std. dev.	105	127	151	118	110	103	138	91	104
	±[+S+L-M]	mean	533	504	500	554	553	513	523	552	578
		std. dev.	123	85	108	123	142	100	142	125	130
	±[-S+L-M]	mean	520	495	512	535	502	507	537	545	561
		std. dev.	98	105	102	117	96	108	111	105	125

We wished to test hypotheses about the differences of the means of the reaction times in different conditions using an analysis of variance (ANOVA) procedure. A requirement of this sort of parametric tests is that the values be distributed normally (approximating a Gaussian distribution). The responses for each of the test directions for each of the adaptation conditions were subjected to a Shapiro-Wilk test to determine whether they were normally distributed. Many of the distributions did significantly differ from the normal distribution. To resolve this, the inverse of each reaction time was taken. This is a common transformation in analyses of RT data, since the inverse of RT is equivalent to reaction speed (RS). 10 (16%) out of the 64 RS distributions were significantly different from a normal distribution. Since the ANOVA test is robust to deviations from normal distributions (Field, 2005), this was deemed acceptable. All further analysis was carried out on these transformed, reaction speed, data.

Figure 29 and Figure 30 show, for observer RJL, plots of mean RS against adaptation condition with 10Hz and 3Hz adaptations stimuli, respectively. Figure 31 and Figure 32 show the same for observer HES. The two lines on each plot correspond to the two test directions. Red lines indicate the test stimuli had chromaticities along the L/M-opponent cardinal colour space direction, blue lines indicate test stimuli along the S-opponent direction. Purple lines indicate test directions along the positive diagonal, while orange lines indicate the negative diagonal. Error bars on each plot show one standard error of the mean, above and below the mean.

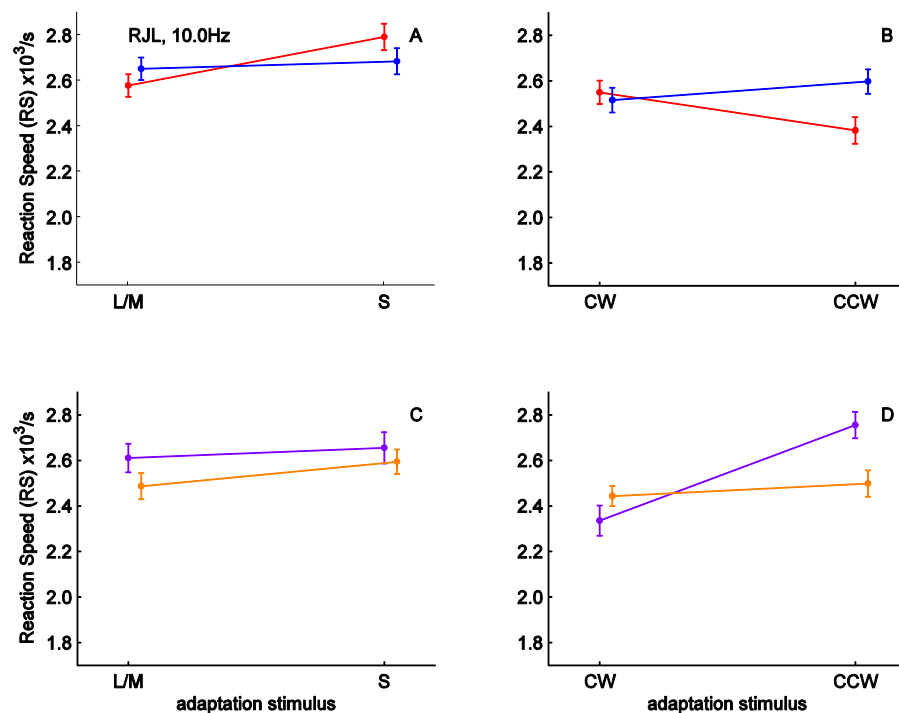


Figure 29. Mean reaction speeds (RS) measured from observer RJL when test stimulus chromaticities varied along the cardinal directions of equiluminant colour space (panels A and B, red lines: $\pm[+L-M]$ direction, blue lines: $\pm[+S-(L+M)]$ direction) and when they varied along intermediate, diagonal directions (panels C and D, purple lines: positive diagonal ($\pm[+S+L-M]$) direction, orange lines: negative diagonal ($\pm[-S+L-M]$) direction). Panels A and C show data from conditions in which adaptation stimuli were

modulations along either of the cardinal directions, and panels B and D show data from conditions when the adaptation stimuli were combinations of modulations in both cardinal directions resulting in CW or CCW processions around a circular chromatic locus. All adaptation stimuli were modulated at 10Hz. Error bars show one standard error above and below the mean ($\pm\sigma/\sqrt{n}$). The horizontal displacement of the blue and orange lines is simply for clarity.

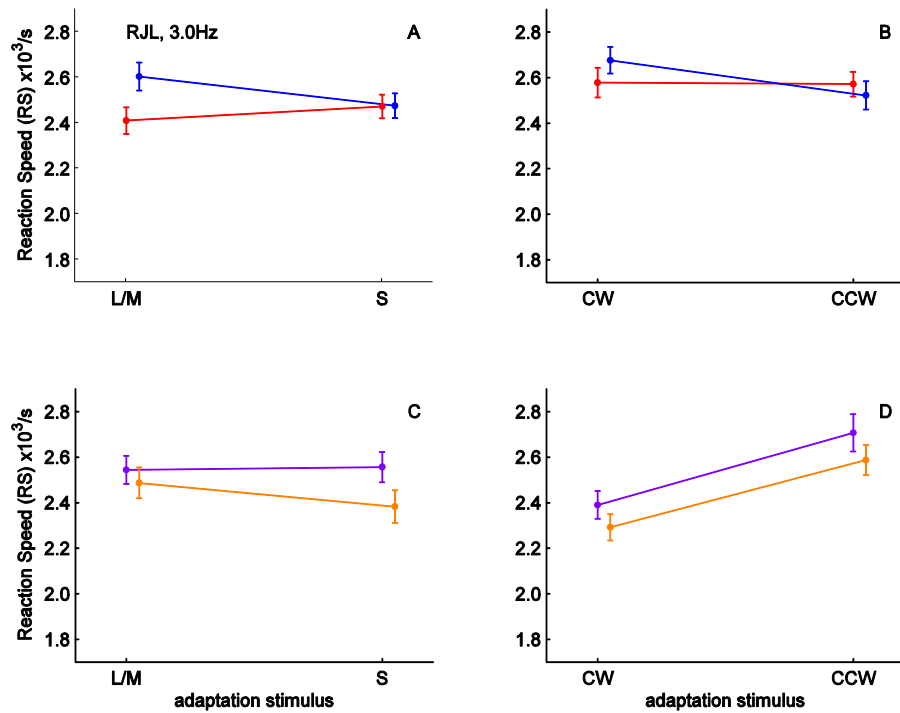


Figure 30. As Figure 29, with 3Hz adaptation modulations.

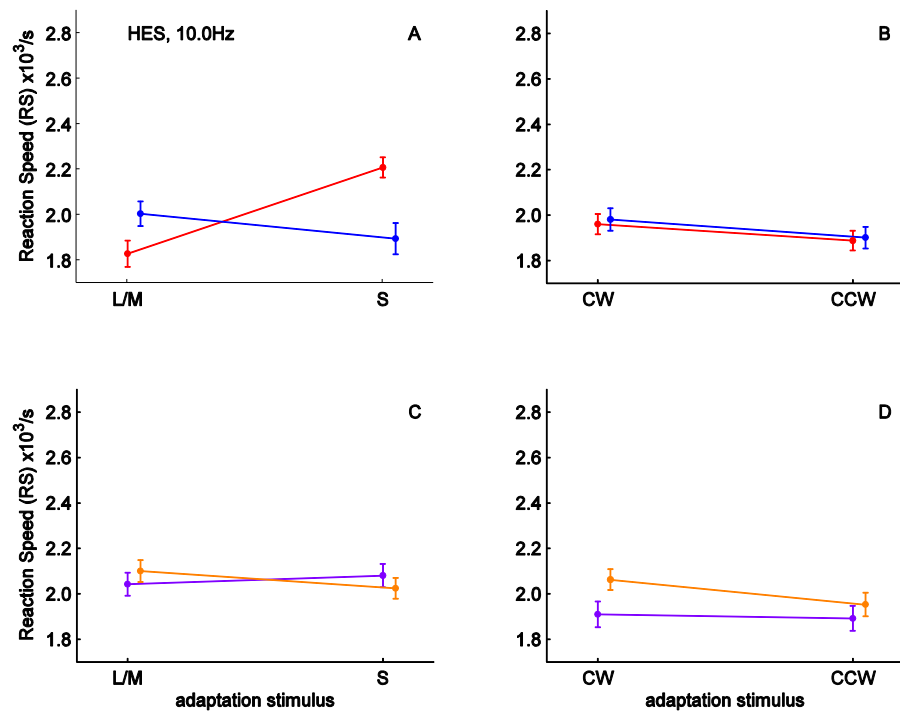


Figure 31. As Figure 29, for observer HES.

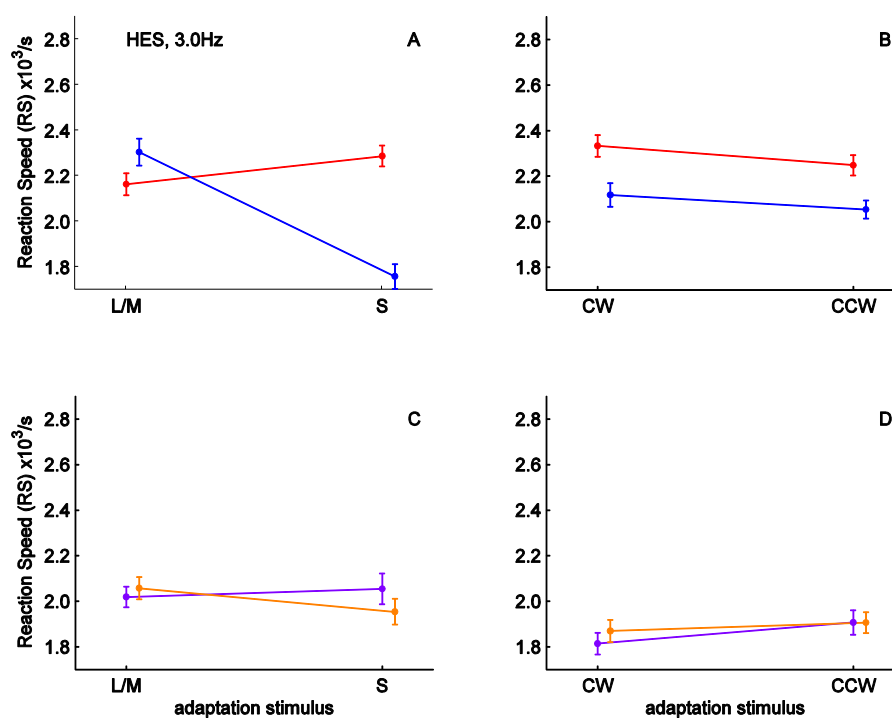


Figure 32. As Figure 31, with 3Hz adaptation modulations.

As discussed previously, we make two predictions about differences in reaction times, and further discussion will take place with regard to those predictions:

1. We predict that after adaptation to 3Hz modulations along either of the cardinal directions, thresholds (and hence RTs) for discriminations of chromaticities in the same direction will be increased, meaning RSs will be decreased. So, after L/M-opponent modulation adaptation, RSs for discriminating L/M test stimuli should be lower than after adaptation to S-opponent stimuli, and vice versa for S-opponent test stimuli.
2. We predict that RSs for discriminating test stimuli along the positive diagonal should be lower after adaptation to the 10Hz CW stimulus than after adaptation to the CCW stimulus, and vice-versa for the RSs for discriminating test stimuli along the negative diagonal.

Panels A in Figure 30 and Figure 32 indicate that, for both observers, prediction 1 is fulfilled. The mean reaction speeds for L/M test stimuli are slower after L/M adaptation than after S adaptation, and mean reaction speeds for S test stimuli are slower after adapting to S modulation adaptation. Similar patterns can be seen in panels A in Figure 29 and Figure 31, indicating that this effect is still obtained for 10Hz adaptation modulations.

Panel D in Figure 29 indicates that, for observer RJL, prediction 2 is fulfilled. The mean RS for positive diagonal test directions is lower after adaptation to the CW stimulus than the CCW stimulus, and although the mean RS for negative diagonal test directions is still lower after adaptation to the CW stimulus than the CCW stimulus, it is higher than for positive diagonal

directions after CW adaptation and lower than for positive diagonals after CCW adaptation. The same is true for observer HES (Figure 31 panel D), however the differences are much smaller. Also consistent with the theory is that the interaction effect is not present when 3Hz adaptation stimuli are used (Panels D in Figure 30 and Figure 32).

Any other differences in means that are contingent on both the direction of the test stimuli and on the adaptation stimulus, which might be indicated on other plots, appear small. This is consistent with our predictions that suggest there should be no interaction effects on the condition combinations described.

We look for interaction effects in these results, and so do not draw conclusions on any main effects. This is because we cannot assume that, regardless of adaptation, chromatic sensitivities are the same for each pair of directions we test. Differences in these sensitivities would lead to differences in reaction time between test directions, regardless of the adaptation stimulus. Also, while we assume that since the CW and CCW stimuli are equivalent in their adaptive strength on the opponent mechanisms, we cannot assume that the L/M- and S-opponent adapting stimuli have equivalent effects on their respective mechanisms, and this may also lead to main effects.

For each of the combinations of test direction conditions (cardinal and intermediate axes) and adaptation condition (L/M- and S-opponent modulations vs. CW vs. CCW modulations), a two-way ANOVA was performed to determine whether there was any effect of the test direction (within the possible directions, given that either cardinal or intermediate directions were being tested), modulation direction (within the possible directions, given that either cardinal or intermediate directions were being tested) or more importantly an interaction of the two.

A significant interaction effect was found for observer HES when discriminations for cardinal direction chromaticities were being tested after adaptation to 3Hz cardinal axis modulations ($F_{(1,259)}=40.28$, $p<0.001$), and also with 10Hz modulations ($F_{(1,250)}=19.08$, $p<0.001$). This is in full support of prediction 1. These interaction effects were not, however, significant for observer RJL, despite the differences in means being in the predicted directions (3Hz: $F_{(1,271)}=2.8$, $p=0.0985$. 10Hz: $F_{(1,279)}=2.83$, $p=0.094$).

A significant interaction effect, was found for observer RJL when diagonal test directions and 10Hz CW and CCW adaptation directions were used ($F_{(1,281)}=9.98$, $p<0.01$). This, along with the observation above that these mean RSs show differences in the correct directions, is in support of prediction 2. However, the corresponding interaction was not found for observer HES ($F_{(1,237)}=0.78$, $p=0.377$).

There are several other significant effects revealed by the ANOVAs. Several main effects of the test stimulus direction were found, but these do not affect prediction since we know that thresholds for discriminating chromaticities in different directions of colour space are different under neutral adaptation conditions (see section 5.3). Also, scaling test stimuli as proportions of threshold units (as was done here) may not be appropriate for RT experiments (Zele, Cao, & Pokorny, 2007). It is, however, difficult to explain why these differences were found only in some experimental conditions. There were also some significant main effects of the adaptation stimulus condition, indicating that RSs were different after L/M- vs. S-opponent adaptation, or CW vs. CCW adaptation, not contingent on the test chromaticity direction. These differences could be explained by the fact that it was this variable that was changed between sessions, so is likely to be affected by alertness states. Also, for the case of the differences of L/M- vs. S-opponent adaptation, another explanation could be that the amplitudes of the adapting stimulus modulations were not appropriately scaled. The ratio of the amplitudes was set by the unadapted thresholds (section 5.2.1.2), but this procedure may not equate the suprathreshold adaptation stimuli. The only significant interaction effects found were the ones already mentioned, which are explained by predictions 1 and 2.

6.3.2 Analysis of error rates

In addition to the response times/speeds, the numbers of incorrect responses made in different conditions were investigated. The number of incorrect responses made, which includes the occasions on which the observer did not make a response within the allowed time interval, are plotted on Figure 33 to Figure 36. The notation used is the same as in the previous figures.

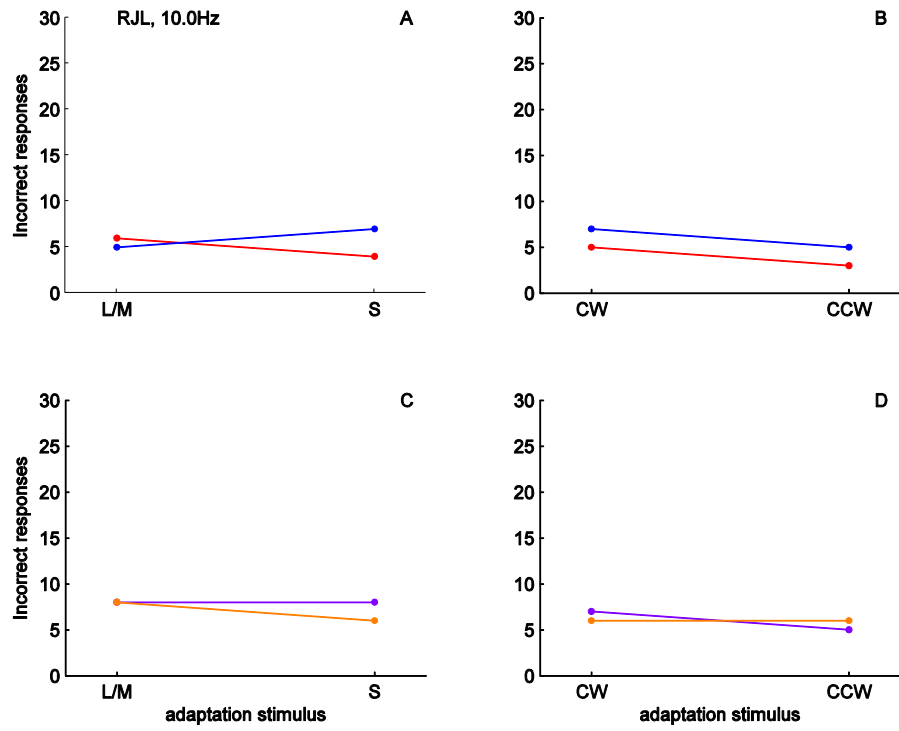


Figure 33. Incorrect responses made by observer RJL, when 10Hz adaptation stimuli were used. The colours of the lines represent the same as in Figure 29.

