

# CUE COMBINATION OF COLOUR AND LUMINANCE IN EDGE DETECTION

REBECCA JANE PARKER SHARMAN, MSc.

Thesis submitted to the University of Nottingham  
for the degree of Doctor of Philosophy

SEPTEMBER 2013

## Abstract

Much is known about visual processing of chromatic and luminance information. However, less is known about how these two signals are combined. This thesis has three aims to investigate how colour and luminance are combined in edge detection. 1) To determine whether presenting colour and luminance information together improves performance in tasks such as edge localisation and blur detection. 2) To investigate how the visual system resolves conflicts between colour and luminance edge information. 3) To explore whether colour and luminance edge information is always combined in the same way.

It is well known that the perception of chromatic blur can be constrained by sharp luminance information in natural scenes. The first set of experiments (Chapter 3) quantifies this effect and demonstrates that it cannot be explained by poorer acuity in processing chromatic information, higher contrast of luminance information or differences in the statistical structure of colour and luminance information in natural scenes. It is therefore proposed that there is a neural mechanism that actively promotes luminance information.

Chapter 4 and Experiments 5.1 and 5.3 aimed to investigate whether the presence of both chromatic and luminance information improves edge localisation performance. Participant performance in a Vernier acuity (alignment) task was compared to predictions from three models; 'winner takes all', unweighted averaging and maximum likelihood estimation (a form of weighted averaging). Despite several attempts to differentiate the models

we failed to increase the differences in model predictions sufficiently and it was not possible to determine whether edge localisation was enhanced by the presence of both cues.

In Experiment 5.4 we investigated how edges are localised when colour and luminance cues conflict, using the method of adjustment. Maximum likelihood estimation was used to make predictions based on measurements of each cue in isolation. These predictions were then compared to observed data. It was found that, whilst maximum likelihood estimation captured the *pattern* of the data, it consistently over-estimated the weight of the luminance component. It is suggested that chromatic information may be weighted more heavily than predicted as it is more useful for detecting object boundaries in natural scenes.

In Chapter 6 a novel approach, perturbation discrimination, was used to investigate how the spatial arrangement of chromatic and luminance cues, and the type of chromatic and luminance information, can affect cue combination. Perturbation discrimination requires participants to select the grating stimulus that contains spatial perturbation. If one cue dominated over the other it was expected that this would be reflected by masking and increased perturbation detection thresholds. We compared perturbation thresholds for chromatic and luminance defined line and square-wave gratings in isolation and when presented with a mask of the other channel and other grating type. For example, the perturbation threshold for a luminance line target alone was compared to the threshold for a luminance

line target presented with a chromatic square-wave target. The introduction of line masks caused masking for both combinations. Introduction of an achromatic square-wave mask had no effect on perturbation thresholds for chromatic line targets. However, the introduction of a chromatic square-wave mask to luminance line targets *improved* perturbation discrimination performance. This suggests that the perceived location of the chromatic edges is determined by the location of the luminance lines.

Finally, in Chapter 7, we investigated whether chromatic blur is constrained by luminance information in bipartite edges. Earlier in the thesis we demonstrated that luminance information constrains chromatic blur in *natural scenes*, but also that chromatic information has more influence than expected when colour and luminance edges conflict. This difference may be due to differences in the stimuli or due to differences in the task. The luminance masking effect found using natural scenes was replicated using bipartite edges. Therefore, the finding that luminance constrains chromatic blur is not limited to natural scene stimuli. This suggests that colour and luminance are combined differently for blur discrimination tasks and edge localisation tasks.

Overall we can see that luminance often dominates in edge perception tasks. For blur discrimination this seems to be because the mechanisms differ. For edge localisation it might be simply that luminance cues are often higher contrast and, when this is equated, chromatic cues are actually a good indicator of edge location.

## Acknowledgements

I would like to thank my supervisor Jon Peirce, I am eternally grateful for his help, support and guidance. This thesis would be a poor shadow of itself without his white-board diagrams and 'direct' comments. A thank you also goes to my PhD cohort for group commiseration and congratulation at the pub.

Finally, I would like to thank Bonnie, Pete and my parents, for always believing in me and being so very proud of me, even when they had no idea what I was talking about.

## Contents

|       |   |    |
|-------|---|----|
| 1.    | Introduction .....  | 1  |
| 1.1.  | The colour pathways.....  | 2  |
| 1.2.  | Chromatic and spatial selectivity.....  | 10 |
| 1.3.  | Models of cue combination .....   | 12 |
| 1.4.  | Combining colour and luminance in shape from shading and contrast detection tasks.....        | 22 |
| 1.5.  | The dominance of luminance information in edge detection tasks.....                           | 28 |
| 1.6.  | The current thesis .....  | 31 |
| 2.    | General methods .....   | 36 |
| 2.1.  | Participants .....  | 36 |
| 2.2.  | Apparatus.....  | 36 |
| 2.3.  | Bits ++ Digital Video Processor .....   | 36 |
| 2.4.  | Gamma Correction .....  | 37 |
| 2.5.  | DKL Space.....  | 38 |
| 2.6.  | Generating isoluminant stimuli .....  | 40 |
| 2.7.  | Defining contrast .....   | 42 |
| 2.8.  | Adaptive staircase procedures .....   | 43 |
| 2.9.  | Method of adjustment.....   | 44 |
| 2.10. | Curve fitting .....   | 44 |
| 2.11. | Using the last six reversals as an estimate of threshold.....                                 | 46 |
| 2.12. | Bootstrapping .....   | 49 |
| 3.    | Luminance information constrains chromatic blur discrimination in natural scene stimuli ..... | 51 |
| 3.1.  | Blur discrimination.....  | 53 |
| 3.2.  | Equating effective contrast.....  | 59 |
| 3.3.  | Controlling for statistical regularities .....  | 62 |
| 3.4.  | Ruling out luminance artefacts.....   | 65 |
| 3.5.  | Discussion .....  | 68 |
| 4.    | Cue combination of colour and luminance in aligned synthetic edges.....                       | 74 |
| 4.1.  | The effect of Gaussian white noise.....   | 76 |
| 4.2.  | The effect of increasing viewing distance.....  | 83 |