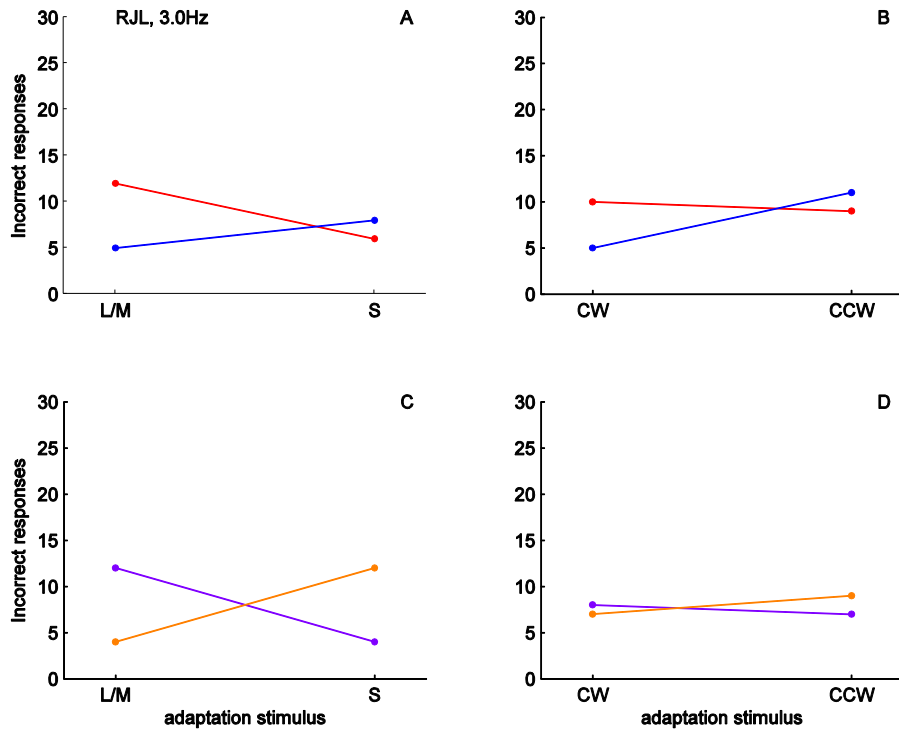


Figure 34. As Figure 33, with 3Hz adaptation modulations.

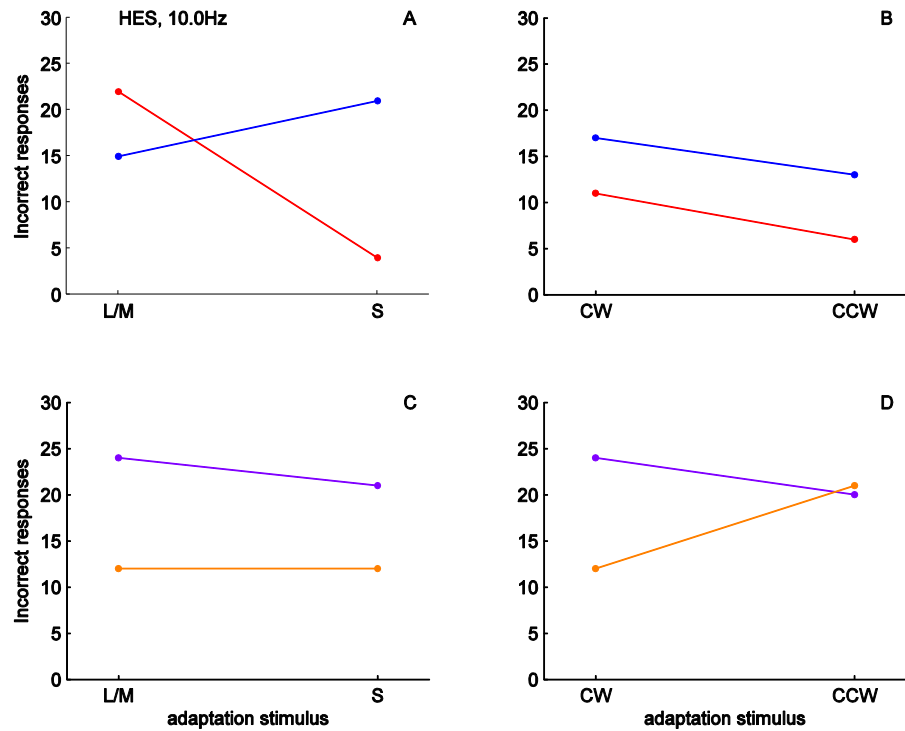


Figure 35. As Figure 33, for observer HES.

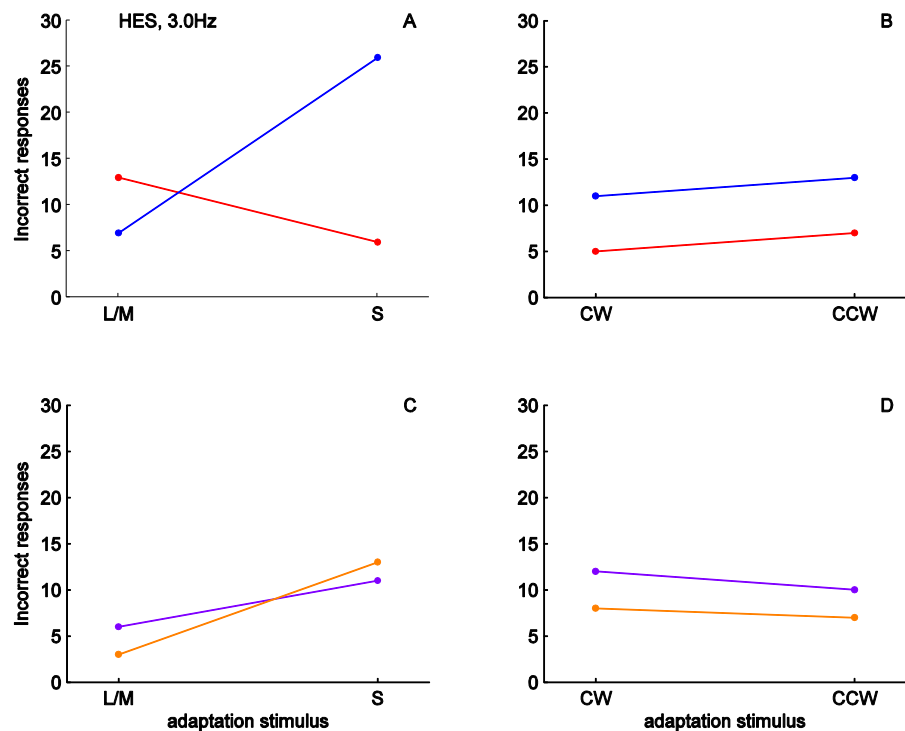


Figure 36. As Figure 35, with 3Hz adaptation modulations.

Since the number of errors should correlate with discrimination threshold, and therefore with RT and inversely with RS, we can extend predictions 1 and 2 above to error rates:

1. We predict that after adaptation to 3Hz modulations along either of the cardinal directions, thresholds for discrimination of chromaticities in the same direction will be increased, so more incorrect responses will be made. So, after L/M-opponent modulation

adaptation, more errors should be made when discriminating L/M test stimuli than after adaptation to S-opponent stimuli, and vice versa for S-opponent test stimuli.

2. We predict that more errors should be made when discriminating stimuli along the positive diagonal after adaptation to the 10Hz CW stimulus than after adaptation to the CCW stimulus, and vice-versa for the test stimuli along the negative diagonal.

Panels A in Figure 34 and Figure 36 are consistent with prediction 1: more errors are made when discriminating test stimuli with chromaticities in the same direction of colour space as the adaptation stimulus modulation, and this is true for both observers. The same differences, although smaller, are shown for 10Hz, modulations in Figure 33 and Figure 35, which lends more support to the prediction, and indicates that the improvements in response speed seen above were not a result of a speed-accuracy trade-off.

Panels D in Figure 33 and Figure 35 are consistent with prediction 2: more errors are made when discriminating test stimuli with chromaticities on the positive diagonal of colour space after adaptation to the CW stimulus than after adaptation to the CCW stimulus. As when analysing RSs, this interaction disappears when 3Hz adaptation stimuli are used with observer HES, (Figure 36 panel D) and with RJL it is small (Figure 34 panel D).

To test whether the differences in error rates were statistically significant, χ^2 contingency tests were carried out on the frequencies of incorrect and correct responses to determine whether this was dependent on the test stimuli chromatic direction within a given combination of test and adaptation conditions, with the null hypothesis that the number of errors made in a given condition combination was equally distributed between the two test directions. Several of these tests indicated significant contingencies, however there is little correspondence between these significances and the predictions we have made, perhaps because the absolute numbers of errors we are dealing with are small. The only significant differences that support the predictions are the following, all for observer HES: more errors were made to S-opponent test stimuli than L/M stimuli after adaptation to 3Hz and 10Hz S-opponent modulation (3Hz: $\chi^2=20.83$, $p<0.01$, 10Hz: $\chi^2=16.81$, $p<0.01$) and more errors were made to positive diagonal test stimuli than negative diagonal stimuli after adaptation to 10Hz CW modulation ($\chi^2=7.27$, $p<0.01$). This last result seems inconsistent with the lack of significant differences between RSs in the same condition for this observer, but we can conclude that the effect of adaptation was revealed primarily as an accuracy deficit, rather than a speed deficit. Against our predictions are the significant differences between the errors made by RJL to positive and negative diagonal stimuli, which are contingent on adaptation either to 3Hz L/M- and S-opponent modulations, and similar differences for HES after the corresponding 10Hz modulations.

6.4 Discussion

The results of this experiment provide tentative support for the predictions. Disappointingly, not all the differences in reaction time between adaptation/test stimulus conditions that we observed, and that were in line with those we predicted, proved statistically significant. Furthermore, there were some significant differences in RT that we did not predict, and run contrary to our theory. All the differences in reaction time that we observe are quite small. It might be possible to obtain less noisy results with more measurements, or by averaging the result of several observers.

The plots in the previous section show that reaction times for discriminating chromaticities differing along either of the cardinal directions of colour space are longer after adaptation to 3Hz - and perhaps to a lesser extent 10Hz - modulations in the same direction, compared to adaptation in the orthogonal direction. This is consistent with results of Krauskopf et al. (1982) and Parry et al. (2008). These previous studies used 1Hz adapting modulations, while ours were 3Hz. We notice that the colours of the adapting stimulus appear less saturated at higher frequencies than at lower ones, which might be expected since the chromatic visual system can be modelled as a low-pass temporal frequency filter (Swanson, Ueno, Smith, & Pokorny, 1987). Sensitivity to modulations decreases above a certain modulation frequency (approximately 4Hz in the study by Swanson et al.) and therefore the adaptive effects of stimuli above this frequency are smaller. This explains why the adaptive effects of our 3Hz cardinal-directions modulation stimuli do not appear large, and the 10Hz stimuli even less effective, but the finding that we can achieve similar results to previous studies with higher frequency stimuli is interesting. As discussed previously, using appreciably lower-frequency modulations for our CW and adaptation stimuli is unlikely to have caused significant adaptation, and indeed this is what we found.

Previous experiments (Nissen & Pokorny, 1977) have demonstrated that differences in RTs in a detection task arise when isoluminant stimuli are used, but when stimuli appear on a dimmer background, the effect disappears. They reason that this is because the luminance cue provides an additional cue to the onset of the stimulus. The sections of our test stimuli were each randomly offset in luminance from the background, and the task required discrimination between sections with different chromaticities that will often have varied in luminance, however the luminance differences did not correspond to the chromaticity differences and the luminance changes were not a reliable cue to the task. We assume that this should allow us to measure differences in RT due to chromaticity.

Several studies (e.g. McKeefry, et al., 2003; Smithson & Mollon, 2004) have shown that RTs to S-cone stimulation isolation stimuli are longer than those to L/M-cone isolating stimuli. Thus, we might expect main effects of the test stimulus chromaticity on RTs in our cardinal axes test

conditions. This was only the case for observer HES and 3Hz stimuli, where these effects are clear and significant.

The results of the previous experiment, particularly for HES, suggest a further reason why the effects of 10Hz CW and CCW adaptation on diagonal test chromaticity RTs might be weak. The ellipses fitted to the threshold data after CW adaptation are not aligned with their major axis along a diagonal direction of colour space, as they are after CCW adaptation. This means that thresholds along the diagonal directions that we test in this experiment are not as different as they are for observer RJL. Even for RJL, the major axes are not perfectly aligned with the diagonals, or directions at 45° from the cardinal axes in these diagrams, which represents the colour space as it was scaled in both of these experiments. It is possible that cleaner, statistically significant results might be obtained if we used the results of the adaptation experiment to target directions to probe after adaptation to the CW and CCW stimuli.

Another possibility is that differences in RTs are primarily determined by mechanisms early in the visual processing system, such as the classical L/M- and S-opponent mechanisms. This is suggested by McKeefry et al. (2003) whose measured RTs are minimal to L/M stimuli and maximal to S stimuli, and intermediate at intermediate chromatic directions, where the response might be dependent on responses in both mechanisms. In our experiment, with the CW and CCW stimuli, we aim to adapt mechanisms that are later in the system than the opponent process stage and that are tuned to intermediate chromatic directions. It may be the case that these mechanisms are more resistant to adaptation (although the previous experiment indicates the same stimuli can cause measureable threshold changes) or that RTs are less affected by adaptation in these mechanisms. As discussed in the previous experiment, the CW and CCW stimuli are both designed to be matched in the adaptation they cause in the opponent mechanisms.

7 Experiment 4: The time-course of colour constancy

7.1 Introduction

The purpose of this study was to investigate two related questions. Firstly, how much time is required, following a change in illuminant chromaticity, for the visual system to adjust? Secondly, when classifying the colour of a stimulus, how much information is used from the stimuli seen in the time preceding that stimulus?

As discussed in section 1.5, there has been much research into colour constancy and possible models to explain how it is achieved. However, very few studies have sought to determine the amount of time that any of these models take to process the change of illuminant and make the perceptual adjustment to maintain colour constancy. Most studies that have considered temporal effects have focussed on lightness constancy. However, some studies indicate that some degree of colour constancy can be achieved instantaneously (e.g. Barbur & Spang, 2008; Foster, et al., 2001 discussed later).

We perform an experiment in which we measure colour constancy, as calculated by a colour constancy index that will be explained later, at several intervals after a simulated shift of illuminant. We aim to determine if the almost perfect colour constancy that is demonstrated under prolonged exposure to a single illuminant is achieved immediately following the change, or if some time is needed to achieve it. We look for a relationship between the time, or number of samples seen since an illuminant shift, and the colour constancy index. We relate our findings to a model, such as those that are discussed in section 1.5.2, in which the visual system estimates the chromaticity of the background in order to recover the surface chromaticity. We show how the chromaticity of the achromatic point moves from coordinates close to those of the illuminant before a shift to coordinates closer to those of the illuminant following the shift, and this happens gradually over time. We also investigate the amount to which extra scene information, such as simulated background surfaces with illumination consistent with or different to that of the test surfaces. We predict that, more scene information with illumination conditions that are consistent with the illumination on the test stimulus will lead to increased constancy, or constancy increasing at a faster rate, compared with less scene information or scene illumination inconstant with the test stimulus (Cataliotti & Bonato, 2003).

In this experiment, we ask observers to make classifications (rather than make colour matches) of a sequence of simulated reflecting surfaces, each having the reflectance function of a natural material, rendered under one of two simulated illuminants at any given time. The test materials are presented against a background which, in two of three experimental conditions was presented against a variegated background of other simulated materials. The simulated illuminant

of the background region and test surface region could be independently specified so that we could investigate the effects information in the surround on classifying the colour of a stimulus. In the other condition, the test material was presented in isolation, against a black background.

From these classification results we calculate the chromaticity coordinates of the point considered achromatic at each temporal position after the illuminant shift. We make predictions about how the achromatic point moves, and therefore the degree of constancy achieved, as function of time assuming that constancy is achieved perfectly and instantaneously. In our first experimental condition the test patch and background always have the same illuminant (global condition) and this alternates between two spectra. If colour constancy were perfect and instantaneous, the achromatic point will shift from the chromaticity of the illuminant before the shift, to the chromaticity of the illuminant after the switch, immediately.

When the test patch was presented on a black background (isolated condition) and the test illuminant alternated as before, we expect that any change in the position of the achromatic point will occur gradually. This condition cannot strictly be considered a test of colour constancy. Presented on a black background there is very little information (certainly no spatial information, and only a little temporal information) to suggest that the test patch is a simulated material under a particular illuminant. Reflectance changes are indistinguishable from illuminant changes; however the mean of a given number of previous test chromaticities will change gradually after an illuminant shift. At first, there will be a bias in the observers' responses but this will decrease as the internal standard is renormalized (Morgan, Watamaniuk, & McKee, 2000). This is a process that could contribute towards colour constancy, but it is not an explicit demonstration of colour constancy. This is similar reasoning to the classical literature on successive judgements by Helson and others (e.g. Helson, 1948).

In our third condition (inconsistent condition) the illuminant on the test patch alternates as in the previous conditions while illumination of the background region remains constant at an illuminant spectrum that has a chromaticity in between those used for the test patch. In this case we predict that, if all judgements are made relative to the constant spatial context, the achromatic point will remain at the coordinates of the background illuminant.

We also look at the classification results from the point of view not of changing illuminants, but of the effect of the sequence of chromaticities seen previously on the classification of the next. We predict that the chromaticity seen on a given trial will have the most predictive effect on the classification of that stimulus, and that chromaticities seen on previous trials will have less effect. We may also see the chromaticities on previous stimuli having the opposite predictive effect on classification, so for example seeing a blue stimuli (one with a high $S/(L+M)$ coordinate) that the

observer classifies as blue) will increase the probability of classifying the next stimulus as yellow and this may be through an adaptive effect which serves to shift the neutral point in the blue direction. Benucci, Ringach & Carandini (2009) measured cat V1 responses to stimuli varying in spatial orientation (not chromaticity) and compared responses of single stimuli to those of successive stimuli. They found these response patterns were different, but successfully predicted the responses to sequences of stimuli with simple summation of responses. Although we don't employ such a model here, it suggests that previous stimuli seen could affect the classification of the next. Nikolić, Häusler, Singer and Maass (2009) also made recordings from cats presented with sequences of letters of the alphabet, and found that responses to successive stimuli contained were highly dependent on the previous stimulus, and that information about previous stimuli was available for some time after their offset.

7.2 Methods

7.2.1 Observers

Three observers participated in the study. Observers RJL and WS were experienced psychophysical observers; observer KAD was not. Observer WS was naive to the purpose of the study; observers RLJ and KAD were not. Observers RJL and WS are male, observer KAD is female.

7.2.2 Apparatus

Stimuli were presented on a CRT monitor (Mitsubishi DiamondPro 2070SB). The monitor display had size 0.386×0.295m, a spatial resolution of 1024×768 pixels (each pixel measured 0.4mm) and a refresh rate of 100Hz. The stimuli were generated by a Cambridge Research Systems ViSaGe system controlled by custom software. The monitor was surrounded by a box whose interior was painted matt black and which served to prevent light from the monitor reflecting off other objects in the room and then reaching the observer. The observer viewed the monitor binocularly through holes in the opposite side of the box, and was stabilised with a chin rest. Goggles were placed around the holes to ensure that as little light as possible from the room entered the eyes. Matt black baffles were placed around the monitor screen and viewing holes. The purpose of this arrangement was to isolate the stimulus from possible light sources in the room that might affect judgement of the simulated illuminant, and also to prevent multiple reflections from the monitor inside the box that might provide the observer with a spatial average of the stimulus chromaticity. The distance between the observer and the monitor face was 0.56m, and the horizontal dimension of the monitor subtended a visual angle of 38°. A schematic of the apparatus arrangement, with more measurements, is shown in Figure 37.




Figure 37. Arrangement of monitor and observer in this experiment, viewed from the top (top diagram) and side (side diagram).