|      |  |     |
|------|--|-----|
| 4.3. | Discussion .....   | 90  |
| 5.   | Cue combination of conflicting colour and luminance edges .....                            | 92  |
| 5.1. | Piloting method of adjustment.....   | 94  |
| 5.2. | Equating performance using contrast .....  | 97  |
| 5.3. | Cue combination in aligned synthetic edges .....   | 101 |
| 5.4. | Edge localisation with conflicting stimuli .....   | 105 |
| 5.5. | Discussion .....   | 108 |
| 6.   | Detection of perturbation in chromatic and luminance stimuli is modulated by context. .... | 112 |
| 6.1. | Aligned and orthogonal stimuli. ....   | 116 |
| 6.2. | Square-wave and line stimuli.....  | 125 |
| 6.3. | Discussion .....   | 136 |
| 7.   | Luminance information constrains chromatic blur in bipartite edges .....                   | 141 |
| 7.1. | Method .....   | 143 |
| 7.2. | Results.....   | 146 |
| 7.3. | Discussion .....   | 148 |
| 8.   | General Discussion.....  | 151 |
| 8.1. | Summary of findings .....  | 152 |
| 8.2. | Implications of findings.....  | 154 |
| 8.3. | Relationship to previous literature .....  | 156 |
| 8.4. | Unanswered questions .....   | 160 |
| 8.5. | Conclusion.....  | 163 |
| 9.   | References .....   | 164 |

## 1. Introduction

Edge recognition is a fundamental part of visual perception and is necessary to navigate and interact with the environment. We can easily and accurately localise edges under a variety of conditions including when vision is obscured for example, in low light conditions, in a crowded scene or through a rain soaked window. However, the majority of edges in natural scenes are comprised of both colour and luminance information. Colour and luminance information enter the cortex in separate pathways, but we do not see two overlaid percepts, therefore, this information must be combined. Furthermore, in situations when the two cues conflict, this must be resolved to give an unambiguous and accurate percept of the world around us.

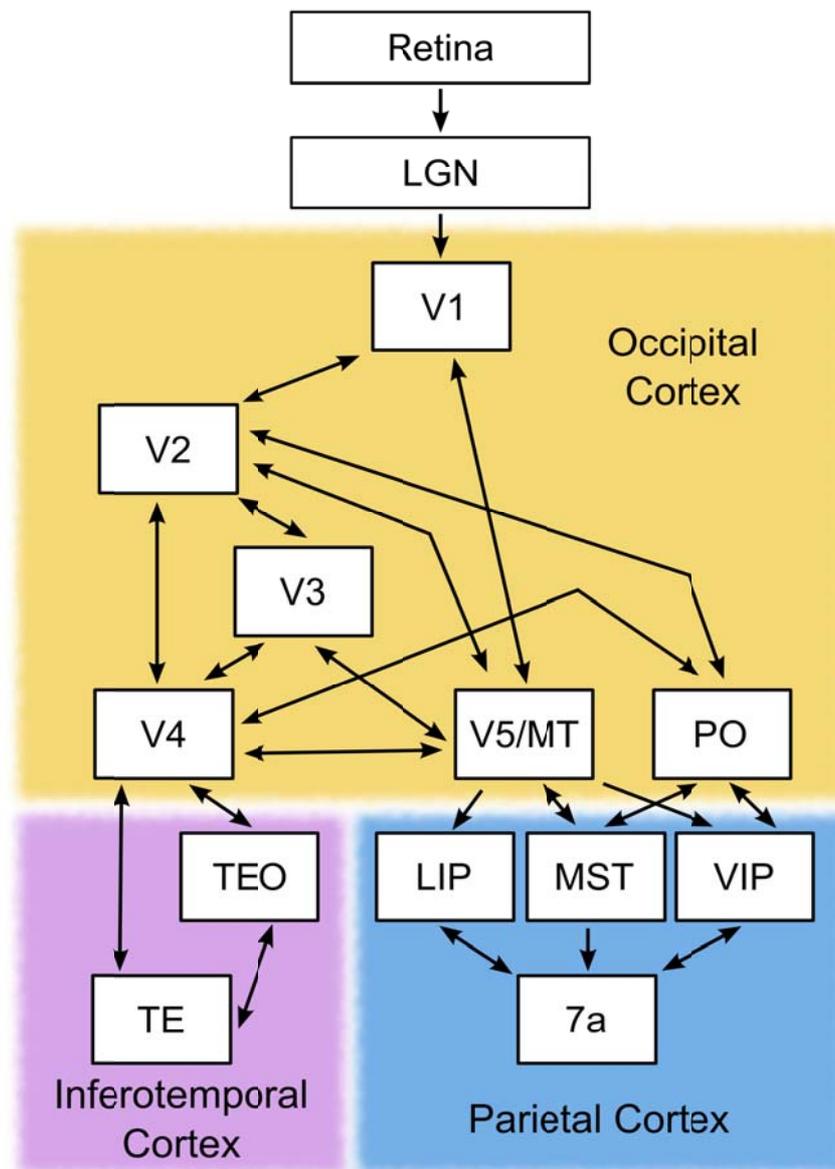
Previously the question of how colour and luminance cues are combined has been considered using behavioural, physiological and neuro-imaging techniques. The research suggests that colour, like luminance, is represented throughout the cortex as evidenced by neurons in the early visual areas that are selective for both colour and other form attributes (orientation, edges). Psychophysical data suggest that colour and luminance are not combined linearly and may be subject to specific priors under specific conditions.

This introduction will review current knowledge of the colour pathways and areas in the visual system, focusing on whether these areas are segregated from those processing luminance information. We will then consider models of cue combination and how they can be applied to the

current research. Following this existing research into how colour and luminance are combined will be reviewed. Finally, the aims of the current thesis will be outlined.

### 1.1. The colour pathways

The perception of colour begins with the absorption of light in the retinal cone photoreceptors; this is converted into electrical voltages that are then converted into action potentials by the cells in the retina. The information from the retinal ganglion cells is then sent to the lateral geniculate nucleus (LGN) and then on to the visual cortex (see Gegenfurtner & Kiper, 2003 for a review of colour vision). See Figure 1.1 for a schematic of the visual hierarchy.



**Figure 1.1.** Schematic of the visual hierarchy.

There are three cone classes; long- (L-), medium- (M-) and short-wavelength sensitive (S-). These are sometimes referred to as red, green and blue although, in reality, all cone types are sensitive to a large range of wavelengths. The signal produced by each of the three cone classes increases either as the light's wavelength becomes closer to the peak of their absorption spectrum (the colour they are most sensitive to), as the intensity

of the light increases and becomes brighter or a combination of both of these things. As a result of this the colour of a stimulus can only be determined by comparing the output signals of the three cone classes. This comparison is performed by the horizontal and ganglion cells in the retina. The axons of the retinal ganglion cells in the optic nerve then pass this information to the LGN (in the thalamus) and the superior colliculus.

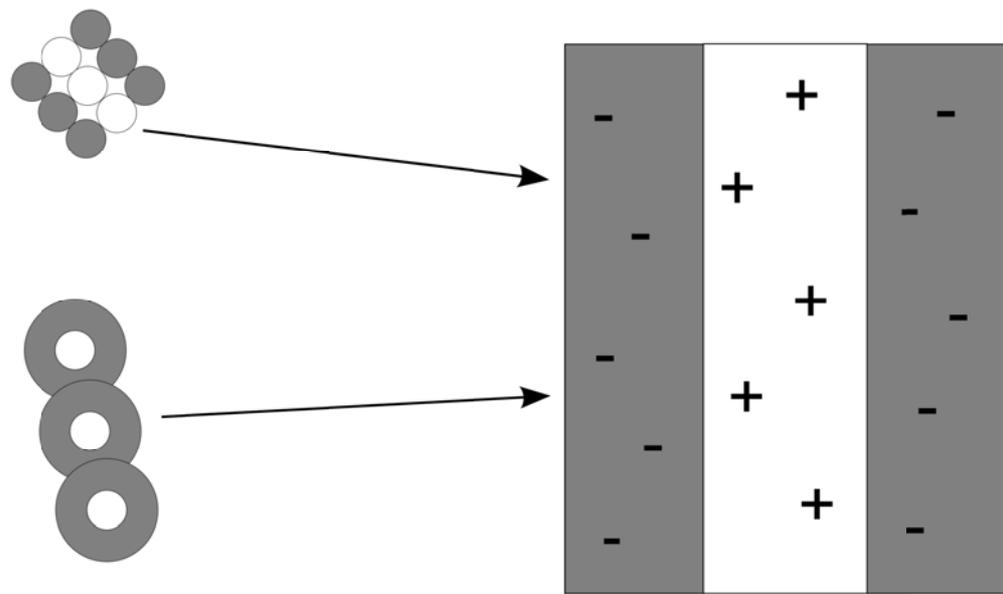
There is very little variation in the peak sensitivities of the three cone classes across all Old World primate species (Jacobs & Deegan, 1999), potentially indicating that they are tuned to the wavelengths of light in the natural environment and that this is an evolutionary adaptation. However, these cone spectral sensitivity curves do not provide the maximum possible amount of colour information. If L-cones were sensitive to even longer wavelengths a significant increase in colour information could be achieved, but this would be at the expense of spatial acuity, due to increased diffraction and chromatic aberration (Lewis & Zhaoping, 2006). It is therefore likely that human cone sensitivities represent a compromise between maximising colour information and maximising spatial information.

There are three major subclass of retinal ganglion cell; parasol, midget and bistratified. These cells project via three pathways to separate layers in the LGN, each of which contains a different type of cell, respectively; magnocellular (M-), parvocellular (P-) and koniocellular (K-) cells. These cells have a classic centre-surround organisation and can be described as either ON-centre or OFF-centre, the firing rate of an ON-centre cell increases when

light hits the centre and decreases when light hits the surround and OFF-centre is the reverse. The areas that fire in response to light are the cells' receptive fields. A receptive field can be defined as the region in visual space where a specific type of visual stimulus can drive electrical responses.