7.2.3 Stimuli

Each scene comprised a background and test patch. The background was made up of 252 ellipses, each randomly oriented and positioned across the display. The size of each ellipse was also randomised, so that one axis of the ellipse was fixed at 50 pixels (approximately 2°) and the other varied between 25 and 100 pixels (approximately 1° to 4°). The ellipses were randomly placed sequentially, so many overlapped. The random positions, sizes and orientations were generated each trial. Each ellipse was assigned a spectral reflectance function, from the set of 120 measurements of materials (see below), and rendered under the specified illuminant. All the ellipses in the background were illuminated by the same simulated light. The test patch was a square of side length 100 pixels (approximately 4°), placed in the centre of the screen. The test patch was not overlapped by any of the background ellipses. The patch was assigned a spectral reflectance function from a pre-determined sequence that will be described later. The illumination of the background and test patch could be different, depending on the test condition used, and could be changed every 10 trials.

The chromaticity of each material under each illuminant, and therefore the colours to be displayed by the monitor, were calculated using the spectral reflectance and radiance functions of the materials and illuminants, respectively. The procedure is described below. The material reflectances were from spectrophotometric measurements of various man-made and natural

surfaces, obtained from various sources (Chittka, Shmida, Troje, & Menzel, 1994; Hiltunen, 1996; Marshall, 2000; Vrhel, Gershon, & Iwan, 1994). The simulated illuminants were measurements of sunlight and skylight (measured by Taylor & Kerr (1941)), as well as the standard spectrum of CIE Illuminant C, which has a chromaticity intermediate between sunlight and skylight. Examples of the stimuli are shown in Figure 38.




Figure 38. Examples of stimuli in this experiment. The left panel shows background and test under sunlight, the middle shows background under skylight and test patch under sunlight, and the right panel shows the test under sunlight and the background unilluminated (the simulated illuminant was a zero-energy-spectrum). Note that display and printing properties make it impossible to exactly reproduce here the colours seen in the experimental conditions.

To give an impression of the chromaticities of the test material reflectances as used in the experiment, in a colour space that might be familiar to most, the average, minimum and maximum CIE 1931 (x,y) coordinates (CIE, 1932) of the set of materials rendered under the sunlight and skylight illuminants is given in Table 4. The complete set of reflectances is plotted in Figure 39.

Table 4. Mean, minimum and maximum CIE 1931 (x,y) coordinates of the set of materials used in the experiment, when rendered under the sunlight and skylight illuminant.

Illuminant	mean	min	max
Sunlight	(0.334,0.349)	(0.256,0.244)	(0.422,0.486)
Skylight	(0.263,0.281)	(0.206,0.171)	(0.337,0.444)




Figure 39. The chromaticities of the set of material reflectances used in the experiment when rendered under the sunlight (yellow crosses) and skylight (blue crosses), on the CIE 1931 colour space diagram. The black curve represents the spectrum locus.

7.2.4 Generation of chromaticities

The desired colour of any patch of the stimulus display was generated by combining the spectral power distribution of the desired illuminant, with the spectral reflectance function of the desired material. The two distributions were multiplied together to give the spectral power distribution of the light from the illuminant after reflection from the material.

This distribution was then multiplied by each of the three cone sensitivity fundamentals. This gives a triplet of values, representing the relative cone excitations of this spectral distribution, for the standard observer. These desired cone excitations were then converted to monitor primary output levels by multiplying by the inverse of a matrix constructed by multiplying the cone fundamentals by measured spectral power distributions of the monitor primaries, as discussed in detail in section 3. This complete procedure is described in the following set of equations:

$$\begin{pmatrix} p_R \\ p_G \\ p_B \end{pmatrix} = A^{-1} \begin{pmatrix} F_L(\lambda) \\ F_M(\lambda) \\ F_S(\lambda) \end{pmatrix} I(\lambda)R(\lambda)$$

Where:

$$A = \begin{pmatrix} P_R(\lambda) \\ P_G(\lambda) \\ P_B(\lambda) \end{pmatrix} \begin{pmatrix} F_L(\lambda) \\ F_M(\lambda) \\ F_S(\lambda) \end{pmatrix}'$$

$P_R(\lambda)$, $P_G(\lambda)$, $P_B(\lambda)$ are the spectral output distributions of the red, green and blue monitor primaries, respectively, as functions of wavelength λ . $F_L(\lambda)$, $F_M(\lambda)$, $F_S(\lambda)$ are the spectral sensitivities of the L-, M- and S-cones respectively (Stockman & Sharpe, 2000; Stockman, et al., 1999), $I(\lambda)$ is the spectral power distribution of the illuminant to be simulated, as a function of λ , and $R(\lambda)$ is the reflectance function of the material to be simulated, as a function of λ . p_R , p_G , p_B are the fractions of maximum output of the red, green and blue monitor primaries.

As can be seen from the above equations, conversion to cone excitations is not strictly necessary. However, it does allow a device-independent representation of all the illuminant and reflectance spectra combinations, in terms of three values rather than whole spectra. These calculations can be done once, and cone excitations stored, requiring less processing and memory than would be required to do it on-line during the experiment. Matrix A can then be calculated with measurements of the primary outputs of the apparatus to be used.

In practice, all the spectra are stored as tabulated values every 5nm, not functions, so products of spectra are obtained by multiplying the two tables element-by-element. The column vectors shown above are, in fact, $3 \times n$ matrices, where n is the number of elements in the list of values representing each spectra, and must be the same across two matrices that are to be multiplied.

The illuminant spectra were individually linearly scaled so that after the above calculations were carried out, all materials under all illuminants resulted in chromaticities that were within gamut of the monitor.

7.2.5 Procedure and task

The data collected should allow classification of all the simulated materials in the test set, under each of the illuminants and at regular intervals after changes in illumination conditions, as either red or green and either blue or yellow.

At the start of each trial, the screen changed to show the appropriate stimulus. The observer had 1200ms, beginning immediately, to make a response via a button box. The observers' task was to decide if the test patch appeared red or green, in one set of sessions, or blue or yellow in another set of sessions. Responses were given by pressing one two buttons, one corresponding to red or blue, the other corresponding to green or yellow. The same stimuli were repeated in separate sessions (see below) so that red vs. green and blue vs. yellow responses could be obtained for all stimuli. Audio tones signalled the onset of each trial and the acquisition of a response. The computer required several milliseconds to draw the next stimulus, so this was done after the

response period, before the next trial, and was done in off-screen video memory. The inter-trial interval was 2100ms.

In the first three of four experimental conditions, the test patch illuminant changed every ten trials. In the first condition (global), the background illumination was always the same as the test patch illumination, and this changed from sunlight to skylight or vice-versa after every 10th trial. In the second condition (isolated), the background illumination was always a zero-energy-spectrum, so the entire background field appeared black, while the test patch illumination changed as before. In the third condition (inconsistent), the background illumination was always Illuminant C while the test patch illumination alternated between sunlight and skylight as before.

Within each session, the presentation order of stimuli was pseudo-randomised after being arranged to satisfy two requirements: each material had to be presented in the test patch position under both of the two illuminants, and in each of ten temporal positions (trials) after an illuminant change. Since this would require a lot of trials, the set of 120 materials was split into four sets labelled A, B, C and D, each run in a different session and each fulfilling the two requirements above. All the sessions were run twice, so that data could be averaged later to reduce noise. This leads to a total of sixteen sessions in each of the three conditions: four sets of materials, two decisions and two repeats. The requirement to have each material under each illuminant in each temporal position means that each session contained 600 trials and lasted for approximately 40 minutes.

The final experimental condition was designed to determine the observer's classification of each material under each illuminant, under constant illumination conditions. The test patch and background illuminants were constant throughout each session. The test patch illuminant was either sunlight, skylight or Illuminant C, and the background illuminant was either sunlight, skylight, Illuminant C, or a zero-energy-spectrum. Every combination of background and test patch illuminant was used, except that the test patch under Illuminant C was only tested a background that shared the same illuminant. Each session was again repeated so that red vs. green and blue vs. yellow responses could be collected, and then to collect more data for averaging. This means there were 36 sessions required in this condition. Since there was now only one illuminant present in each session and there was now no requirement to show each simulated material in every temporal position after an illuminant change, all the materials could be presented in the test patch position in the same session and in any random order, which therefore meant that each of these sessions required 120 trials and lasted for 10 minutes.

7.3 Results

The data gathered in this experiment can be analysed in a number of ways. The first stage in many of these is to determine the boundaries in colour space that divide chromaticities classified as red from those classified as green, and those classified as blue from those classified as yellow.

7.3.1 Classification boundaries and achromatic points

Each of these boundaries is found by performing a regression analysis on the classification probabilities. That is, the probability of classifying the material as red when looking for the red/green boundary or blue for the blue/yellow boundary. The regressors used were the chromaticity and luminance coordinates of each material, and those coordinates squared. This generates a polynomial function predicting classification probability from chromaticity coordinates, and is done separately for red/green and blue/yellow classifications.

We assume that a classification probability of 0.5 (an equal chance of classifying a material as e.g. red or green) defines the boundary we seek. This probability is represented by a surface in the three dimensional space whose axes are defined by the two chromatic axes and luminance (MB-DKL space), or considering any equiluminant plane, a line created by the intersection of the surface and this plane. Examples are shown in Figure 40.

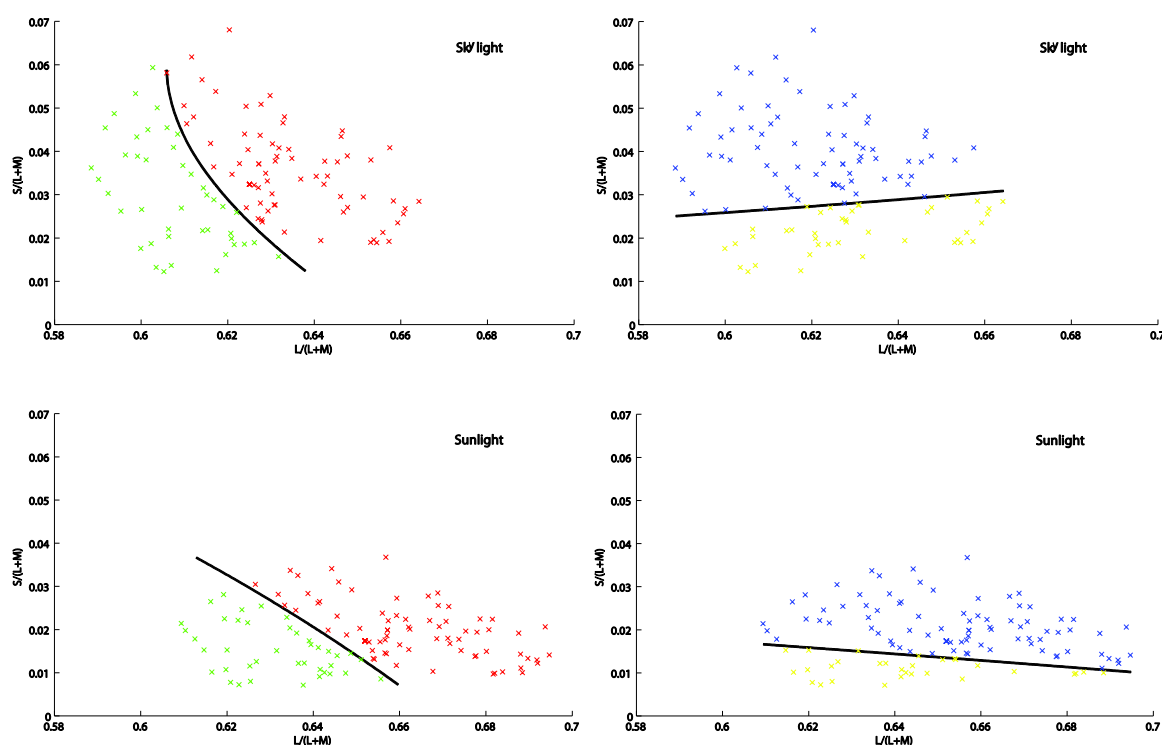


Figure 40. Examples of boundaries (black lines) fitted between chromaticities classified as red and green (shown with correspondingly coloured x symbols on the left plots) or blue and yellow (on the right plots). The intersection of these lines is taken as the achromatic point. Each plot represents a portion of the MacLeod-Boynton chromaticity diagram. Classifications were made by observer RJL in conditions when the illumination conditions were constant. The top plots show the chromaticities and their classifications of the set of reflectances under skylight and the bottom plots show the chromaticities and classifications of the same reflectances under sunlight.

We plot these two lines, taking the equiluminant plane through the mean luminance of all those in the set of materials, rendered under the given illumination conditions. The intersection of these lines should represent the point in chromaticity space judged to be neither red, nor green, nor blue, nor yellow: the achromatic point. It is this single point that we use as a reference when comparing the effect that changes in simulated illuminant have on perception. The data we collected allows us to plot the position of the achromatic point, for each of the two illuminant transitions (sunlight to skylight and skylight to sunlight) at ten temporal intervals after the illuminant switch.

Figure 41, Figure 42, and Figure 43 show these points for the global, isolated and inconsistent conditions respectively. In each figure, the left hand panel shows data from the observer KAD, the middle panel shows data from the observer RJL, and the right panel from observer WS. Each panel represents a portion of the Macleod-Boynton chromaticity diagram, and plots the path of the achromatic point as it moves over time after an illuminant change from sunlight to skylight (blue line) or skylight to sunlight (yellow line). In all cases, the square symbol represents the position of the achromatic point measured on the trial immediately after the illuminant switch, the jagged line joins achromatic points at successive temporal intervals until the 10th trial after the illuminant switch, represented by the circular symbol. The crosses represent the achromatic points measured in the fourth condition described in the methods (section 7.2.5) when all the materials are viewed under an illuminant which does not change and when both test patch and background were rendered under the same illuminant (as in our first, global, condition). We refer these points as the ‘anchor points’, and as described in the introduction to this experiment, we might expect that after an illuminant change, the position of the achromatic point shifts over time from its chromaticity under the illuminant before the shift to its chromaticity under the illuminant after the shift. This can be seen in the data from the global condition, and less clearly in the isolated and inconsistent conditions.




Figure 41. A section of the Macleod-Boynton Chromaticity Diagram showing the path taken by the achromatic point as it moved following a change of illuminant from skylight to sunlight (yellow lines) and sunlight to skylight (blue lines). The squares show the position measured on the trial immediately after the change, and the circles show the position on the tenth trial after the change. Data are show in separate plots for all three observers, for the global condition. The crosses show the position of the

achromatic point measured during prolonged viewing under one illuminant (yellow: sunlight, blue: skylight).




Figure 42. As Figure 41, for the isolated condition.




Figure 43. As Figure 41, for the inconsistent condition.

7.3.2 Colour constancy index

A common way of quantifying colour constancy is to calculate a constancy index. This is usually a value of the relationship between the change in chromaticity of the achromatic point after a change of illuminant, and the change in chromaticity of a simulated material of uniform spectral reflectance (equivalent to the chromaticity of the illuminant) (Brainard, 1998). In this thesis, we will use indices similar to those by Smithson & Zaidi (2004). These constancy indices can be calculated independently for each class of cone signal or dimension in colour space: as discussed in section 1.5, changes of illuminant result in approximately multiplicative scaling of cone coordinates and S-opponent signals, but translational scaling of L/M opponent signals. Different methods of calculating indices are required for the two types of classification. Smithson & Zaidi's (2004) constancy index used for signals undergoing multiplicative scaling (S opponent coordinate) is:

$$C_S = \frac{\log(a_1/a_2)}{\log(b_1/b_2)} \quad (54)$$

and for signals undergoing translational scaling (L/M opponent coordinate) it is:

$$C_{LM} = \frac{|a_1 - a_2|}{|b_1 - b_2|} \quad (55)$$

where a_1 and a_2 are the coordinates (on the appropriate axis) of the achromatic point measured under the first and second illuminants, respectively, and b_1 and b_2 are the coordinates of the first and second illuminants. $C=0$ indicates no colour constancy: a_1/a_2 (or $|a_1-a_2|$) is small compared to b_1/b_2 (or $|b_1-b_2|$), suggesting that the neural transformation of the achromatic point is very small compared to the shift of the illuminant. $C=1$ indicates perfect colour constancy: the neural transformation is equivalent to the illuminant shift.

The index we use in this thesis instead compares the difference in position of the achromatic point following the illuminant shift and after prolonged viewing under the previous illuminant (the anchor point), with the difference in positions of the anchor points for the two illuminants:

$$C_S = \frac{\log(a_1/g_2)}{\log(g_1/g_2)} \quad (56)$$

and

$$C_{LM} = \frac{|a_1 - g_2|}{|g_1 - g_2|} \quad (57)$$

where g_1 and g_2 are the coordinates of the anchor points of the first and second illuminants.

In the equations used by Smithson and Zaidi, the constancy index represents the amount of compensation in appearance (a_1 to a_2) relative to the physical change in illuminant chromaticity (b_1 to b_2). The equations used here represent the compensation in appearance (i.e. achromatic point on a given trial under the new illuminant, a_1 compared the achromatic point under prolonged exposure to the previous illuminant, g_2) relative to the perceptual difference that accompanies the physical change in illuminant (i.e. the change in position of the achromatic point, measured under prolonged exposure each illuminant, g_1 and g_2).

Figure 44, Figure 45 and Figure 46 show, for all observers, how the calculated constancy indices vary with temporal position after illuminant change in the same experimental conditions as above. The top row of panels plot the constancy index for the L/M opponent direction, as a function of the temporal position after illuminant switch, again when this switch is from sunlight to skylight (blue line) or skylight to sunlight (yellow line). The bottom row of panels shows the constancy indices for the S opponent direction, with the same notation as in the top row.




Figure 44. Colour constancy index against number of trials since an illuminant shift (t), with linear regression lines fitted. Data are for all observers in the global condition. The top row show the constancy indices for the L/M opponent chromatic direction, and the bottom row show the constancy indices for the S opponent chromatic direction. The yellow symbols show indices calculated after a transition from skylight to sunlight, and the blue symbols show indices calculated after a transition from sunlight to skylight. The lines are fitted by least-squares linear regression.

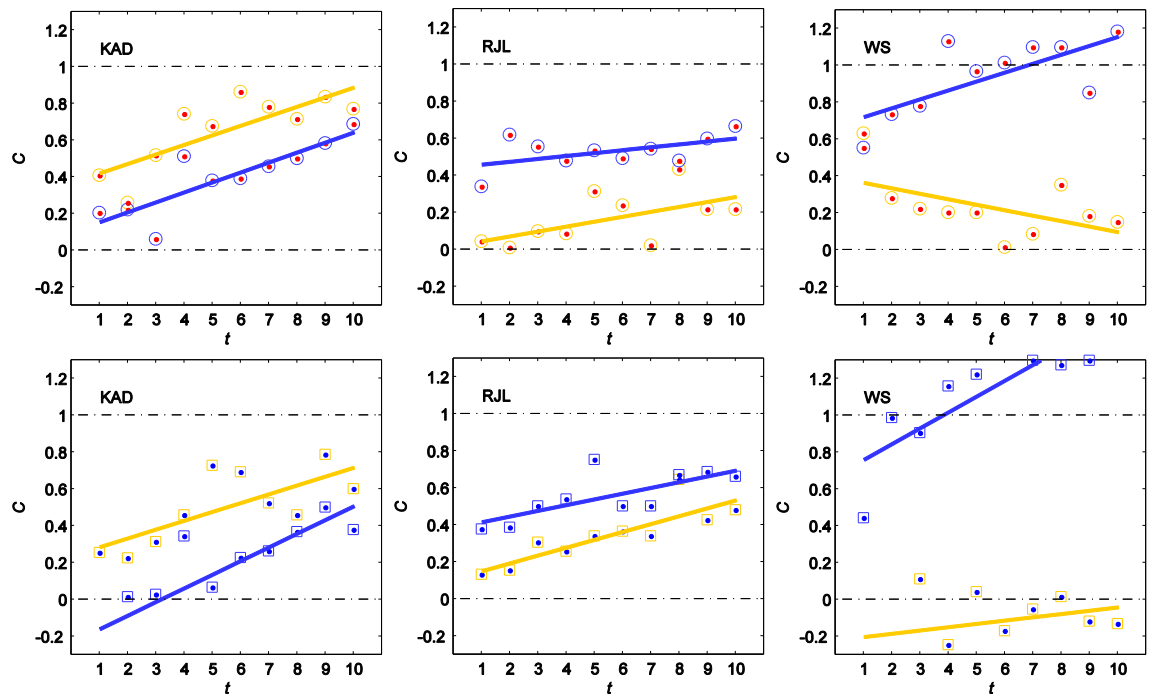


Figure 45. As Figure 44, for the isolated condition.

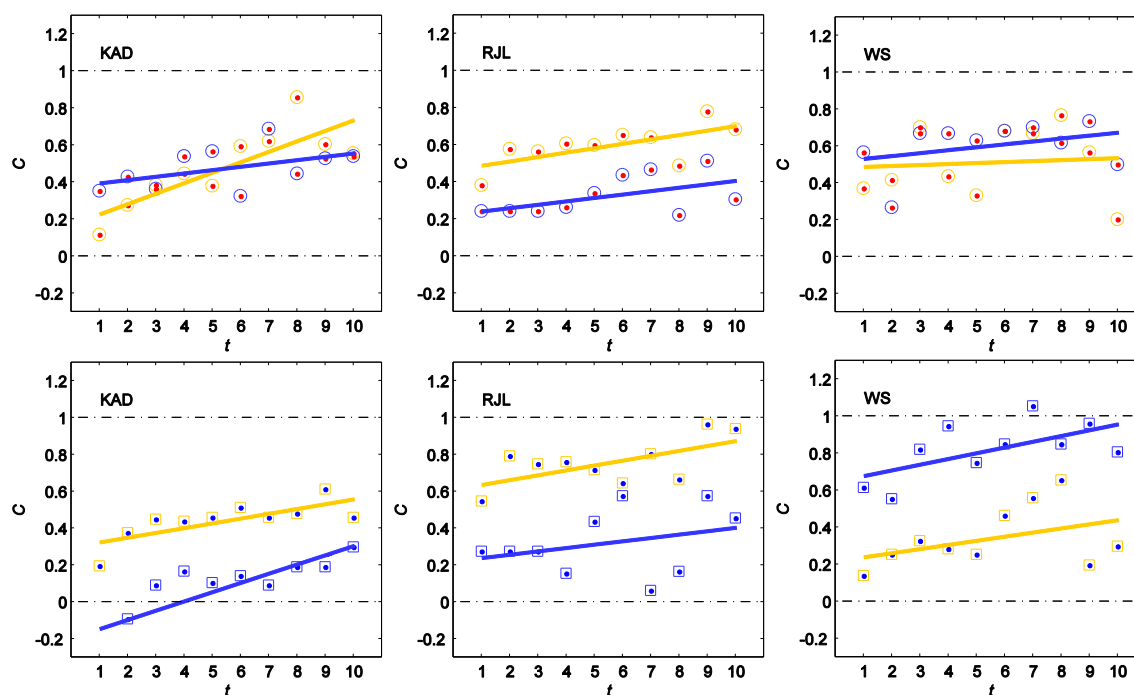


Figure 46. As Figure 44, for the inconsistent condition.

7.3.3 Linear regression

The middle and bottom row of panels Figure 44, Figure 45 and Figure 46 have been fitted with linear least-squares regression lines, for both the sunlight to skylight and skylight to sunlight data. Each equation takes the form

$$C = mt + c \quad (58)$$

and the values for parameters m and c are given in Table 6. t-tests were then performed to test the hypothesis that the gradients of these lines are different from zero, and therefore that the constancy index is dependent on temporal location. Those values of m that are significant are highlighted with an asterisk in Table 6.

For observer KAD, for all except one regression line, there is sufficient evidence to reject the null hypothesis and conclude that there is a significant relationship between constancy index and temporal position. Furthermore, the gradients of all the regression lines are positive. This suggests, as expected, that colour constancy increases with time after an illumination shift. For observers RJL and WS, there is insufficient evidence to suggest that the majority of the regression lines have a gradient significantly different from zero.

7.3.4 Nonlinear model fits

It is clear from looking at the plots of colour constancy index against temporal position, however, that a linear regression is not completely appropriate. A function whose slope is greatest at smaller values of the time since the illuminant switch, and decreases as the time increases, might provide a better fit than a linear one.

The function

$$C = p \ln(t + q) \quad (59)$$

where t is the temporal position since the illuminant shift, C is the constancy index, and p and q are parameters, was fit to the data in each of the conditions by adjusting p and q to minimise the sum-squared-error. This was done with the simplex search algorithm implemented by the MATLAB (The Mathworks, Inc.) `fminsearch` function. The values of parameters m and c are given in Table 6. The data in Figure 44, Figure 45 and Figure 46 are re-plotted Figure 47, Figure 48 and Figure 49 with these logarithmic curves fitted.

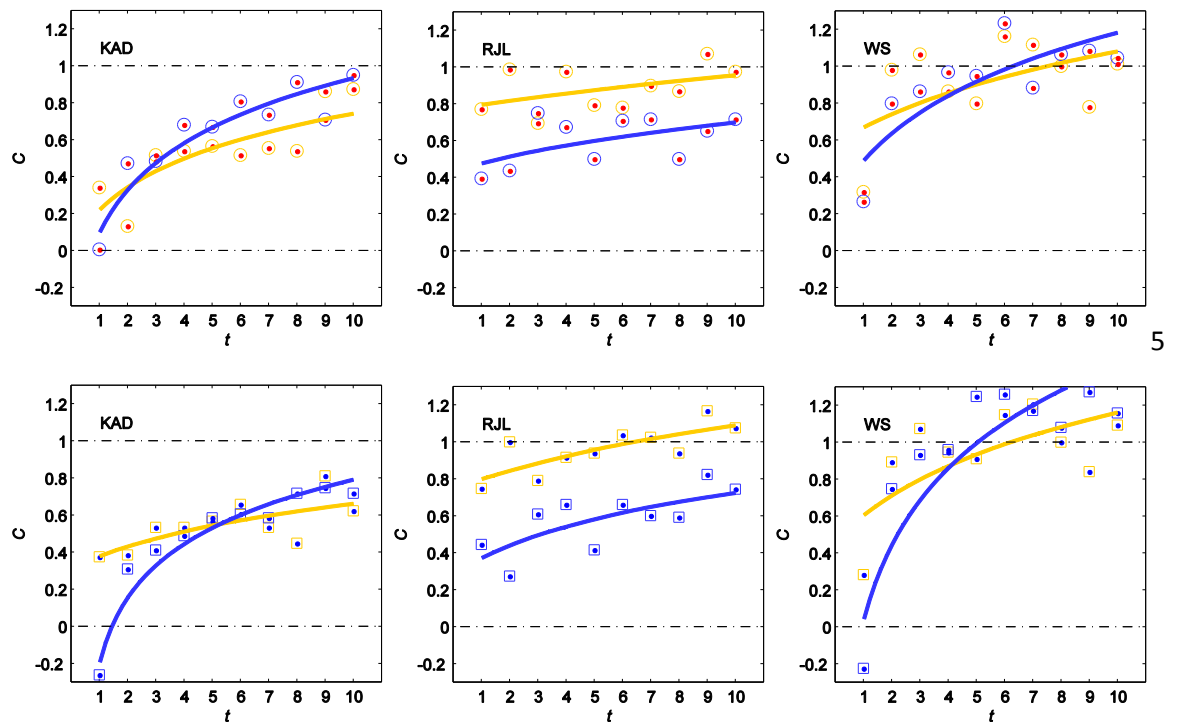


Figure 47. The data from Figure 44, re-plotted with logarithmic curves of the form of Equation 59 fitted by least-squares regression.

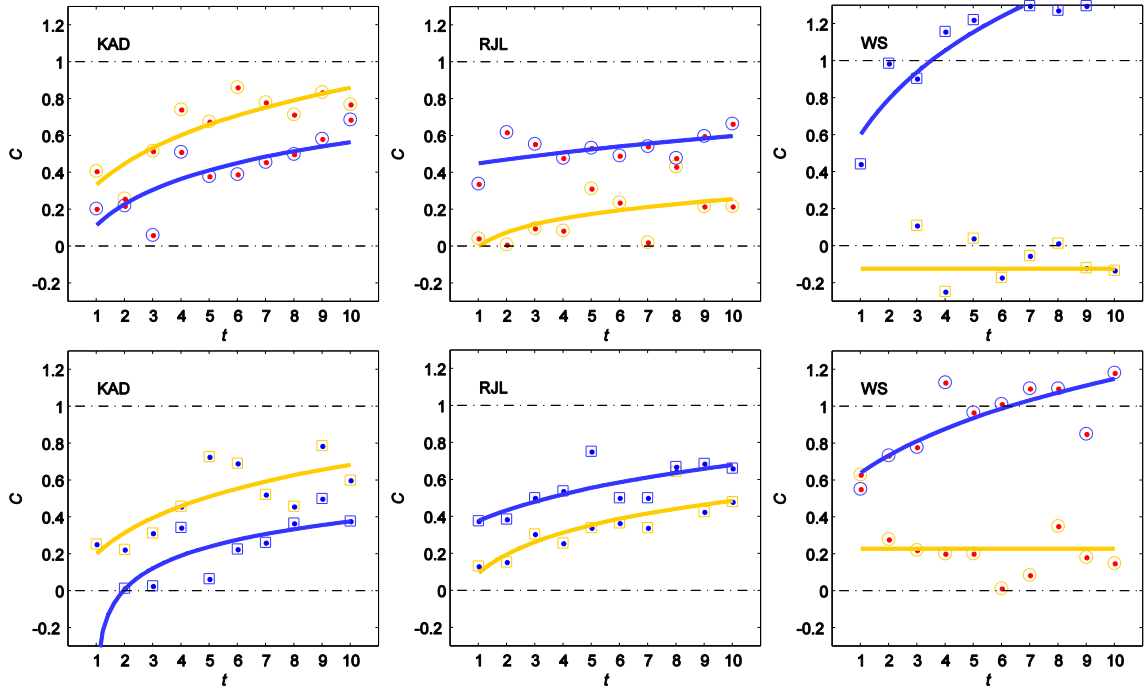


Figure 48. The data from Figure 45, re-plotted with logarithmic curves of the form of Equation 59 fitted by least-squares regression.

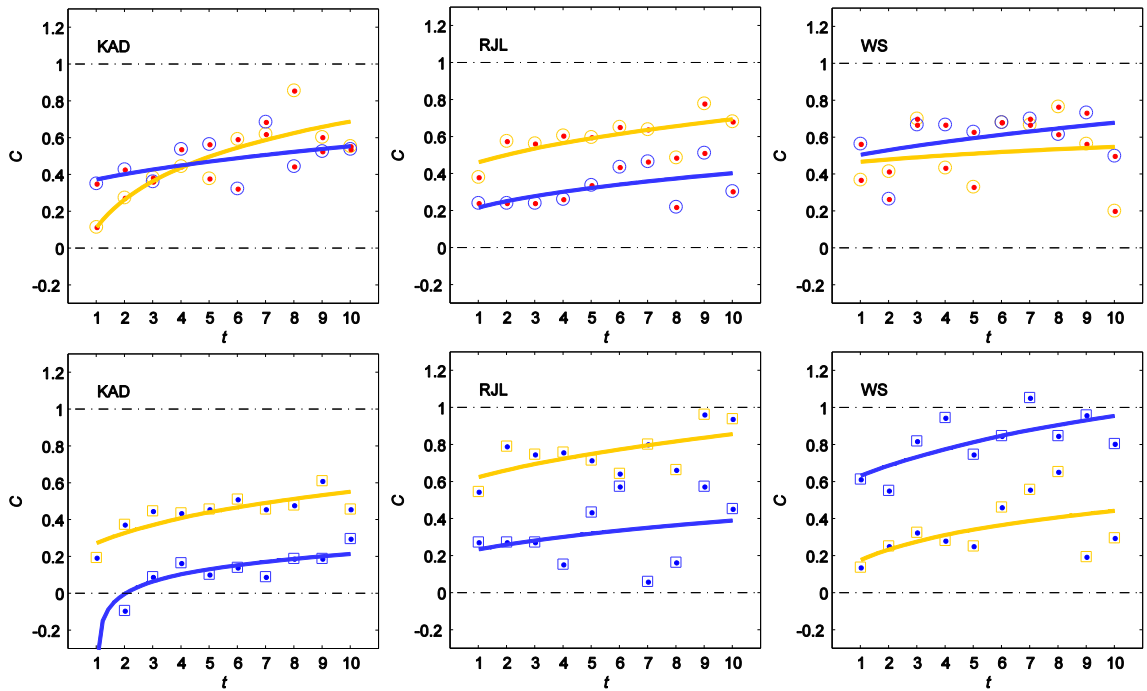


Figure 49. The data from Figure 46, re-plotted with logarithmic curves of the form of Equation 59 fitted by least-squares regression.

Again, the equations of the fitted line are on the plots. By eye, these curves appear to provide much better fits to the data, although when comparing the root-mean-squared (RMS) error of any given fitted curve and that of the linear regression model fitted to the same data, the reduction appears very small. In fact, in some cases, the RMS error is greater for the logarithmic fit than for the linear one. However, a two-way ANOVA, on the RMS values, comparing the effects of

observer and regression type on RMS error, did not indicate a significant difference between the RMS errors from the linear and logarithmic fits.

7.3.5 Projected colour constancy values

A simpler way of representing the data above is to combine the separate sets of chromatic constancy values for the L/M and S opponent chromatic directions into one set of values. This can be done by calculating the projection of the constancy values for the two directions on to a vector joining the two chromaticities of the achromatic points as measured under prolonged viewing of stimuli under the two illuminants (the ‘anchor points’ as described as above). This index represents the amount by which the achromatic point has moved after the illuminant change, relative to the amount by which the illuminant chromaticity (or at least the achromatic point obtained under prolonged viewing, which is close to the illuminant chromaticity) has moved, considered only in the chromatic direction of the illuminant shift. This is calculated with equation 60.

$$C = \frac{(\mathbf{a} - \mathbf{g}_2) \cdot (\mathbf{g}_2 - \mathbf{g}_1)}{(\mathbf{g}_1 - \mathbf{g}_2) \cdot (\mathbf{g}_2 - \mathbf{g}_1)} \quad (60)$$

where \mathbf{a} is the chromaticity coordinate vector ($L/(L+M)$, $S/(L+M)$) of the achromatic point measured under the first illuminant, and \mathbf{g}_1 and \mathbf{g}_2 are the chromaticity coordinate vectors of the achromatic point measured under prolonged viewing under the first and second illuminants, respectively.

Firstly, we predict that when we measured the achromatic points under prolonged viewing under a constant illuminant, with both background and test patch under the same illuminant (our ‘anchor points’), the chromaticity of the achromatic point should be very close to that of the illuminant. The chromaticities of the illuminants and these achromatic points are shown in Figure 50.




Figure 50. Chromaticity coordinates of the simulated illuminants and of the achromatic points measured under prolonged viewing under constant illumination conditions over the whole stimulus.

Using equation 60 (although replacing g_1 and g_2 with b_1 and b_2 , the chromaticity coordinate vectors of the illuminants), we calculate a constancy index for each achromatic point relative to the chromaticity of the illuminant under which it was measured, projected onto a direction joining the chromaticities of the sunlight and skylight illuminants. If this prediction is correct, all these indexes should be high. The indexes are given in Table 5:

Observer	$C_{Sunlight}$	$C_{Skylight}$
KAD	0.88	0.78
RJK	0.90	0.94
WS	0.86	0.68

Table 5. Constancy indices expressing the coordinates of achromatic points under prolonged viewing under constant illumination conditions over the whole stimulus, relative to the chromaticity of the illuminant, projected onto a direction joining the chromaticities of the sunlight and skylight.

The data from Figure 47, Figure 48 and Figure 49 are now plotted again in Figure 51, Figure 52 and Figure 53 but this time showing this projected colour constancy index. The data points now appear less noisy than before, and the logarithmic curves appear to fit well. Since we are now fitting models to different values, we cannot show statistically that these fits account for more variability than curves fitted to the constancy indices separately along the two chromatic axes.

The logarithmic nature of the model suggests that as the colour constancy index C increases, the time since the illuminant change t required to achieve this index increases by a constant proportion. Alternatively, increasing the value of t by a constant proportion increases C by a constant amount, or, the increase in t required to achieve a given increase in C increases as t increases. For all three observers, the data points and model fits suggest that the constancy index reaches a higher magnitude in a shorter time in the global and isolated illumination conditions than in the inconsistent condition.




Figure 51. Colour constancy index, calculated along a vector joining the anchor points for the two illuminants, against number of trials since an illuminant shift (t), with logarithmic regression lines fitted. Data are for subject all observers in the global condition. The colours of symbols and lines indicate the same illuminants as above.




Figure 52. As Figure 51, for the isolated condition.




Figure 53. As Figure 51, for the inconsistent condition.

Table 6. Values of the parameters of the curves fitted to colour constancy indices as a function of time since an illuminant shift. Definitions of the parameters are given in the text.