The geniculate M-cells are sensitive to luminance stimuli but exhibit null responses to some L-M combinations (Shapley & Hawken, 2002). For this reason the M-pathway is considered to largely convey achromatic (L+M) information about motion. The P-pathway receives inputs from L- and M-cones and is considered to convey information about colour and edges. P-cells respond well to high-spatial frequency achromatic gratings and low-spatial frequency chromatic gratings (Ingling & Martinezuriegas, 1983). The K-pathway predominantly carries signals from the S-cones and appears to convey information about colour (Casagrande, 1994). The three pathways terminate in different layers of V1 (Callaway, 1998).

In the cortex, LGN luminance inputs are combined to create simple cells with orientation tuned receptive fields (Figure 1.2). Simple cells are phase sensitive and their response changes dependent on where a stimulus bar is in their receptive field.



**Figure 1.2.** Combining LGN single opponent cells to create a V1 orientation tuned simple cell. The illustrated receptive field will respond best to a vertical bright bar on a dark background.

Simple cells with the same orientation tuning but different spatial positions can be combined to form complex cells; these are phase insensitive. Complex cells respond to a bar of their preferred orientation at any position in their receptive field. There is strong evidence that luminance signals are combined in this way, however, it is not clear if the same is true for chromatic signals. It has been proposed that there are double-opponent cells that create oriented receptive fields for chromatic stimuli, in the same way as described above. These cells would respond to chromatic patterns, but not full field or low-spatial frequency stimuli (see Shapley & Hawken, 2011 for a review).

Within both V1 (E. N. Johnson, Hawken, & Shapley, 2001) and V2 (Gegenfurtner, Kiper, & Fenstemaker, 1996) about 50% of the cells populations are selective for colour. A majority of these cells sum their inputs

linearly and linear combination principles can be used to model them (Lennie, Krauskopf, & Sclar, 1990). In V2 a greater percentage of cells are more selective i.e. respond to a narrower range of colours, than would be predicted by a linear combination model and a non-linear stage is necessary (Kiper, Fenstemaker, & Gegenfurtner, 1997).

There are, however, V1 cells that do not combine cone signals linearly. Horwitz and colleague (2005) excited V1 neurons in awake macaques with dynamic randomly coloured stimuli and analysed the stimulus sequences that preceded spikes in two steps. First they computed the average stimulus that preceded the spike, identifying a group of S-cone dominated colour opponent neurons. If these neurons had combined cone signals linearly this would have characterised the colour tuning. However, the second stage of their analysis showed that approximately half of neurons received a rectified non-opponent signal from the L- and M-cones that was combined with the opponent signal. The result is a receptive field structure that might respond to both a luminance edge and the presence of chromatic contrast.

It was initially believed that colour and form were physiologically and functionally segregated within the visual system. Hubel and Livingstone (1987) posited that colour and form sensitive cells were physically separated, with colour sensitive cells confined to the cytochrome oxidase (CO) blobs. However, more recent electrophysiology work has found that there are not discrete regions for orientation/direction and colour and that in the upper layers of V1, neurons that are colour sensitive are orientation selective as

often as those that were not sensitive to colour (Leventhal, Thompson, Liu, Zhou, & Ault, 1995). In V2 there is a tendency for colour selective cells to be found in the thin CO stripes, but they are still frequently found in the thick and inter stripe areas. There is also not a relationship between colour selectivity and selectivity for other attributes in V2 and many cells can encode information along more than one dimension (Gegenfurtner, et al., 1996), this has led to the suggestion that V2 integrates information about different attributes.

Evidence from psychophysical and fMRI studies suggest that, whilst there may be some segregation, the relationship between colour and luminance in the cortex is more complicated than originally thought. Curvature integration functions in the same fashion for achromatic, (L-M)- and S-defined stimuli, however the mechanisms underlying this seem to be separate as integration is disrupted when the components are chromatically different (McIlhagga & Mullen, 1996; Mullen, Beaudot, & McIlhagga, 2000). fMRI work has found selectivity in V1, V2 and V3 for chromatically defined orientation signals and it is possible to discriminate between luminance, L-M and S defined stimuli based on the activity patterns (Sumner, Anderson, Sylvester, Haynes, & Rees, 2008). This also suggests that whilst there are neurons which are jointly selective for orientation and colour there may be some segregation in processing streams if not in physical location.

V4 has been suggested as the 'colour centre' of the monkey brain (Zeki, 1983a, 1983b), however, there may not be a single area responsible for

colour processing. Neurophysiological studies in monkeys show that lesions of V4 lead to mild colour vision deficits but, problematically, they also lead to a variety of other deficits (Schiller, 1993; Walsh, Kulikowski, Butler, & Carden, 1992). In macaques, lesions of the infero-temporal (IT) cortex, the next processing stage, produce effects similar to cerebral achromatopsia (acquired colour-blindness caused by damage to the cortex). However the entire IT area must be removed for this to occur which once again has a dramatic effect on other areas of vision (Heywood, Gaffan, & Cowey, 1995).

There are colour sensitive neurons throughout the early visual cortex with spectral sensitivities no narrower than those found in V4 (de Monasterio & Schein, 1982). Therefore, it may be that V4 is involved in higher order processing of colour information such as colour constancy (Kulikowski, Walsh, McKeefry, Butler, & Carden, 1994) or illuminant discounting (Bartels & Zeki, 2000) with lower level processing of colour occurring throughout earlier areas. In support of this argument, V4 is involved with the ratio-taking operations necessary for illuminant discounting (Bartels & Zeki, 2000).

Evidence from fMRI demonstrates how colour is represented differently throughout the visual system. Brouwer & Heeger (2009) were able to decode stimulus colour from activity in human V1, V2, V3, V4 and VO1 but not LO1, LO2, V3A/B or MT+. They found that in areas V4 and VO1 responses were similar to colour perception; perceptually similar colours evoked similar responses. In contrast, V1 responses appear to be organised according to a cone-opponency model, demonstrating that colour representation changes

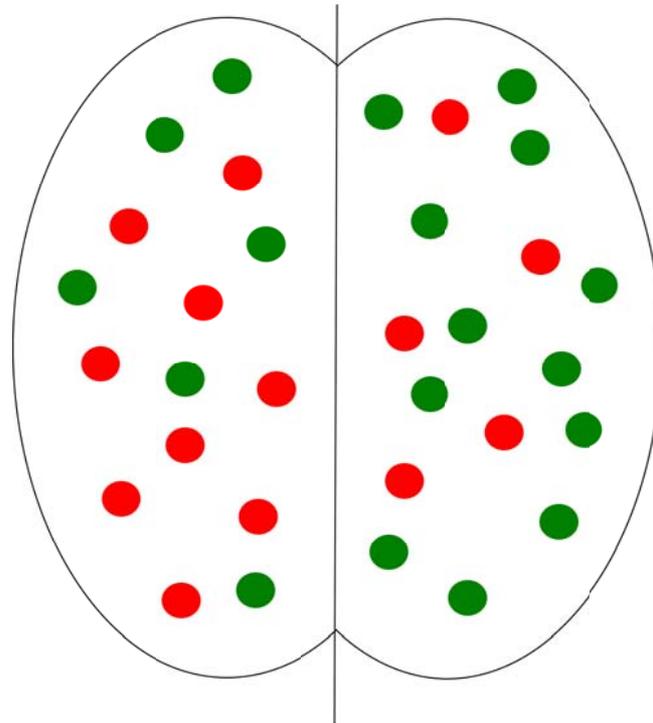
through the visual system; transforming from a cone-opponency pattern into perceptual colour space.

Spatial sensitivities are different for chromatic and luminance stimuli. Chromatic contrast sensitivity is low-pass whereas achromatic contrast sensitivity is band-pass. This means that high-spatial frequency chromatic stimuli are not resolvable by the visual system but high-spatial frequency achromatic stimuli are. There are also differences in chromatic and luminance processing in the periphery. Chromatic sensitivity falls off more quickly than achromatic in the periphery and, more specifically, sensitivity for green stimuli decreases faster than for red (Newton & Eskew, 2003). It has been suggested that there may be proportionally fewer M-cones in the periphery, leading to a weaker 'green' response, however this does not appear to be the case (Newton & Eskew, 2003). Currently it is believed that there are post-receptoral limitations on peripheral colour resolution, potentially based in the double opponent cells of the cortex (Anderson, Mullen, & Hess, 1991).

## 1.2. Chromatic and spatial selectivity

In Figure 1.3 the left side of the receptive field has more L-cones and the right has more M-cones. Let us consider how a neuron might use these cones as potential inputs depending on various different functional roles. If maximum sensitivity to changes in colour is required, then the L-cones should be subtracted from M-cones at all points in the space. This is non-spatially selective and will only give information about the chromaticity. If maximum spatial selectivity is required, the responses from all the cones in the left

receptive field should be subtracted from the right. The imbalance in the cone distribution will result in this neuron having a chromatic preference, in this case a slight preference for changes from red to green across the receptive field. If spatial selectivity based *only* on luminance information is required, then the number of L- and M-cones in each receptive field need to be balanced and so some must be disregarded. This would lead to lower sensitivity, meaning that the inclusion of some chromatic information improves spatial selectivity. Maximum sensitivity to chromatic change and maximum spatial selectivity are mutually exclusive (Peirce, et al., 2008) and so if the task requires both chromatic and spatial information then the receptive field will be less sensitive.



**Figure 1.3.** Schematic representation of the cone mosaic. The central line represents an edge to be detected by the human visual system, the two areas either side represent receptive fields. Subtracting one side from the other to determine edge location would also include some chromatic information.

### 1.3. Models of cue combination

A considerable amount is known about how early visual signals are combined; how non-oriented LGN receptive fields are combined to form oriented V1 receptive fields. However, less is known about how these signals are subsequently combined in extra-striate areas. Colour and luminance signals are the fundamental building blocks of vision and we know that they must be combined to form the whole percept, yet, it is unclear how (or where) this combination takes place and what combinatorial rules may exist.

We can only infer the state of the world based on sensory neural activity which is corrupted by noise, so the brain is forced to create our perceptions of the world under conditions of uncertainty. Perception is an ill-posed 'inverse problem' i.e. an image on the retina can be caused by an infinite number of physical realities, this means that a computational strategy must be employed to allow us to perceive the world unambiguously. We can model potential strategies that the visual system may use to combine colour and luminance and test these models against psychophysical findings to infer the way cue combination may occur.

There are demonstrations that are suggestive of how colour and luminance signals are combined to form edges. The Spanish castle and Boynton illusions (see Section 1.5 for more details, Kaiser, 1996; Sadowski, 2006) both show luminance information appearing to constrain the spread of chromatic information. This suggests that luminance edge information is more important to the visual system for edge localisation and that chromatic edge information is effectively ignored.