Condition	Global						Isolated						Inconsistent								
	Model			logarithmic			linear			logarithmic			linear			logarithmic					
	sky>sun	sun>sky		sky>sun	sun>sky		sky>sun	sun>sky		sky>sun	sun>sky		sky>sun	sun>sky		sky>sun	sun>sky				
Observer	m	0.06	*	0.08	*	p	0.31	0.40m	0.05	*	0.05	*	p	0.35	0.24m	0.06	*	0.02	p	0.29	0.20
	c	0.21		0.22		q	1.08	0.30c	0.37		0.10		q	1.63	0.65c	0.17		0.38		0.51	5.46
KAD	m	0.03	*	0.08	*	p	0.26	0.35m	0.05	*	0.07	*	p	0.29	0.17m	0.03	*	0.05	*	0.22	0.10
	c	0.38		0.03		q	3.46	-0.40c	0.24		-0.23		q	1.06	-0.90c	0.30		-0.10		2.54	-0.90
projected	m	0.05	*	0.08	*	p	0.30	0.39m	0.05	*	0.06	*	p	0.34	0.23m	0.05	*	0.02	p	0.28	0.19
	c	0.26		0.18		q	1.41	0.15c	0.35		0.04		q	1.56	0.20c	0.21		0.27		0.76	2.86
RG	m	0.02		0.02		p	0.31	0.26m	0.03		0.02		p	0.11	0.20m	0.02	*	0.02	p	0.26	0.16
	c	0.78		0.49		q	12.64	5.54c	0.02		0.44		q	0.04	8.15c	0.47		0.22		5.19	3.14
YV	m	0.03	*	0.04	*	p	0.39	0.29m	0.04	*	0.03	*	p	0.21	0.26m	0.03		0.02	p	0.30	0.15
	c	0.79		0.37		q	7.00	2.72c	0.11		0.39		q	0.62	3.23c	0.61		0.22		6.91	4.11
projected	m	0.02		0.02		p	0.32	0.26m	0.04	*	0.02		p	0.12	0.21m	0.02	*	0.02	p	0.26	0.15
	c	0.79		0.45		q	11.90	4.61c	-0.03		0.42		q	-0.09	6.48c	0.50		0.22		5.81	3.17
RG	m	0.03		0.06	*	p	0.41	0.48m	-0.03		0.05		p	0.02	0.45m	0.01		0.02	p	0.17	0.24
	c	0.72		0.58		q	4.21	1.80c	0.39		0.67		q	7971.00	3.25c	0.48		0.52		14.30	7.60
YV	m	0.05		0.11		p	0.46	0.61m	0.02		0.09		p	-0.01	0.61m	0.02		0.03	p	0.18	0.35
	c	0.69		0.38		q	2.80	0.08c	-0.20		0.67		q	31095.16	1.72c	0.22		0.65		1.73	5.05
projected	m	0.04		0.07		p	0.41	0.51m	0.06		0.06		p	0.11	0.48m	0.01		0.02	p	0.18	0.25
	c	0.72		0.55		q	4.03	1.38c	-0.20		0.67		q	-1.00	2.76c	0.45		0.53		9.76	6.77

7.3.6 Within observer differences between illumination conditions

In order to test predictions about the three different spatial illumination conditions, we calculate the difference between the values of constancy indices between conditions. These data are shown below. Figure 54 shows the differences between the indices in the global and isolated conditions ($C_{global} - C_{isolated}$) and Figure 55 shows the differences between the indices in the global and inconsistent conditions ($C_{global} - C_{inconsistent}$).




Figure 54. Differences between calculated colour constancy indices in the global and isolated conditions, as a function of the time since an illuminant shift. Yellow symbols show differences after a change from skylight to sunlight, and blue symbols show differences after a change from sunlight to skylight.




Figure 55. As Figure 54, for the differences between calculated colour constancy indices in the global and inconsistent conditions.

These plots indicate that many of the differences between indices are constant over time. t-tests were performed on the gradients of regression lines fitted to the data, to test the hypotheses that the gradients were different from zero. All the gradients were not significantly different from zero except that of the differences between the global and inconsistent conditions, for illuminant changes from sunlight to skylight for observer KAD ($t_{(8)}=2.12, p<0.05$). t-tests were also used to test the hypotheses that the mean differences were different from zero, in order to determine whether there were differences between conditions that were not dependent on time. Those differences that were significant ($p<0.05$) were between the indices in the global and isolated conditions after illuminant shifts from skylight to sunlight for observers RJJ and WS and after shifts from sunlight to skylight for observer KAD, and in the global and isolated conditions after shifts from skylight to sunlight for observers RJJ and WS and after shifts from sunlight to skylight for all three observers.

7.3.7 Logistic regression on chromaticity coordinates

A second approach to interpreting these data is to perform logistic regression to investigate the degree to which the previous chromaticities seen predict the classification of the one on a given current trial (a similar analysis was done by Ludwig, Gilchrist, McSorley, & Baddeley, 2005). Logistic regression predicts the probability of an outcome based on one or more predictors. For this analysis, for each observer, we attempt to predict the probability of classifying a given material-illuminant combination as red or green, or separately, blue or yellow, based on the three colour coordinates ($L/(L+M)$, $S/(L+M)$, and $L+M$) of the previous ten material-illuminant combinations seen in the test patch position. This gives 30 predictors for each classification (three chromaticity coordinates, ten preceding trials). We can only do this for every 10th trial, beginning with trial 10 ($t = 10, t = 20, t = 30\dots$), so that the chromaticity of any one stimulus is not used as more than one predictor. Because of this, there are no illuminant changes during the sequence of ten trials used to predict each classification, and so all conditions are equivalent and we can combine the three conditions. Since all three conditions are identical in their temporal properties, we make the assumption that the effects of preceding stimuli on classifications will not differ between conditions so we can combine data from the three conditions. This assumption would not be valid if the spatial effects were dominant, but we have some evidence from the preceding analysis, and from previous experiments (e.g. Smithson & Zaidi, 2004) that the temporal properties have a dominant effect. The logistic regression procedure (MATLAB's glm function) produces coefficients for each of the predictors, which indicate the degree to which they predict the classification. These regression coefficients are plotted, in Figure 56, as a function of the temporal position t since the illuminant shift. Data for each observer are shown in separate panels. Several curves are on each plot and represent the coefficients for the three DKL chromaticity coordinates for both chromatic decisions (red or green, and blue or yellow). The red lines represent the coefficients for the red vs. green decision and the blue lines represent those for the blue vs. yellow decision. The lines with red symbols represent the coefficients for the L/M chromaticity coordinate, and the lines with blue symbols represent the coefficients for the S chromaticity coordinate. The lines with no symbols represent the coefficients for the luminance (L+M) coordinate.




Figure 56. Logistic regression coefficients when regressors used are the distances of test chromaticity coordinates.

It can be seen from these graphs that the chromaticity of the test patch on a given current trial ($t=10$) has the most predictive power over the classification. The chromaticities of the previous trials have much less predictive power. For the blue vs. yellow chromatic decision, only the S coordinate seems to have more predictive power on the current trial. However for the red vs. green chromatic direction, both the L/M and S coordinate regression coefficients are much higher for the current trial than for the previous one. This can be explained by looking at the classification boundaries used to locate the achromatic point, examples are given in Figure 40. The boundaries between chromaticities classified as blue and those classified as yellow are almost always quite close to a line in colour space that is parallel to the axis, suggesting that blue vs. yellow classification does not depend on the L/M opponent coordinate much. However, the boundaries fitted between chromaticities classified as red and those classified as green are oriented in-between the cardinal axes of colour space. This suggests that red vs. green classification does depend on both the L/M opponent and S chromaticity coordinates. Since the boundaries are often oriented close to 45° in the equiluminant plane, we might expect the L/M and S coordinates to have similar effects on classification and this is indeed what these logistic regression results suggest.

The logistic regression procedure returns significance (p) values for each of the coefficients. Where $p < 0.05$, this suggests a significant contribution to the prediction of classification probability by the corresponding chromaticity and its coefficient. These significant predictors are indicated by star symbols on above each curve in Figure 56. Each star indicates that a chromaticity coordinate at the corresponding temporal position contributes significantly to predicting, and the symbols to the left of each row of stars indicate the coordinate and chromatic decision (using the same symbol as the curves below) to which the significance refers. For each observer, the regression coefficients for the temporal position immediately preceding the illuminant change ($t=10$) for the red vs. green decision, for both the L/M and S chromaticity coordinates, are significant. This gives further support to the observations above about both chromaticity coordinates appearing to affect prediction of the red vs. green outcome. All the regression coefficients for the first temporal position for the blue vs. yellow decision and the S chromaticity

coordinate are significant, which supports the prediction that this chromatic direction should be the main predictor of blue vs. yellow classification. Significant prediction is also indicated for several other chromaticity coordinates at other temporal positions, although the values of the coefficients here are clearly less than at the first temporal position. This is to be expected, since the results of fitting curves to colour constancy indices as a function of temporal position, as described above, indicate that several trials are required before the constancy index approaches unity and so trials previous to the current one do have an effect on classification probability. In general, it seems that more of the chromaticities closer in temporal position to the current trial have a significant predictive effect than those further away, however there is little consistency across observers.

Many of the plots suggest that that on the trial preceding the current one, many or all of the regression coefficients are negative (and are their predictive effect is significant), although of smaller magnitude than the coefficients for the current trial. This indicates that the chromaticities seen on previous trials have the opposite effect on the classification of the chromaticity on the current trial than the chromaticity of the current stimulus. For example, a stimulus seen on trial $n-1$ with a relatively high L/M coordinate, and so predicted to lead be classified as red, would increase the probability of the chromaticity of the stimulus on trial n being classified as green. This effect appears to be present in most conditions, and for both the L/M and S chromaticity coordinates, however, the magnitude of the predictive power of the stimulus chromaticity of trial $n-1$ is considerably smaller than that of trial n .

7.3.8 Logistic regression on distance from boundary

An alternative approach to this logistic regression that may lead to a more intuitive outcome is to use the distance from the measured chromatic classification boundaries (as described above), instead of chromaticity coordinates, as the predictors.

To find these distances, we use the classification boundaries as found for each observer in the fourth experimental condition, which did not include illuminant changes. These boundaries assume prolonged viewing under a constant illuminant, and boundaries are measured for each observer under both sunlight and skylight illuminants. The distance from the boundary of each material, again for each observer and illuminant, is found by taking the perpendicular distance in colour space between the boundary and the chromaticity of the material under the appropriate illuminant. An example plot showing a boundary and lines connecting material/illuminant chromaticities to the boundary is shown in Figure 57.




Figure 57. An example of the measured distances in colour space between the stimulus chromaticities (x symbols) and the fitted boundary between those classified as red and those as green (shown with correspondingly coloured symbols). The black line is the fitted boundary, and the blue circle symbols plot where lines perpendicular to the boundary and meeting each test chromaticity to the boundary (shown in red and green) intersect the boundary. These coloured lines indicate the distances used as predictors in the regression analysis described in the text. The chromaticities and classifications are the same as in the left panel of Figure 40.

We perform a similar logistic regression procedure as before, again separately for each observer and decision type, but now the distance from the appropriate classification boundary replaces the three chromaticity coordinates. This produces simpler plots, with only one regression coefficient per decision type and temporal position. The data are shown in Figure 58.




Figure 58. Logistic regression coefficients when regressors used are the distances of test chromaticities from the measured chromatic decision boundaries.

In a similar way to the logistic regression on chromaticity coordinates, the boundary-distances of the material of the current trial ($t = 10$) have much more predictive power than those of the other temporal positions. Unlike the previous analysis however, few of the chromatic distances in the trials preceding the current trial ($t = 9$, $t = 8$...) contribute significantly to prediction. In contrast to the previous analysis, for observer RJL, values from earlier trials ($t = 1$, $t = 2$, $t = 3$, $t = 4$) do indicate a significant contribution to prediction. For observer WS, few chromatic distances from earlier trials contribute significantly, but for observer KAD there are several values for which a

significant contribution is indicated. Other than the significant and large magnitude regression coefficient for the chromatic boundary distance of the chromaticity of the current stimulus, there are few similarities between the patterns of significant regression coefficients across the three observers. These results are also contrary to the previous logistic regression and curve fitting analyses, which did suggest that the trials immediately preceding the current one may have more predictive effect than those further away in time.

7.4 Discussion

The plots of colour constancy index (C) against the time since an illuminant shift (t) in section 7.3.3, along with the fitted functions, seem to indicate that the calculated colour constancy index increases over the ten trials after the illuminant shift. While the evidence for this is strongest for observer KAD, the plots for the observers seem to indicate the same for the other observers. Although the patterns of data points seem to indicate a nonlinear relationship and that the increase in constancy index is faster immediately after the illuminant shift, there is no consistent indication that the logarithmic model accounts for more variance than the simpler linear one. What is clear from the plots, is that constancy index relatively low at $t=1$, which indicates that constancy is not perfect on the trial immediately after the illuminant change. Over the following trials, the constancy value does approach 1.0, and this happens faster for some observers and in some conditions than others.

In a model in which achieving colour constancy depends on estimating the chromaticity of the illuminant, this non-immediacy would be equivalent to the mechanism taking several seconds, or samples of reflectances, to obtain an accurate estimate of the illuminant. As described previously, models such as this suggest the achromatic point should coincide with the chromaticity of the illuminant. In support of this, Figure 41 shows that in the global condition the achromatic point moves, during the trials following the illuminant shift, from a chromaticity closer to the previous illuminant, to a chromaticity closer to the new one.

The plots of C against t (Figure 44 to Figure 53) show that, in most cases, the value of C in the global condition (left-hand plots) reaches a higher value with a greater gradient (faster increase) than in the isolated (middle plots) and inconsistent (right-hand plots). While there is evidence for colour constancy with scenes with multiple regions, each with different illumination conditions (Khang & Zaidi, 2002), it might be expected that more consistent cues would lead to perfect constancy ($C=1$) being achieved faster. The global condition, in which the central test patch and the background patches are all rendered under the same illuminant, provides the most consistent and greater number of surfaces under the same illuminant, concurrently and in a given number of trials. In the isolated condition there is no spatially distributed information about the test illuminant; information about the test-illuminant must be estimated from the series of

chromaticities presented as test patches. In the inconsistent condition, where the background and test patches are rendered under different illuminants, constancy for the test patches can be achieved only if the cues from the background are ignored and only the test patch sequence used to estimate the illuminant, as in the isolated condition.

The migration of the achromatic point over time would indicate that colour constancy increases over time after an illuminant switch and does not happen instantaneously. Since each trial lasted 2.1 seconds, this could mean that, in the cases where perfect constancy is not achieved in the ten trials, in excess of 21 seconds is needed. However, it is not clear how long would be needed in a natural dynamic scene with continuously changing temporal and spatial illuminant conditions, as opposed to our experimental conditions with regular static stimuli. Even in static scenes, it has been demonstrated that colour constancy is better for natural scenes than for artificial ones (Kraft & Brainard, 1999). It is unclear whether the increase in constancy with trial indicates the time required, or the number of samples needed, to achieve constancy.

As discussed previously in section 1.5, there have not been a great many studies of the time-course of colour constancy. However, Foster et al. (2001) present an experiment in which they show that, within a limited spatial extent, a simulated change in the reflectance of one test patch amongst several that coincides with the time of an instantaneous global illuminant change are noticeable and cause the patch to ‘pop-out’. This would suggest that the adjustment to the new illuminant happens instantaneously, to allow the signal for the change in reflectance to coincide with that of the change in illuminant, in contrast to the results of this experiment. Our results also seem to contradict those of Barbur and colleagues who use both psychophysical and imaging experiments. They demonstrate instantaneous colour constancy and argue that it is mediated mostly by early, monocular chromatic mechanisms in V1 with little contribution from later areas, although greater constancy is achieved with binocular stimuli. They also show that mechanisms comparing spatially local contrasts are more important than those integrating spatially distributed information (Barbur & de Cunha, 2003; Barbur, et al., 2004; Barbur & Spang, 2008).

One possibility is that the stimulus configuration in the global condition is insufficient to simulate a real illuminant change. Certainly, we use only flat patches, with Lambertian reflectance. Maloney (1999) has described such stimuli as representing a “flat world”, which lacks many cues to the illumination conditions, like shadows, mutual reflection, specular highlights that characterise a three dimensional world of illuminated objects. This “shape world” (Maloney, 1999) can be simulated reasonably well in computer graphics, particularly in stereoscopic displays (e.g. Yang & Maloney, 2001). It is possible that an illumination change on a “shape world” environment would yield more rapid compensation for the illuminant change. Constancy indices are higher in static simulation of such environments (e.g. Yang & Maloney, 2001; Yang & Shevell,

2002). The rich array of cues to the illuminant contained in real-world stimuli typically produce even higher constancy (Kraft & Brainard, 1999) and it is possible that such stimuli would produce more immediate compensation to an illuminant switch.

Despite these limitations, it is important to note that the “illuminant change” described by Foster and colleagues and by Barbur and colleagues is very similar to the one we use here: a simulated illuminant falling on a variegated array of patches of simulated reflectance. To account for the difference in our results, we must consider the characteristics of studies that imply that colour constancy is instantaneous, and discuss how our task differs from these.

In the studies that show instantaneous recognition of an illuminant change, the task depends on identifying a change in cone-ratios that is inconsistent with the prevailing change in cone-ratios. Foster et al. (2001) used a constant spatial arrangement of stimuli, so in fact the task required judgement of an inconsistent change in cone-ratios at a particular point on the retina, when the temporal change at all other points was of equivalent magnitude. More recently, Amano and Foster (2004) have shown that performance on their task is not disrupted by a change in the spatial layout of simulated surfaces. So it is not the local temporal change in cone-ratios that is important, but the consistency of changes across the scene. On the other-hand, observers’ ability to select the patch that undergoes a change in cone-excitation that is inconsistent with the illuminant change is impaired when the transition between illuminants is gradual (Linnell & Foster, 1996). So, given that the change is sufficiently rapid, observers can extract in parallel relative cone-excitation ratios, and can identify an aberrant change, of the kind that would be produced by a change in one (or more) of the surface reflectances in the scene. This is clearly related to the phenomenon of colour constancy – the ability to disentangle a change in illumination from a change in surface reflectance. But, it does not test colour constancy explicitly – the stability of the colour appearance of an object under a change in illumination conditions. The limitations of such tests of “colour constancy” have been discussed by Foster (2003).

In our task, however, observers were asked for a judgement of colour appearance. They were not asked to identify which patch has undergone a reflectance change and was inconsistent from the others, though they may be sensitive to this information.

Our requirement that the observers made a series of successive judgements means that we must consider the effects of response bias. The experiment with “isolated” test patches is closely related to work by Helson and others (e.g. Helson, 1948). The re-centring on the mean of the set that we observe in this condition is captured by Helson’s level-of-reference theory, which postulates that observers maintain an internal standard against which judgements are made, and that this internal standard is simply a running average (weighted geometric mean) of preceding

stimuli to which the observer has been exposed. It is likely that this re-centring on the mean plays a role in all conditions. While it is difficult to disentangle perceptual and cognitive components of this effect, it is likely that both types of process play a part. It is known that observers are unwilling to continually produce one of two responses (e.g. redder) and that they will aim to distribute their responses evenly between the options available to them (e.g. Poulton, 1979). It is also true that low-level mechanisms of chromatic adaptation with long-time constants would normalise to the temporal mean of the set, which is a good predictor of the illuminant chromaticity (see section 1.5.3). However, since the temporal statistics of the sequence of chromaticities to which the observer must respond was equivalent in the three spatial conditions (global, isolated and inconsistent), differences between conditions (see section 7.3.6) cannot be accounted for by a process that relies solely on response bias or local adjustments to the set of test chromaticities. In particular, the higher constancy indices obtained in the global condition compared to the isolated condition imply that an additional mechanism is required to explain performance under a global illuminant change.