This could be explained by a 'winner takes all' cue combination strategy. If this method is employed the 'best' cue is selected and everything else is disregarded. The difficulty comes in determining which cue is 'best'. There could be a general over-arching rule; luminance has a higher effective contrast in natural scenes so is always the 'best' cue. However, it could also be defined on a case-by-case basis; of two cues we might select the one with the least variability over a specified time period.

It is unclear how *conflicting* colour and luminance edges are combined in edge localisation (see Chapter 5). If a ‘winner takes all’ strategy as described above is used, one of the cues may be ignored. However, if both cues contribute to the perceived edge location there are two possible ways they could be combined. Firstly, colour and luminance edge cues could be combined using unweighted averaging. In this case both edges would have an equal contribution to perceived edge location, regardless of the cue ‘quality’, and the edge location would be judged to be equidistance between the individual cues. Alternatively, if observers do not always perceive the edge to be equidistant between the two cues this could suggest that the reliability of the two cues contributes to edge localisation. The brain could be performing some form of ‘maximum likelihood estimation’ (MLE) to estimate the edge location. In an MLE model the variability of the localisation judgements of each individual cue are used to generate weights; the contribution that each cue makes to the localisation judgement. A method of using MLE to generate cue combination predictions has been proposed by Hillis and colleagues (2004) and is described in detail below.

If we have unbiased estimates of edge location based on colour ( $\hat{S}_{Col}$ ) and luminance ( $\hat{S}_{Lum}$ ) cues with variances  $\sigma_{Col}^2$  and  $\sigma_{Lum}^2$  respectively. The way to combine these two estimates to produce a prediction with the minimum variances is

$$\hat{S} = w_{Col}\hat{S}_{Col} + w_{Lum}\hat{S}_{Lum}$$

Equation 1.1

where the weights are

$$w_{Col} = \frac{r_{Col}}{r_{Col} + r_{Lum}} \quad \text{and} \quad w_{Lum} = \frac{r_{Lum}}{r_{Lum} + r_{Col}}$$

**Equation 1.2**

and the reliabilities ( $r_{Col}$  and  $r_{Lum}$ ) are the inverse of the respective variances.

The variance of the weighted average  $\hat{S}$  is lower than that for either of the individual cues and is given by

$$\sigma^2 = \frac{\sigma_{Col}^2 \sigma_{Lum}^2}{\sigma_{Col}^2 + \sigma_{Lum}^2} \quad \text{or} \quad r = r_{Col} + r_{Lum}$$

**Equation 1.3**

Equation 1.1 and Equation 1.3 can produce predictions of where a conflicting edge will be judged to be located and the variance of that judgement, respectively.

There is increasing evidence that human perceptual computations are combined optimally according to Bayes Theorem. Bayes Theorem states that, when we try to determine the presence or absence of a signal, we should combine the perceptual data about whether that signal is present (the *likelihood*) with our previous expectation of whether that signal was going to be present (the *prior*). The prior information might be formed in various ways, but is typically assumed to be generated from the statistical history of the signal events i.e. previously observed occurrences of the signal. The prior information is combined with sensory inputs to produce a *posterior* probability distribution. The mean and variance of this distribution represent

what the most probable stimulus is and the probability that it is present respectively, and could represent the unambiguous percept that we actually see. The weight given to priors is likely to depend on the degree of ambiguity in the sensory inputs; in conditions of high perceptual uncertainty more weight is likely to be given to the prior information than the sensory inputs (Kersten & Yuille, 2003).

In order to use this method the visual system needs to have an internal representation of current uncertainty that is always available and changes in response to new information. Psychophysical studies have shown that humans use continuous feedback from the hand to control pointing movements and the relative weights of the different signals are dependent on the expected sensory noise associated with those signals, as would be predicted by Bayesian theory (Saunders & Knill, 2004). The same researchers found that artificial noise can be used to manipulate observers' reliance on the cues, demonstrating the system's ability to adapt to changes in uncertainty in the environment. This adaptation is suggestive of an implicit model of uncertainty that is available at all times, which is also supported by Whiteley and Sahani (2008), who found that observers' decisions were sensitive to current uncertainty even in conditions of minimal feedback.

In a Bayesian system each level of computation maintains representations of all possible values of the parameters and their associated probabilities. This means that the information from different cues and

modalities can be integrated and propagated without the necessity of committing to a particular interpretation too early.

The MLE as described above has ‘flat’ priors, this means that it is assumed that there is no existing or ‘prior’ information in the visual system that will affect the judgement or that the variance of the prior is so large as to have minimal influence. This means that the weights of the signals are solely determined by their variance and are not affected by the statistical history of events. However, Bayesian Theorem can be introduced to the above equations by simply adding a third component to represent the prior

$$\hat{S} = w_{Col}\hat{S}_{Col} + w_{Lum}\hat{S}_{Lum} + w_{Prior}\hat{S}_{Prior}$$

**Equation 1.4**

where

$$w_{Col} = \frac{r_{Col}}{r_{Col} + r_{Lum} + r_{Prior}} \quad \text{and} \quad w_{Lum} = \frac{r_{Lum}}{r_{Col} + r_{Lum} + r_{Prior}} \quad \text{and}$$

$$w_{Prior} = \frac{r_{Prior}}{r_{Col} + r_{Lum} + r_{Prior}}$$

**Equation 1.5**

and the reliabilities are calculated as before.

The question remains as to how priors and sensory information are weighted; how the reliability would be measured in order for the visual system to calculate the appropriate weight. It may be that the final weighting is not static and is determined by the current reliability and availability of the cue dependent on location and time (McGraw, Whitaker, & Badcock, 2000). In

order to support this, the visual system must accommodate dynamic changes in cue weighting similar to the situation dependent models of uncertainty described above (Whiteley & Sahani, 2008).

There are several physiologically plausible models that suggest neuronal representations to account for Bayes or MLE optimal behaviour. One possibility is a binary system, where there are two populations; 'on' and 'off' and the proportional difference between them represents that probability distribution. This could also be achieved with a single population of neurons that responds proportionally creating a likelihood ratio (Knill & Pouget, 2004).

Single cell recordings of the lateral intra-parietal (LIP) area provide evidence for both kinds of system. Platt and Glimcher (1999) found that when monkeys are trained to perform one of two possible saccades two sets of LIP neurons fire proportionally to the probability that the saccade ends in their receptive field. Conversely, Gold and Shadlen (2001) found that when monkeys are trained to distinguish between two possible motion directions a single set of LIP neurons respond by integrating information over time in a manner consistent with computing a likelihood function.

These binary schemes are only suitable when there is a clear dichotomy, for continuous variables different systems must be considered. Convolution codes have been suggested as a way that continuous variables could be encoded, the likelihood functions for observed stimuli would be convolved by the prior distribution. This idea is based on the premise that a

probability density function can be represented by its values at a few points along the ordinate. Each neuron would compute the dot product between the probability density function and its Gaussian tuning curve. For example, to calculate location given colour and luminance position cues, one would multiply the likelihood functions for colour indicating the correct position given the position that is being indicated, luminance indicating the correct position given the position that is being indicated and a prior distribution over position. If there are neurons that represent samples of the likelihood functions and the prior distribution a point-by-point product operation is equivalent to multiplying the functions themselves (Zemel, Dayan, & Pouget, 1998).

It is also possible that the log of the probability density function is encoded, rather than the function itself. In this scenario the point-by-point product required when using convolution codes is replaced by a point-by-point summation (because  $\log(a)+\log(b)=\log(a.b)$ ). This method is consistent with the evidence that LIP neurons integrate by summation (Gold & Shadlen, 2001).

An alternative to convolution coding is gain encoding (Pouget, Dayan, & Zemel, 2003), this uses the near-Poisson nature of neural noise (Tolhurst, Movshon, & Dean, 1983) to code simultaneously the mean and variance of the density function. For example, there are neurons in V1 that have bell-shaped tuning curves for orientation (Hubel & Wiesel, 1968). If these are ranked by their preferred orientations it produces a 'hill' of activation and, for

any given stimulus trial, this activation is distorted by near-Poisson noise. A Bayesian decoder could translate this into a posterior distribution over orientation given the activation hill (Sanger, 1996). In this case the noisy activation hill would be a neuronal representation of the posterior; the position of the peak indicating the mean and gain (amplitude) being the variance. The gain can represent the variance in this scenario because for Poisson noise the variance of the spike count is proportional to the gain, therefore, a high gain indicates a high signal to noise ratio and a narrow distribution (Knill & Pouget, 2004).

None of these possible coding schemes are mutually exclusive. The perceptual uncertainties that the brain is required to process can take many forms and so may use many encoding schemes (Knill & Pouget, 2004). Whilst there is evidence that observers use a Bayesian strategy in many scenarios, there may be other mechanisms that *mimic* Bayesian methodology but, do not require explicit probability representation. For example, probability matching, when participants are asked to make predictions about uncertain events, the probability of them choosing an event typically matches the probability of that event occurring. If the aim is to correctly predict which event will occur, this is a sub-optimal strategy. The optimal strategy would be to always predict the most probable event i.e. if event A occurs 70% of the time the optimal strategy would be to always predict that A will occur. However, participants actually only predict that event A will occur 70% of the time. This can be explained in a Bayesian manner by including the assumption that the sequence of trials contains predictable patterns (Wozny, Beierholm,

& Shams, 2010). However, a simple strategy of 'win-stay, lose-shift' can also explain participant behaviour. If participants stay with an option as long as it offers a reward and switch as soon as it ceases to offer a reward that also creates the pattern observed above.

Bayesian models have been criticised for being so flexible that they can account for a wide range of outcomes and therefore successful predictions made by these models are insufficient to provide evidence that the mind operates in a Bayesian fashion (Bowers & Davis, 2012). Conversely, of course, it might be seen as a benefit that the Bayesian framework is a general model that can be used to conceptualise a large range of challenges, rather than one designed to answer 'one-off' questions (Griffiths, Chater, Norris, & Pouget, 2012). The Bayesian framework has been criticised for being unfalsifiable, but no theoretical framework is directly falsifiable. The success, or failure, of a framework can only be judged by its ability to generate successful models, which are falsifiable, and new lines of research (See Bowers & Davis, 2012; Griffiths, et al., 2012 for a more detailed discussion).

For example, Hillis et al (2004) generated a falsifiable model within the Bayesian framework that successfully predicted participant behaviour in cue combination of texture and disparity in slant perception. This finding led to new lines of research including those reported in Chapters 4 and 5 of this thesis. Therefore, in this case the Bayesian framework is successful; a model based on the framework successfully predicted behaviour and led to new lines of research.