Having considered some caveats regarding the psychophysical method employed in this study, it is interesting to ask again whether our results provide evidence against colour constancy being achieved instantaneously. In a scenario in which colour constancy was perfect and the “flat world” stimuli used in this experiment were sufficient to support perfect constancy, the reflectance functions of the simulated materials would be sufficient to describe perceptual differences between them. So, we are asking our observers to judge a constant set of stimuli, that is, the unchanging set of reflectance functions. If this is the case, a change in response bias (i.e. a re-centring on the mean of the set) should have no dependence on time. If this were how the stimuli were perceived, we should expect to see no change in the achromatic point expressed relative to the reflectance functions (i.e. the same surfaces should be classified as red whatever the illuminant), and we should see an instantaneous change in the achromatic point expressed relative to the test chromaticities when the illuminant changed. This is clearly not an accurate account of observers’ performance in the global condition. If we relax the assumption and say that colour constancy is instantaneous in taking an illuminant change into account, but that constancy is not perfect, the prediction is similar: the achromatic point should fall at a location determined in part by the set of reflectances and in part by the illuminant chromaticity (see Brainard, Brunt and Speigle (1997) for discussion of this equivalent illuminant model), but that again the location of the achromatic point should switch immediately with the illuminant change, alternating between the biased achromatic points (e.g. that are measured in prolonged exposure to a single illuminant (crosses in Figure 41)). If response bias played a part in this process, we would have to assume that it acted to track the difference between a set of reflectances that were contaminated with a residual component from the illuminant preceding a switch, and a set

of reflectances that were contaminated with a residual component from the illuminant following a switch. This difference between the sets to be judged could blur the transition from one set to another (as response-bias renormalized to the change). However, the difference in perceived chromaticities under prolonged exposure to a single illuminant is very small (see Figure 50), and represents only a small fraction of the variance of the distribution of chromaticities. So instantaneous but imperfect constancy might produce a rapid change in achromatic point (accounting for most of the difference between illuminants) followed by much smaller and more gradual changes due to response bias.

A further factor that has been shown to be important in studies of colour and lightness constancy is the presence of a stable anchor. Hansen et al. (2007) have shown for example that judgements are often made relative to a spatially abutting background stimulus. In our “inconsistent” condition, the chromatic statistics of the background remain constant throughout the experiment. As in all conditions, the exact spatial arrangement of the background ellipses, and the arrangement of reflectances used to colour them, changes from trial to trial, but the mean chromaticity, and the chromatic variance, remain constant. On average, since the central test patch touches several ellipses, a spatially-averaged measure of chromatic contrast along this border might be expected to be relatively stable, and important in setting the appearance of the test patch. Indeed Shapiro (2008) has shown strong effects of local contrast on lightness and brightness appearance. However, in the inconsistent condition in our experiment, the appearance of the test-patch is not firmly anchored relative to the unchanging surround: the CI changes over time. The observers’ appearance judgements must have been influenced by the history of preceding chromaticities. Here we have a demonstration that the absolute chromaticity of the test patch is influencing appearance, not its contrast with the immediate surround. This is consistent with Zaidi et al.’s (1997) observation that simultaneous contrast effects are reduced with variegated (compared to uniform) surrounds.

The results of our logistic regression analysis, investigating the extent of the chromaticities of the previous test stimuli on the classification of the current one, are mostly inconclusive. We do find that the predictive effect of the current stimulus chromaticity is far greater than that of the previous trials, and that this chromaticity is best defined relative to the decision boundary appropriate to the classification options (i.e. the blue-yellow probability is best predicted by the distance from the fitted blue-yellow boundary, not the red-green boundary, and vice-versa). However, we don’t find consistent patterns in the predictive effects of previous trials, since the statistically significant predictors vary across observers. This type of analysis requires many repeated independent trials. With only 12 observations per condition our case-to-variable ratio is low, and it is likely that we do not have sufficient power for the analysis.

8 Conclusion

8.1 Summary of hue-circle experiments

Four experiments are presented in this thesis. The first three use variations of a stimulus constructed from modulations of chromaticity over time, which can be considered in terms of the locus of chromaticities visited. In many cases this locus is circular, which is why we often refer to the stimulus as the 'hue-circle'. We investigate how latencies in the signal pathways, from individual cone classes to a site at which the signals from all three cone classes are combined, distort this locus and we use these distortions to make inferences about chromatic processing.

In the first experiment, we replicate and expand a study by Stromeyer et al. (1991), and provide further evidence of a relative latency in the pathway carrying signals from S-cones. Our modification allows us to investigate whether this delay occurs in a mechanism or mechanisms that are aligned to a tritan axis in colour space. Our results suggest that the delay exists in mechanisms that isolate S-cone signals. We interpret this finding in the context of recent physiological findings (with particular reference to Tailby, et al., 2008) suggesting that a variety of opponent mechanisms at post-receptoral stages combine S-cone signals with L- and M-cone signals in such a way that they are tuned to a variety of chromatic directions and not aligned only to the tritan axis. If we assume that these diverse S-cone mechanisms contribute to performance on the task in our experiment, then we must conclude that the delay arises early in the signal pathways, before opponent mechanisms are constructed.

In the second experiment we use the hue-circle stimulus to investigate mechanisms that must combine opponent signals from all three cone types and are therefore tuned to chromatic directions that are not aligned to the cardinal axes of colour space. We implement the hue-circle stimulus in a way that re-creates the conditions leading to the discovery of the Peake effect (see section 1.4.1.5). A sequence of chromaticities following a circular chromatic locus at 10Hz appears different depending on the direction of procession around the circle and this can be explained by a latency difference between mechanisms signalling L/M cone exchanges and mechanisms carrying S-cone signals. Due to the latency difference, the circular locus is distorted at a site at which these mechanisms are combined such that later mechanisms with different tuning receive different stimulation after CW and CCW stimulus modulation. Importantly, however, the CW and CCW hue-circle stimuli present equivalent stimulation to the three cone types, and to the cardinal opponent mechanisms. We demonstrate different chromatic threshold changes after adapting to the two modulation directions, and again relate this to physiological findings. We conclude that the differential adaptation to CW and CCW modulations is taking place at a post receptoral,

perhaps cortical site, where delayed signals from the S-cones are combined with an L/M opponent signal.

In the third experiment, we attempt to produce the same type of differential adaptation as before, however we attempt to measure the effects of this adaptation as changes in reaction time for chromatic discriminations. The reaction time differences we measured were small and there was a great deal of noise in the data, but in general the results of adaptation to CW and CCW hue-circle stimuli were consistent with the previous experiment. In the same experiment, we measure adaptation to modulations in the cardinal directions of colour space, that have previously been used to isolate the putative opponent mechanisms.

In these three experiments we tested the predictions of several existing theories of colour vision. The first experiment provides further support for, and another estimate of, the often-claimed delay in the signals from S-cones (see section 1.4.1). This temporal difference in processing and signal transmission is the first instance of a common theme of these experiments: the time-course of colour vision. Our results are consistent with the delay arising early in the signal pathway, given some assumptions about recently discovered physiological mechanisms contributing to performance in our task. In the literature, there is little agreement on the physiological stage at which the latency is inserted into the pathway. If the delay is early, as our results suggest, it must influence all downstream processing, and is therefore important in any work that relies on the timing of stimuli arising in different cone mechanisms. A delay of the order of 10ms may for example be sufficient to bias the results of studies of attentional allocation and saccadic latencies.

This temporal difference provides the means for the adaptation of chromatic mechanisms that are tuned to non-cardinal directions, shown in the second and third experiments. These results are consistent with physiology that suggests that signals from S-cones remain segregated until some stage in the cortex (Lennie, et al., 1990), and that in the cortex there are units that receive input from all three cone classes (DeValois, et al., 2000). Although the reaction time differences were not statistically significant in all predicted cases, they do lend support to the findings of the threshold experiment.

We also show adaptation of the putative opponent mechanisms, in a similar way to Krauskopf et al. (1982), but using RT measures as Parry et al. (2008). Reaction time differences between S-cone signal carrying mechanisms and mechanisms tuned to other directions are already known at this stage (Smithson & Mollon, 2004) and RTs to purely chromatic signals are generally lower than signals with luminance information (McKeefry, et al., 2003).

These experiments highlight that transmission of chromatic information is not instantaneous and further that different chromatic information is transmitted at different speeds. The Peake effect,

or the observation that CW and CCW hue circle stimuli appear different at frequencies around 10Hz, can be taken as an example of how the relative delay between signals in the putative opponent mechanisms can have a perceptual effect. It is unclear, then, if or how the visual system compensates in any way for the time difference between signals arriving at some central site. The situation is further complicated by differential adaptation in the cone-classes, which could diminish or exaggerate the latency differences. So a simple perceptual calibration is unlikely to work.

8.2 Summary of colour constancy experiment

The final experiment in this thesis explores a different aspect of temporal processing in colour vision. The experiment in section 7 investigates the time-course of colour constancy: specifically the degree of colour constancy achieved as a function of the time after a change in illuminant chromaticity. Colour constancy is the perception of surface colours as invariant under different illuminant chromaticities. The visual system must in some way transform the signals originating from the photoreceptors in order to affect colour constancy, and this transformation depends on the chromaticity of the illuminant. We present a simulated stimulus environment with illumination changes at regular intervals. With a classification task, we make measurements that we can use to calculate a measure of the degree to which this transformation has occurred as a function of the time since an illuminant change. We find that immediately after a simulated illuminant change over the whole scene, colour constancy is not perfect, and over the course of several seconds the constancy index increases.

This result continues the theme of non-instantaneous processing in the colour vision system, and is a result that is timely considering that a great deal of research (summarised in section 1.5) has focussed on colour constancy using static stimuli. It also makes some suggestion about the mechanisms involved in achieving colour constancy: the relatively long times taken to adjust to the new illuminant chromaticity are inconsistent with the process being performed by simple adaptation of early chromatic mechanisms, particularly those in the retina. This suggests that models in which the system integrates information over a longer time period are more likely. An influential theory in the experimental psychology literature of the mid-20th century was Helson's level-of-reference account of perceptual judgements. This sort of process, which is suggested to integrate information over appropriately long time periods, might account for some of our results. However, it cannot account for the differences we observe between the three spatial conditions of our experiment. Nor does the operation of such a mechanism negate the strong conclusion that colour constancy in our global condition is not instantaneous. These results do not, however, explain how in everyday vision we usually notice a rapid change in illuminant but do not perceive surface colours changing over several seconds as the visual system re-adjusts.

8.3 Further experiments

The results from these experiments suggest some possible further work:

The first experiment suggests a general method for investigating phase delays between cone-signals and combinations of cone-signals. The method could be applied using other planes of DKL space to investigate delays between other physiologically significant mechanisms. Some other studies have used sinusoidal modulations to infer phase relationships between mechanisms: Sun, Pokorny & Smith (2001) investigated interactions between rod and cone signals, Stromeyer and colleagues have investigated the temporal properties and effects of adaptation on L- and M-cone inputs to the luminance mechanism (e.g. Stromeyer, Cole, & Kronauer, 1987) and Stockman and colleagues found separate fast and slow inputs to the luminance mechanism (Stockman & Plummer, 2005). These methods generally rely on finding a stimulus phase difference that nulls a perceptual difference, rendering the stimuli indiscriminable. The method we use here additionally relies on a failure to discriminate the temporal order, or sense, in which stimuli are visited.

If we construct similar stimuli to those used in 'cardinal axes' condition, with one of the component modulations in a luminance direction we may be able to measure a phase difference and therefore a latency between signals in either of the cardinal opponent mechanisms and a luminance mechanism. We would expect that chromatic signals lag the luminance signals (Smithson & Mollon, 2004). Using a modification equivalent to our 'intermediate axes' condition we might be able to investigate the independence of chromaticity and luminance mechanisms, which would be interesting given the debate over S-cone signal contributions to luminance (e.g. Chatterjee & Callaway, 2002; Stockman, et al., 1987; Sun, et al., 2006a). Also, by constructing stimuli from independent modulations of L- and M-cone signals (rather than in a plane of DKL space) and maintaining S-cone signals constant, we might be able to obtain measurements of latencies between L- and M-cone signals (Tsujimura, et al., 2000). All of these experiments would depend on the latencies we seek having a similar perceptual effect as the Peake effect: CW and CCW stimuli must be distinguishable at stimulus phase differences at which any phase difference resulting from inherent latency is not nulled. Initial observations suggest that this is indeed the case. CW and CCW modulations around a circular locus in a plane defined by independent L- and M-cone modulations are perceptually different, as are analogous modulations in a plane defined by a $L/(L+M)$ vs. luminance plane.

We show in the first experiment that we can make measurements of the latency of S-cone signals relative to L/M opponent signals for an individual observer. We could apply these measurements to our second experiment. We calculate that the latency of the putative S-opponent mechanism leads to a phase delay between our component modulations and this results in elliptical chromatic loci at a central site. We could use this in conjunction with the methods of our second

experiment to provide insight into the number and tuning of higher-order mechanisms, which is the subject of debate (for review, see Eskew Jr, 2009). By knowing the phase difference caused by latencies, we could calculate stimulus phase differences designed to create linear loci aligned to either the positive or negative diagonal in the equiluminant plane (rather than CW and CCW elliptical loci), corrected for the latency difference. Adaptation to these stimuli should cause more specific desensitisation effects than uncorrected modulations along the diagonal directions, if we assume a model in which there are many higher-order mechanisms sensitive to different chromatic directions and with narrow tuning. Uncorrected modulations in the diagonal colour space directions would result in elliptical loci at a central site (due to the inherent latency difference), meaning that adaptation effects would be distributed over several mechanisms.

In the second experiment presented in this thesis, we demonstrate adaptation of mechanisms tuned to non-cardinal chromatic directions. Observers viewed adapting stimuli for two minutes at the start of each session, before threshold measurements began, and then for four seconds in between threshold measurements. These stimulus durations were chosen after some preliminary work, but it is possible that shorter adaptation durations would produce the same magnitude of adaptation, or that longer durations would produce more adaptation. Given the amount of time required to collect data in our experiment, we were limited in the number of observers we could test. If shorter adaptation stimuli are sufficient, we could increase the number of observers and look for individual differences. Also, we might be able to measure the amount of adaptation of these mechanisms as a function of stimulation duration. This would continue the theme of the time course of chromatic processing in this thesis, and possibly reveal details of the underlying mechanism that we adapt. Different time courses have previously been used to dissociate different underlying mechanisms (Mollon, Stockman, & Polden, 1987; Pugh & Mollon, 1979; Stockman & Smithson, 2006).

There are several modifications that could be made to the final experiment presented in this thesis that would allow us to investigate aspects of colour constancy. As mentioned before, varying the stimulus presentation rate will allow us to distinguish whether the degree of colour constancy is dependent on the number of samples or on time. If the constancy index increases with the number of stimuli presented (under a constant illuminant) then increasing the presentation rate would mean that the same constancy index would be achieved in a shorter time. If we find that constancy is dependent on the number of samples seen, this might suggest that the process is achieved by a mechanism maintaining an internal standard based on some running average of the chromaticities seen, and that this can be better updated with more samples. By the same reasoning we might also expect to see constancy achieved faster with more background ellipses of different reflectances in each stimulus. If the degree of constancy depends

only on the amount of time in which stimuli are seen under consistent illumination condition, a low-level chromatic mechanism might be implicated.

If we adjust the proportion of test stimuli to which the observer is required to respond, we can investigate whether the relevance to the task has an effect on the degree to which stimulus chromaticity affects classification as suggested by Smithson & Zaidi (2004).

Other unanswered questions about colour constancy include the degree to which the visual system can maintain colour constancy in a scene comprising of multiple regions with different illuminants. Our stimuli could easily be adapted to simulate this, and we have pilot data that suggest that two regions, differing in illuminant, can be monitored independently. This is interesting, since natural scenes often include multiple regions with different illuminations because of shadows, secondary reflections from objects within the scene, or clouds creating regions of different amounts of sunlight and skylight illumination. In these situations, assuming a model in which the visual system must maintain a standard or level of reference by which to judge each stimulus, then a separate standard would need to be maintained for each region. Morgan (1992) showed results consistent with multiple standards, each in different spatial locations, and Morgan, Watamaniuk and McKee (2000) showed that in a line-separation discrimination task, an explicitly given standard was not necessary and observers could maintain a running-average over as many as 20 trials. If similar capabilities exist in judgements of colour appearance, then we might expect that stable constancy could be achieved in several scene regions without explicit cues to the illuminant chromaticities.

Rendering more complex simulated environments, including three dimensional ones presented stereoscopically, would allow us to investigate environmental cues such as shadows and specularities that might increase the rate at which constancy can be achieved. In our experiment, colour constancy was not achieved instantaneously after an illuminant switch, whereas in other experiments, cues similar to ours were sufficient (e.g. Foster, et al., 2001) for observers to dissociate change in illumination and in reflectance. We may find that rendering more complex scenes, that might include stronger cues to the illuminant, leads to judgements of chromaticity being updated to take account of the new illuminant immediately.

9 References

- Amano, K., & Foster, D. H. (2004). Colour constancy under simultaneous changes in surface position and illuminant. *Proceedings of the Royal Society of London Series B: Biological Sciences*, 271(1555), 2319-2326.
- Annan, V., & Gilchrist, A. (2004). Lightness depends on immediately prior experience. *Perception & Psychophysics*, 66(6), 943-952.
- Arend, L. E., & Reeves, A. (1986). Simultaneous color constancy. *Journal of the Optical Society of America A: Optics Image Science and Vision*, 3(10), 1743-1751.
- Arend, L. E., Reeves, A., Schirillo, J., & Goldstein, R. (1991). Simultaneous color constancy - papers with diverse munsell values. *Journal of the Optical Society of America A: Optics Image Science and Vision*, 8(4), 661-672.
- Barbur, J. L., & de Cunha, D. (2003). The location of mechanisms for instantaneous colour constancy in human vision. *ARVO Annual Meeting Abstract Search and Program Planner*, 2003, Abstract No. 4202.
- Barbur, J. L., de Cunha, D., Williams, C. B., & Plant, G. (2004). Study of instantaneous color constancy mechanisms in human vision. *Journal of Electronic Imaging*, 13(1), 15-28.
- Barbur, J. L., & Spang, K. (2008). Colour constancy and conscious perception of changes of illuminant. *Neuropsychologia*, 46(3), 853-863.
- Barlow, H. B., & Hill, R. M. (1963). Evidence for a physiological explanation of waterfall phenomenon and figural after-effects. *Nature*, 200(491), 1345-&.
- Benucci, A., Ringach, D. L., & Carandini, M. (2009). Coding of stimulus sequences by population responses in visual cortex. *Nature Neuroscience*, 12(10), 1317-U1145.
- Berson, D. M. (2007). Phototransduction in ganglion-cell photoreceptors. *Pflugers Archiv-European Journal of Physiology*, 454(5), 849-855.
- Blake, Z., Land, T., & Mollon, J. (2008). Relative latencies of cone signals measured by a moving vernier task. *Journal of Vision*, 8(16), 11.
- Blakemore, C., & Campbell, F. W. (1969). On existence of neurones in human visual system selectively sensitive to orientation and size of retinal images. *Journal of Physiology-London*, 203(1), 237-&.
- Bompas, A., & Sumner, P. (2008). Sensory sluggishness dissociates saccadic, manual, and perceptual responses: An S-cone study. *Journal of Vision*, 8(8).
- Boycott, B. B., & Wässle, H. (1991). Morphological classification of bipolar cells of the primate retina. *European Journal of Neuroscience*, 3(11), 1069-1088.
- Boynton, R. M., Ikeda, M., & Stiles, W. S. (1964). Interactions among chromatic mechanisms as inferred from positive and negative increment thresholds. *Vision Research*, 4(1-2), 87-117.
- Brainard, D. H. (1998). Color constancy in the nearly natural image. 2. Achromatic loci. *Journal of the Optical Society of America A: Optics Image Science and Vision*, 15(2), 307-325.
- Brainard, D. H., Brunt, W. A., & Speigle, J. M. (1997). Color constancy in the nearly natural image .1. Asymmetric matches. *Journal of the Optical Society of America A: Optics Image Science and Vision*, 14(9), 2091-2110.
- Brainard, D. H., & Wandell, B. A. (1992). Asymmetric color matching - how color appearance depends on the illuminant. *Journal of the Optical Society of America A: Optics Image Science and Vision*, 9(9), 1433-1448.
- Brenner, E., & Cornelissen, F. W. (2005). A way of selectively degrading colour constancy demonstrates the experience dependence of colour vision. *Current Biology*, 15(21), R864-R866.
- Brill, M. H. (1995). The relation between the color of the illuminant and the color of the illuminated object - commentary. *Color Research and Application*, 20(1), 70-72.
- Brindley, G. S., Ducroz, J. J., & Rushton, W. A. H. (1966). The flicker fusion frequency of the blue-sensitive mechanism of colour vision. *Journal of Physiology-London*, 183(2), 497-&.
- Broackes, J. (1992). The autonomy of color. In K. Lennon & D. Charles (Eds.), *Reduction, explanation and realism* (pp. 421-465): Oxford University Press.

- Buchsbaum, G. (1980). A spatial processor model for object color-perception. *Journal of the Franklin Institute-Engineering and Applied Mathematics*, 310(1), 1-26.
- Buchsbaum, G., & Gottschalk, A. (1983). Trichromacy, Opponent Colors Coding and Optimum Color Information-Transmission in the Retina. *Proceedings of the Royal Society of London Series B: Biological Sciences*, 220(1218), 89-113.
- Buck, S. L. (2004). Rod-cone interactions. In L. M. Chalupa & J. S. Werner (Eds.), *The Visual Neurosciences* (pp. 863-878). Boston: MIT Press.
- Calkins, D. J., & Sterling, P. (1999). Evidence that circuits for spatial and color vision segregate at the first retinal synapse. *Neuron*, 24(2), 313-321.
- Cao, D. C., Zele, A. J., & Pokorny, J. (2005, Jul). *Dark-adapted rod suppression of cone flicker detection: Evaluation of receptor and postreceptor interactions*. Paper presented at the 18th Biennial Meeting of the International-Colour-Vision-Society, Lyon, FRANCE.
- Cardinal, R. N., & Aitken, M. R. F. (2006). *ANOVA For The Behavioral Sciences Researcher* Mahwah, NJ: Lawrence Erlbaum Associates.
- Cataliotti, J., & Bonato, F. (2003). Spatial and temporal lightness anchoring. *Visual Cognition*, 10(5), 621-635.
- Cataliotti, J., & Gilchrist, A. (1995). Local and global processes in surface lightness perception. *Perception & Psychophysics*, 57(2), 125-135.
- Cavanagh, P., Macleod, D. I. A., & Anstis, S. M. (1987). Equiluminance - Spatial and Temporal Factors and the Contribution of Blue-Sensitive Cones. *Journal of the Optical Society of America A: Optics, Image Science & Vision*, 4(8), 1428-1438.
- Chatterjee, S., & Callaway, E. M. (2002). S cone contributions to the magnocellular visual pathway in macaque monkey. *Neuron*, 35(6), 1135-1146.
- Chatterjee, S., & Callaway, E. M. (2003). Parallel colour-opponent pathways to primary visual cortex. *Nature*, 426(6967), 668-671.
- Chen, C. C., Foley, J. M., & Brainard, D. H. (2000). Detection of chromoluminance patterns on chromoluminance pedestals II: model. *Vision Research*, 40(7), 789-803.
- Chichilnisky, E. J., & Baylor, D. A. (1999). Receptive-field microstructure of blue-yellow ganglion cells in primate retina. *Nature Neuroscience*, 2(10), 889-893.
- Chittka, L., Shmida, A., Troje, N., & Menzel, R. (1994). Ultraviolet as a component of flower reflections, and the color-perception of hymenoptera. *Vision Research*, 34(11), 1489-1508.
- CIE. (1932). *Commission internationale de l'Eclairage proceedings, 1931*. Cambridge: Cambridge University Press.
- Conway, B. R. (2001). Spatial structure of cone inputs to color cells in alert macaque primary visual cortex (V-1). *Journal of Neuroscience*, 21(8), 2768-2783.
- Conway, B. R. (2009). Color Vision, Cones, and Color-Coding in the Cortex. *Neuroscientist*, 15(3), 274-290.
- Conway, B. R., Moeller, S., & Tsao, D. Y. (2007). Specialized color modules in macaque extrastriate cortex. *Neuron*, 56, 560-573.
- Cottaris, N. P., & De Valois, R. L. (1998). Temporal dynamics of chromatic tuning in macaque primary visual cortex. *Nature*, 395(6705), 896-900.
- Craven, B. J., & Foster, D. H. (1992). An operational approach to color constancy. *Vision Research*, 32(7), 1359-1366.
- Crawford, B. H. (1947). Visual adaptation in relation to brief conditioning stimuli. *Proceedings of the Royal Society of London Series B: Biological Sciences*, 134(875), 283-302.
- Curcio, C. A., Allen, K. A., Sloan, K. R., Lerea, C. L., Hurley, J. B., Klock, I. B., et al. (1991). Distribution and morphology of human cone photoreceptors stained with anti-blue opsin. *Journal of Comparative Neurology*, 312(4), 610-624.
- Curcio, C. A., Sloan, K. R., Kalina, R. E., & Hendrickson, A. E. (1990). Human Photoreceptor Topography. *Journal of Comparative Neurology*, 292(4), 497-523.
- D'Zmura, M. (1992). Color constancy: surface color from changing illumination. *Journal of the Optical Society of America A: Optics Image Science and Vision*, 9, 490-493.

- D'Zmura, M., & Iverson, G. (1993a). Color constancy. I. Basic theory of 2-stage linear recovery of spectral descriptions for lights and surfaces. *Journal of the Optical Society of America A: Optics Image Science and Vision*, 10(10), 2148-2163.
- D'Zmura, M., & Iverson, G. (1993b). Color constancy. II. Results for 2-stage linear recovery of spectral descriptions for lights and surfaces. *Journal of the Optical Society of America A: Optics Image Science and Vision*, 10(10), 2166-2180.
- D'Zmura, M., & Lennie, P. (1986). Mechanisms of color constancy. *Journal of the Optical Society of America A: Optics Image Science and Vision*, 3(10), 1662-1672.
- Dacey, D. M. (1993). The mosaic of midget ganglion-cells in the human retina. *Journal of Neuroscience*, 13(12), 5334-5355.
- Dacey, D. M., & Lee, B. B. (1994). The blue-on opponent pathway in primate retina originates from a distinct bistratified ganglion-cell type. *Nature*, 367(6465), 731-735.
- Dacey, D. M., Lee, B. B., Stafford, D. K., Pokorny, J., & Smith, V. C. (1996). Horizontal cells of the primate retina: Cone specificity without spectral opponency. *Science*, 271(5249), 656-659.
- Dacey, D. M., Liao, H. W., Peterson, B. B., Robinson, F. R., Smith, V. C., Pokorny, J., et al. (2005). Melanopsin-expressing ganglion cells in primate retina signal colour and irradiance and project to the LGN. *Nature*, 433(7027), 749-754.
- Dacey, D. M., & Packer, O. S. (2003). Colour coding in the primate retina: diverse cell types and cone-specific circuitry. *Current Opinion in Neurobiology*, 13(4), 421-427.
- Dacey, D. M., Packer, O. S., Diller, L., Brainard, D., Peterson, B., & Lee, B. B. (2000). Center surround receptive field structure of cone bipolar cells in primate retina. *Vision Research*, 40(14), 1801-1811.
- Dacey, D. M., Peterson, B. B., & Robinson, F. R. (2002). Identification of an S-cone opponent OFF pathway in the macaque monkey retina: Morphology, physiology and possible circuitry. *Investigative Ophthalmology & Visual Science*, 43, U840-U840.
- Dacey, D. M., Peterson, B. B., Robinson, F. R., & Gamlin, P. D. (2003). Fireworks in the primate retina: In vitro photodynamics reveals diverse LGN-projecting ganglion cell types. *Neuron*, 37(1), 15-27.
- Dannemiller, J. L. (1993). Simultaneous color constancy revisited - an analysis of viewing strategies. *Vision Research*, 33(1), 131-140.
- Dartnall, H. J. A., Bowmaker, J. K., & Mollon, J. D. (1983). Human visual pigments - microspectrophotometric results from the eyes of 7 persons. *Proceedings of the Royal Society of London Series B: Biological Sciences*, 220(1218), 115-130.
- Davidoff, J. B., Aspinall, P. A., & Hill, A. R. (1978). A new colour flicker phenomenon. *Mod Probl Ophthalmol*, 19, 187-188.
- Deeb, S. S. (2005). The molecular basis of variation in human color vision. *Clinical Genetics*, 67(5), 369-377.
- Derrington, A. M., Krauskopf, J., & Lennie, P. (1984). Chromatic mechanisms in lateral geniculate nucleus of macaque. *Journal of Physiology-London*, 357(Dec), 241-265.
- DeValois, R. L., Abramov, I., & Jacobs, G. H. (1966). Analysis of Response Patterns of LGN Cells. *Journal of the Optical Society of America*, 56(7), 966.
- DeValois, R. L., Cottaris, N. P., Elfar, S. D., Mahon, L. E., & Wilson, J. A. (2000). Some transformations of color information from lateral geniculate nucleus to striate cortex. *Proceedings of the National Academy of Sciences of the United States of America*, 97(9), 4997-5002.
- Eskew Jr, R. T. (2009). Higher order color mechanisms: A critical review. *Vision Research*, 49(22), 2686-2704.
- Fairchild, M. D., & Lennie, P. (1992). Chromatic adaptation to natural and incandescent illuminants. *Vision Research*, 32(11), 2077-2085.
- Field, A. (2005). *Discovering Statistics Using SPSS* (2 ed.). London: SAGE Publications.
- Finlayson, G. D., Hordley, S. D., & Hubel, P. M. (2001). Color by correlation: a simple, unifying framework for colour constancy. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 23, 1209-1221.
- Foster, D. H. (2003). Does colour constancy exist? *Trends in Cognitive Sciences*, 7(10), 439-443.

- Foster, D. H., & Nascimento, S. M. C. (1994). Relational color constancy from invariant cone-excitation ratios. *Proceedings of the Royal Society of London Series B: Biological Sciences*, 257(1349), 115-121.
- Foster, D. H., Nascimento, S. M. C., Amano, K., Arend, L., Linnell, K. J., Nieves, J. L., et al. (2001). Parallel detection of violations of color constancy. *Proceedings of the National Academy of Sciences of the United States of America*, 98(14), 8151-8156.
- Gilchrist, A., Kossyfidis, C., Bonato, F., Agostini, T., Cataliotti, J., Li, X. J., et al. (1999). An anchoring theory of lightness perception. *Psychological Review*, 106(4), 795-834.
- Giulianini, F., & Eskew, R. T. (2007). Theory of chromatic noise masking applied to testing linearity of S-cone detection mechanisms. *Journal of the Optical Society of America A: Optics Image Science and Vision*, 24, 2604-2621.
- Golz, J., & MacLeod, D. I. A. (2002). Influence of scene statistics on colour constancy. *Nature*, 415(6872), 637-640.
- Hadjikhani, N., Liu, A. K., Dale, A. M., Cavanagh, P., & Tootell, R. B. H. (1998). Retinotopy and color sensitivity in human visual cortical area V8. *Nature Neuroscience*, 1(3), 235-241.
- Hammond, B. R., Johnson, E. J., Russell, R. M., Krinsky, N. I., Yeum, K. J., Edwards, R. B., et al. (1997). Dietary modification of human macular pigment density. *Investigative Ophthalmology & Visual Science*, 38(9), 1795-1801.
- Hansen, T., Walter, S., & Gegenfurtner, K. R. (2007). Effects of spatial and temporal context on color categories and color constancy. *Journal of Vision*, 7(4).
- Hartridge, H. (1949). *Colours and how we see them*. London: Bell.
- He, S., & MacLeod, D. I. A. (1998). Local nonlinearity in S-cones and their estimated light-collecting apertures. *Vision Research*, 38(7), 1001-1006.
- Helson, H. (1947). Adaptation-level as frame of reference for prediction of psychophysical data. *American Journal of Psychology*, 60(1), 1-29.
- Helson, H. (1948). Adaptation-level as a basis for a quantitative theory of frames of reference. *Psychological Review*, 55(6), 297-313.
- Hendry, S. H. C., & Reid, R. C. (2000). The koniocellular pathway in primate vision. *Annual Review of Neuroscience*, 23, 127-153.
- Hering, E., Hurvich, L. M., & Jameson, D. (1964). *Outlines of a theory of the light sense*. Cambridge, Mass: Harvard University Press.
- Hill, A. R., Rodger, G., & Smalridge, L. (1980). Some further observations on the colour flicker phenomenon. In G. Verriest (Ed.), *Colour vision deficiencies V*. Bristol: Adam Hilger.
- Hiltunen, J. (1996). Munsell colors matt (spectrophotometer measurements by Hiltunen). Retrieved September 10, 1999: <http://www.it.lut.fi/ip/research/color/database/download.html>
- Hofer, H., Carroll, J., Neitz, J., Neitz, M., & Williams, D. R. (2005). Organization of the human trichromatic cone mosaic. *Journal of Neuroscience*, 25(42), 9669-9679.
- Horwitz, G. D., Chichilnisky, E. J., & Albright, T. D. (2005). Blue-yellow signals are enhanced by spatiotemporal luminance contrast in macaque V1. *Journal of Neurophysiology*, 93(4), 2263-2278.
- Hubel, D. H., & Wiesel, T. N. (1966). Effects of varying stimulus size and color on single lateral geniculate cells in rhesus monkeys. *Proceedings of the National Academy of Sciences of the United States of America*, 55(6), 1345.
- Hurlbert, A. C. (1998). Computational models of color constancy. In V. Walsh & J. Kulikowski (Eds.), *Perceptual constancies: why things look as they do* (pp. 283-322). Cambridge: Cambridge University Press.
- Hurlbert, A. C., Bramwell, D. I., Heywood, C., & Cowey, A. (1998). Discrimination of cone contrast changes as evidence for colour constancy in cerebral achromatopsia. *Experimental Brain Research*, 123(1-2), 136-144.
- Hurvich, L. M., & Jameson, D. (1957). An Opponent-Process Theory of Color-Vision. *Psychological Review*, 64(6), 384-404.
- Ives, H. E. (1912). The relation between the color of the illuminant and the color of the illuminated object. *Transactions of the Illuminating Engineering Society*, 7, 62-72.

- Johnson, E. N., Hawken, M. J., & Shapley, R. (2001). The spatial transformation of color in the primary visual cortex of the macaque monkey. *Nature Neuroscience*, 4(4), 409-416.
- Katz, D. (1935). *The world of colour*. London: K. Paul, Trench, Trubner.
- Kentridge, R. W. (2005). Constancy, illumination and the whiteness of the moon. *Clinical and Experimental Ophthalmology*, 33, 572-573.
- Kentridge, R. W., Heywood, C. A., & Weiskrantz, L. (2007). Color contrast processing in human striate cortex. *Proceedings of the National Academy of Sciences of the United States of America*, 104(38), 15129-15131.
- Khang, B. G., & Zaidi, Q. (2002). Cues and strategies for color constancy: perceptual scission, image junctions and transformational color matching. *Vision Research*, 42(2), 211-226.
- Klug, K., Herr, S., Ngo, I. T., Sterling, P., & Schein, S. (2003). Macaque retina contains an S-cone OFF midget pathway. *Journal of Neuroscience*, 23(30), 9881-9887.
- Kouyama, N., & Marshak, D. W. (1992). Bipolar cells specific for blue cones in the macaque retina. *Journal of Neuroscience*, 12(4), 1233-1252.
- Kraft, J. M., & Brainard, D. H. (1999). Mechanisms of color constancy under nearly natural viewing. *Proceedings of the National Academy of Sciences of the United States of America*, 96(1), 307-312.
- Krauskopf, J., Williams, D. R., & Heeley, D. W. (1982). Cardinal Directions of Color Space. *Vision Research*, 22(9), 1123-1131.
- Krauskopf, J., Williams, D. R., Mandler, M. B., & Brown, A. M. (1986). Higher-Order Color Mechanisms. *Vision Research*, 26(1), 23-32.
- Krauskopf, J., & Zaidi, Q. (1986). Induced desensitization. *Vision Research*, 26(5), 759-762.
- Land, E. H. (1964). Retinex. *American Scientist*, 52(2), 247.
- Land, E. H. (1983). Recent advances in retinex theory and some implications for cortical computations - color-vision and the natural image. *Proceedings of the National Academy of Sciences of the United States of America-Physical Sciences*, 80(16), 5163-5169.
- Land, E. H., Hubel, D. H., Livingstone, M. S., Perry, S. H., & Burns, M. M. (1983). Color-generating interactions across the corpus-callosum. *Nature*, 303(5918), 616-618.
- Land, E. H., & McCann, J. J. (1971). Lightness and retinex theory. *Journal of the Optical Society of America*, 61(1), 1-&.
- Lee, B. B., Martin, P. R., & Valberg, A. (1989). Nonlinear summation of m-cone and l-cone inputs to phasic retinal ganglion-cells of the macaque. *Journal of Neuroscience*, 9(4), 1433-1442.
- Lee, H. C. (1986). Method for computing the scene-illuminant chromaticity from specular highlights. *Journal of the Optical Society of America A: Optics Image Science and Vision*, 3(10), 1694-1699.
- Lee, J., & Stromeyer, C. F. (1989). Contribution of human short-wave cones to luminance and motion detection. *Journal of Physiology-London*, 413, 563-593.
- Lee, S. C. S., Telkes, I., & Grunert, U. (2005). S-cones do not contribute to the OFF-midget pathway in the retina of the marmoset, *Callithrix jacchus*. *European Journal of Neuroscience*, 22(2), 437-447.
- Lennie, P., Haake, P. W., & Williams, D. R. (1991). The design of chromatically opponent receptive fields. In L. M. S. & M. J. A. (Eds.), *Computational models of visual processing* (pp. 71-82). Cambridge, MA: MIT Press.
- Lennie, P., Krauskopf, J., & Sclar, G. (1990). Chromatic Mechanisms In Striate Cortex Of Macaque. *Journal of Neuroscience*, 10(2), 649-669.
- Linnell, K. J., & Foster, D. H. (1996). Dependence of relational colour constancy on the extraction of a transient signal. *Perception*, 25(2), 221-228.
- Livingstone, M. S., & Hubel, D. H. (1984). Anatomy and Physiology of a Color System in the Primate Visual-Cortex. *Journal of Neuroscience*, 4(1), 309-356.
- Ludwig, C. J. H., Gilchrist, I. D., McSorley, E., & Baddeley, R. J. (2005). The temporal impulse response underlying saccadic decisions. *Journal of Neuroscience*, 25(43), 9907-9912.
- MacLeod, D. I. A., & Boynton, R. M. (1979). Chromaticity diagram showing cone excitation by stimuli of equal luminance. *Journal of the Optical Society of America*, 69(8), 1183-1186.

- MacLeod, D. I. A., & He, S. (1993). Visible flicker from invisible patterns. *Nature*, *361*(6409), 256-258.
- Maloney, L. T. (1999). Physics-based models of surface color perception. In K. R. Gegenfurtner & L. T. Sharpe (Eds.), *Color vision: from genes to perception* (pp. 387-416). Cambridge: Cambridge University Press.
- Mariani, A. P. (1984). Bipolar cells in monkey retina selective for the cones likely to be blue-sensitive. *Nature*, *308*(5955), 184-186.
- Marshall, N. J. (2000). Communication and camouflage with the same 'bright' colours in reef fishes. *Philosophical Transactions of the Royal Society of London Series B: Biological Sciences*, *355*(1401), 1243-1248.
- Masland, R. H. (2001). The fundamental plan of the retina. *Nature Neuroscience*, *4*(9), 877-886.
- McCann, J. J., McKee, S. P., & Taylor, T. H. (1976). Quantitative Studies in Retinex Theory - Comparison between Theoretical Predictions and Observer Responses to Color Mondrian Experiments. *Vision Research*, *16*(5), 445-&.
- McCollough, C. (1965). Color adaptation of edge-detectors in the Human visual system. *Science*, *149*, 1115-1116.
- McKeefry, D. J., Parry, N. R. A., & Murray, I. J. (2003). Simple reaction times in color space: The influence of chromaticity, contrast, and cone opponency. *Investigative Ophthalmology & Visual Science*, *44*(5), 2267-2276.
- McKeefry, D. J., & Zeki, S. (1997). The position and topography of the human colour centre as revealed by functional magnetic resonance imaging. *Brain*, *120*, 2229-2242.
- McLellan, J. S., & Eskew, R. T. (2000). ON and OFFS-cone pathways have different long-wave cone inputs. *Vision Research*, *40*(18), 2449-2465.
- Mollon, J. D. (1980). Post-receptoral processes in colour-vision. *Nature*, *283*(5748), 623-624.
- Mollon, J. D. (1989). "Tho' she kneel'd in that place where they grew..." The uses and origins of primate color-vision. *Journal of Experimental Biology*, *146*, 21-&.
- Mollon, J. D. (2003). Origins of modern color science. In S. K. Shevell (Ed.), *The science of color* (2 ed., pp. 1-39). Amsterdam/London: Elsevier/Optical Society of America.
- Mollon, J. D., & Bowmaker, J. K. (1992). The spatial arrangement of cones in the primate fovea. *Nature*, *360*(6405), 677-679.
- Mollon, J. D., & Jordan, G. (1997). On the nature of unique hues. In C. Dickinson, M. I. & C. D. (Eds.), *John Dalton's Colour Vision Legacy* (pp. 391-403). London: Taylor and Francis.
- Mollon, J. D., & Krauskopf, J. (1973). Reaction time as a measure of the temporal response properties of individual colour mechanisms. *Vision Research*, *13*(1), 27-40.
- Mollon, J. D., Regan, B. C., & Bowmaker, J. K. (1998). What is the function of the cone-rich rim of the retina? *Eye*, *12*, 548-552.
- Mollon, J. D., Stockman, A., & Polden, P. G. (1987). Transient tritanopia of a second kind. *Vision Research*, *27*(4), 637-650.
- Moreland, J. D., & Bhatt, P. (1984). Retinal distribution of macular pigment. In G. Verriest (Ed.), *Colour vision deficiencies Vol. VII: Dr W Junk Publishers*.
- Morgan, M. J. (1992). On the scaling of size judgments by orientational cues. *Vision Research*, *32*(8), 1433-1445.
- Morgan, M. J., Watamaniuk, S. N. J., & McKee, S. P. (2000). The use of an implicit standard for measuring discrimination thresholds. *Vision Research*, *40*(17), 2341-2349.
- Murray, I. J., Parry, N. R. A., & McKeefry, D. J. (2002). RTs in colour space; effects of equiluminance and adaptation. *Perception*, *31*, 70-70.
- Murray, I. J., & Plainis, S. (2003). Contrast coding and magno/parvo segregation revealed in reaction time studies. *Vision Research*, *43*(25), 2707-2719.
- Nascimento, S. M. C., Ferreira, F. P., & Foster, D. H. (2002). Statistics of spatial cone-excitation ratios in natural scenes. *Journal of the Optical Society of America A: Optics Image Science and Vision*, *19*(8), 1484-1490.
- Neitz, J., Carroll, J., Yamauchi, Y., Neitz, M., & Williams, D. R. (2002). Color perception is mediated by a plastic neural mechanism that is adjustable in adults. *Neuron*, *35*(4), 783-792.

- Newson, L. J. (1958). Some principles governing changes in the apparent lightness of test surfaces isolated from their normal backgrounds. *Quarterly Journal of Experimental Psychology*, *10*(2), 82-95.
- Nikolić, D., Häusler, S., Singer, W., & Maass, W. (2009). Distributed fading memory for stimulus properties in the primary visual cortex. *PLoS Biology*, *7*(12), e1000260.
- Nissen, M. J., & Pokorny, J. (1977). Wavelength effects on simple reaction-time. *Perception & Psychophysics*, *22*(5), 457-462.
- Parry, N. R. A., Murray, I. J., & McKeefry, D. J. (2008). Reaction time measures of adaptation to chromatic contrast. *Visual Neuroscience*, *25*(3), 405-410.
- Piéron, H. (1931). La sensation chromatique. Données sur la latence propre et l'établissement des sensations de couleur. *Année Psychol.*, *32*, 1-29.
- Plainis, S., & Murray, I. J. (2000). Neurophysiological interpretation of human visual reaction times: effect of contrast, spatial frequency and luminance. *Neuropsychologia*, *38*(12), 1555-1564.
- Pokorny, J., Smithson, H., & Quinlan, J. (2004). Photostimulator allowing independent control of rods and the three cone types. *Visual Neuroscience*, *21*(3), 263-267.
- Poulton, E. C. (1979). Models for biases in judging sensory magnitude. *Psychological Bulletin*, *86*(4), 777-803.
- Pugh, J. E. N., & Mollon, J. D. (1979). A theory of the π_1 and π_3 color mechanisms of stiles. *Vision Research*, *19*(3), 293-312.
- Rabin, J., Switkes, E., Crognale, M., Schneck, M. E., & Adams, A. J. (1994). Visual-evoked potentials in 3-dimensional color space - correlates of spatiochromatic processing. *Vision Research*, *34*(20), 2657-2671.
- Rayleigh, L. (1881). Experiments on colour. *Nature*, *25*, 64-66.
- Reid, R. C., & Shapley, R. M. (1992). Spatial structure of cone inputs to receptive-fields in primate lateral geniculate-nucleus. *Nature*, *356*(6371), 716-718.
- Robson, A. G., & Kulikowski, J. J. (1998). Objective specification of tritanopic confusion lines using visual evoked potentials. *Vision Research*, *38*(21), 3499-3503.
- Rodieck, R. W. (1991). Which cells code for color? In A. Valberg & B. B. Lee (Eds.), *From Pigments to Perception: Advances in Understanding Visual Processes* (pp. 83-93). New York: Plenum Press.
- Roorda, A., & Williams, D. R. (1999). The arrangement of the three cone classes in the living human eye. *Nature*, *397*(6719), 520-522.
- Ruttiger, L., Braun, D. I., Gegenfurtner, K. R., Petersen, D., Schonle, P., & Sharpe, L. T. (1999). Selective color constancy deficits after circumscribed unilateral brain lesions. *Journal of Neuroscience*, *19*(8), 3094-3106.
- Schiller, P. H., & Lee, K. M. (1991). The Role of the Primate Extrastriate Area V4 in Vision. *Science*, *251*(4998), 1251-1253.
- Schnapf, J. L., Nunn, B. J., Meister, M., & Baylor, D. A. (1990). Visual transduction in cones of the monkey *Macaca fascicularis*. *Journal of Physiology-London*, *427*, 681-713.
- Schultz, S., Doerschner, K., & Maloney, L. T. (2006). Color constancy and hue scaling. *Journal of Vision*, *6*(10), 1102-1116.
- Shapiro, A. (2008). Separating color from color contrast. *Journal of Vision*, *8*(1), -.
- Shapiro, A., Pokorny, J., & Smith, V. (1996). Cone-rod receptor spaces with illustrations that use CRT phosphor and light-emitting-diode spectra. *Journal of the Optical Society of America A: Optics, Image Science & Vision*, *13*(12), 2319-2328.
- Sharpe, L. T., Stockman, A., Jagla, W., & Jägle, H. (2005). A luminous efficiency function, $V^*(\lambda)$, for daylight adaptation. *Journal of Vision*, *5*(11), 948-968.
- Shinomori, K., Spillmann, L., & Werner, J. S. (1999). S-cone signals to temporal OFF-channels: asymmetrical connections to postreceptoral chromatic mechanisms. *Vision Research*, *39*(1), 39-49.
- Shinomori, K., & Werner, J. S. (2008). The impulse response of S-cone pathways in detection of increments and decrements. *Visual Neuroscience*, *25*(3), 341-347.

- Smithson, H. E. (2005). Sensory, computational and cognitive components of human colour constancy. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 360(1458), 1329-1346.
- Smithson, H. E., Henning, G. B., MacLeod, D. I. A., & Stockman, A. (2009). The effect of notched noise on flicker detection and discrimination. *Journal of Vision*, 9(5), 1-18.
- Smithson, H. E., & Mollon, J. D. (2004). Is the S-opponent chromatic sub-system sluggish? *Vision Research*, 44(25), 2919-2929.
- Smithson, H. E., & Zaidi, Q. (2004). Colour constancy in context: Roles for local adaptation and levels of reference. *Journal of Vision*, 4(9), 693-710.
- Solomon, S. G., & Lennie, P. (2007). The machinery of colour vision. *Nature Reviews Neuroscience*, 8(4), 276-286.
- Solomon, S. G., Peirce, J. W., & Lennie, P. (2004). The impact of suppressive surrounds on chromatic properties of cortical neurons. *Journal of Neuroscience*, 24(1), 148-160.
- Spillmann, L. (1990). Wahrgenommene Bunttonunterschiede auf gegenläufigen Farbkreiseln. *Farbe und Design*, 49/50, 18-24.
- Spillmann, L., & Neumeyer, C. (1984). Änderung der Farbwahrnehmung bei gegensinniger Abfolge der Farbtöne im Farbkreis. In L. Spillmann & B. R. Wooten (Eds.), *Sensory experience, adaptation and perception: Festschrift for Ivo Kohler*. Hillsdale, N.J.: Lawrence Erlbaum.
- Stockman, A., Langendorfer, M., Smithson, H. E., & Sharpe, L. T. (2006). Human cone light adaptation: From behavioral measurements to molecular mechanisms. *Journal of Vision*, 6(11), 1194-1213.
- Stockman, A., MacLeod, D. I. A., & DePriest, D. D. (1987). An inverted S-cone input to the luminance channel: evidence for two processes in S-cone flicker detection. *Investigative Ophthalmology & Visual Science (suppl.)*, 28.
- Stockman, A., Macleod, D. I. A., & DePriest, D. D. (1991). The temporal properties of the human short-wave photoreceptors and their associated pathways. *Vision Research*, 31(2), 189-208.
- Stockman, A., Macleod, D. I. A., & Lebrun, S. J. (1993). Faster than the eye can see - blue cones respond to rapid flicker. *Journal of the Optical Society of America A: Optics, Image Science & Vision*, 10(6), 1396-1402.
- Stockman, A., & Plummer, D. J. (2005). Long-wavelength adaptation reveals slow, spectrally opponent inputs to the human luminance pathway. *Journal of Vision*, 5(9), 702-716.
- Stockman, A., & Sharpe, L. T. (2000). The spectral sensitivities of the middle- and long-wavelength-sensitive cones derived from measurements in observers of known genotype. *Vision Research*, 40(13), 1711-1737.
- Stockman, A., Sharpe, L. T., & Fach, C. (1999). The spectral sensitivity of the human short-wavelength sensitive cones derived from thresholds and color matches. *Vision Research*, 39(17), 2901-2927.
- Stockman, A., & Smithson, H. E. (2006). Transient tritanopia of a second kind redux: Delayed loss of S-cone sensitivity after long-wavelength field onset is consistent with the sluggish generation of an active photoproduct within the L- and M-cones. *Journal of Vision*, 6(13), 5-5.
- Stoughton, C. M., & Conway, B. R. (2008). Neural basis for unique hues. *Current Biology*, 18(16), R698-R699.
- Stromeyer, C. F., Cole, G. R., & Kronauer, R. E. (1987). Chromatic suppression of cone inputs to the luminance flicker mechanism. *Vision Research*, 27(7), 1113-1137.
- Stromeyer, C. F., Eskew, R. T., Kronauer, R. E., & Spillmann, L. (1991). Temporal Phase Response of the Short-Wave Cone Signal for Color and Luminance. *Vision Research*, 31(5), 787-803.
- Sugita, Y. (2004). Experience in early infancy is indispensable for color perception. *Current Biology*, 14(14), 1267-1271.
- Sun, H., Pokorny, J., & Smith, V. C. (2001). Brightness induction from rods. *Journal of Vision*, 1, 32-41.
- Sun, H., Smithson, H. E., Zaidi, Q., & Lee, B. B. (2006a). Do magnocellular and parvocellular ganglion cells avoid short-wavelength cone input? *Visual Neuroscience*, 23(3-4), 441-446.

- Sun, H., Smithson, H. E., Zaidi, Q., & Lee, B. B. (2006b). Specificity of cone inputs to macaque retinal ganglion cells. *Journal of Neurophysiology*, *95*(2), 837-849.
- Sung, C. H., Davenport, C. M., Hennessey, J. C., Maumenee, I. H., Jacobson, S. G., Heckenlively, J. R., et al. (1991). Rhodopsin Mutations in Autosomal Dominant Retinitis-Pigmentosa. *Proceedings of the National Academy of Sciences of the United States of America*, *88*(15), 6481-6485.
- Swanson, W. H., Ueno, T., Smith, V. C., & Pokorny, J. (1987). Temporal modulation sensitivity and pulse-detection thresholds for chromatic and luminance perturbations. *Journal of the Optical Society of America A: Optics, Image Science and Vision*, *4*(10), 1992-2005.
- Tailby, C., Solomon, S. G., & Lennie, P. (2008). Functional Asymmetries in visual pathways carrying S-Cone signals in macaque. *Journal of Neuroscience*, *28*(15), 4078-4087.
- Taylor, A. H., & Kerr, G. P. (1941). The Distribution of Energy in the Visible Spectrum of Daylight. *Journal of the Optical Society of America*, *31*(1), 3-8.
- Troost, J. M., & Deweert, C. M. M. (1991). Naming versus matching in color constancy. *Perception & Psychophysics*, *50*(6), 591-602.
- Tsujimura, S., Shioiri, S., Hirai, Y., & Yaguchi, H. (2000). Technique to investigate the temporal phase shift between L- and M-cone inputs to the luminance mechanism. *Journal of the Optical Society of America A: Optics, Image Science and Vision*, *17*(5), 846-857.
- Ueno, T., Pokorny, J., & Smith, V. C. (1985). Reaction-times to chromatic stimuli. *Vision Research*, *25*(11), 1623-1627.
- Valberg, A., Lee, B. B., & Tigwell, D. A. (1986). Neurons with strong inhibitory S-cone inputs in the macaque lateral geniculate-nucleus. *Vision Research*, *26*(7), 1061-1064.
- Vassilev, A., Mihaylova, M. S., Racheva, K., Zlatkova, M., & Anderson, R. S. (2003). Spatial summation of S-cone ON and OFF signals: Effects of retinal eccentricity. *Vision Research*, *43*(27), 2875-2884.
- von Kries, J. (1878). Beitrag zur Physiologie der Gesichtsempfindungen. [Transl. Physiology of visual sensations]. In D. L. MacAdam (Ed.), *Sources of Color Science* (pp. 101-108). Cambridge, MA: MIT Press.
- Vrhel, M. J., Gershon, R., & Iwan, L. S. (1994). Measurement and analysis of object reflectance spectra. *Color Research and Application*, *19*(1), 4-9.
- Webster, M. A., & Mollon, J. D. (1994). The influence of contrast adaptation on color appearance. *Vision Research*, *34*(15), 1993-2020.
- Webster, M. A., & Mollon, J. D. (1997). Adaptation and the color statistics of natural images. *Vision Research*, *37*(23), 3283-3298.
- Wiesel, T. N., & Hubel, D. H. (1966). Spatial and chromatic interactions in lateral geniculate body of rhesus monkey. *Journal of Neurophysiology*, *29*(6), 1115-&.
- Winderickx, J., Lindsey, D. T., Sanocki, E., Teller, D. Y., Motulsky, A. G., & Deeb, S. S. (1992). Polymorphism in red photopigment underlies variation in color matching. *Nature*, *356*(6368), 431-433.
- Wisowaty, J. J., & Boynton, R. M. (1980). Temporal modulation sensitivity of the blue mechanism: measurements made without chromatic adaptation. *Vision Research*, *20*(11), 895-909.
- Wyszecki, G., & Stiles, W. S. (2000). *Color science : concepts and methods, quantitative data and formulae* (2nd ed.). New York ; Chichester: Wiley.
- Xiao, Y. P., Wang, Y., & Felleman, D. J. (2003). A spatially organized representation of colour in macaque cortical area V2. *Nature*, *421*(6922), 535-539.
- Yang, J. N., & Maloney, L. T. (2001). Illuminant cues in surface color perception: tests of three candidate cues. *Vision Research*, *41*(20), 2581-2600.
- Yang, J. N., & Shevell, S. K. (2002). Stereo disparity improves color constancy. *Vision Research*, *42*(16), 1979-1989.
- Yeh, T., Lee, B. B., & Kremers, J. (1995). Temporal response of ganglion-cells of the macaque retina to cone-specific modulation. *Journal of the Optical Society of America A: Optics, Image Science & Vision*, *12*(3), 456-464.
- Zaidi, Q. (1998). Identification of illuminant and object colors: heuristic-based algorithms. *Journal of the Optical Society of America A: Optics Image Science and Vision*, *15*(7), 1767-1776.

- Zaidi, Q., & Halevy, D. (1993). Visual Mechanisms That Signal the Direction of Color Changes. *Vision Research*, 33(8), 1037-1051.
- Zaidi, Q., & Shapiro, A. G. (1993). Adaptive orthogonalization of opponent-color signals. *Biological Cybernetics*, 69(5-6), 415-428.
- Zaidi, Q., Spehar, B., & DeBonet, J. (1997). Color constancy in variegated scenes: role of low-level mechanisms in discounting illumination changes. *Journal of the Optical Society of America A: Optics, Image Science and Vision*, 14(10), 2608-2621.
- Zeki, S. (1973). Color coding in rhesus-monkey prestriate cortex. *Brain Research*, 53(2), 422-427.
- Zeki, S., Aglioti, S., McKeefry, D., & Berlucchi, G. (1999). The neurological basis of conscious color perception in a blind patient. *Proceedings of the National Academy of Sciences of the United States of America*, 96(24), 14124-14129.
- Zeile, A. J., Cao, D. C., & Pokorny, J. (2007). Threshold units: A correct metric for reaction time? *Vision Research*, 47(5), 608-611.