#### 1.4. Combining colour and luminance in shape from shading and contrast detection tasks

There has been much debate on the purpose of colour vision. It may be used to discriminate ripe fruit from foliage (Mollon, 1989), assist with recognition and memory (Gegenfurtner & Rieger, 2000), facilitate shadow recognition (Kingdom, Beauce, & Hunter, 2004) and perceive 3D shapes (Kingdom, 2003). Previously it was believed that colour information was represented and processed separately from other types of information, however, as seen above, it now seems more likely that colour is represented and processed together with other types of form information.

Luminance signals provide ambiguous information about surfaces because these signals are a combination of reflectance and illumination. The ideal way to estimate the real world properties of a luminance signal is to combine the inputs to the visual system with prior knowledge about reflectance changes and non-uniform illumination. Kingdom (2008) considered that prior knowledge would be necessary to make these kinds of judgements and suggested a list of heuristics necessary to determine whether a luminance discontinuity is a change in reflectance or illumination.

Colour vision is useful in distinguishing reflectance from illumination changes because chromatic changes typically occur at object, but not shadow, boundaries (Kingdom, et al., 2004). Shadows are important for perceiving the spatial arrangement of stimuli, and are processed within the confines of certain priors. In an illusion created by Mamassian and colleagues (1998) the

position of a shadow can cause a sphere to appear to roll to the bottom of a box or rise in a frontal plane. The shadow movements that led to this perception of depth could be caused by any combination of changes to viewpoint, light source, the object causing the shadow or the background surface, but it is reliably interpreted as movement in depth. This once again supports the idea that the visual system processes information not only on the basis of sensory information but also prior experience with the world.

Chromatic changes that are aligned with shadow borders suppress the identification of shadows and those that are not aligned facilitate this process (Kingdom, et al., 2004). This supports the suggestion by Kingdom (2008) that there are prior assumptions about the meanings of the relationship between colour and luminance signals. In this case the prior would be that achromatic edges suggest inhomogeneous illumination whereas combined chromatic and luminance edges suggest surface changes. Chromatic variations can also 'unmask' transparent achromatic targets (shadows) in densely variegated achromatic backgrounds (Kingdom & Kasrai, 2006). This could suggest that colour processing is suppressing luminance noise that is impairing detection of luminance-defined targets.

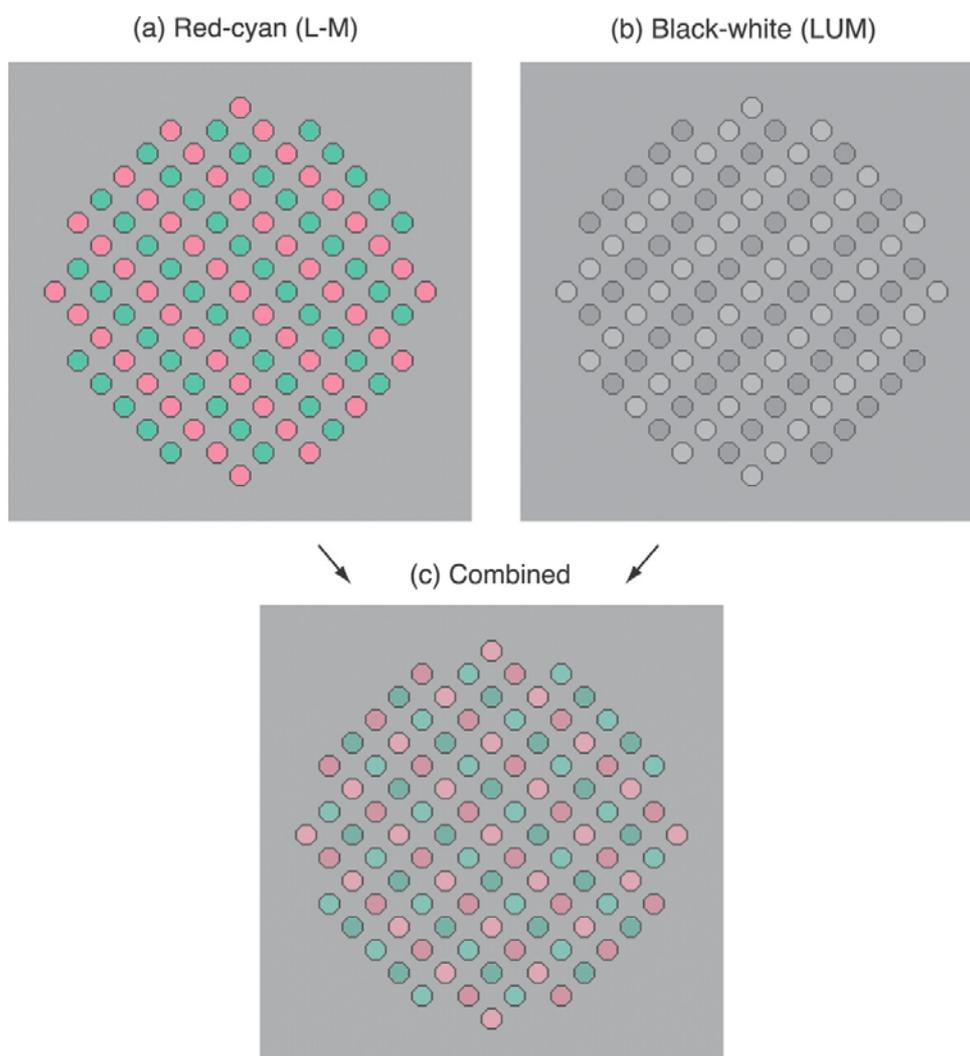
When achromatic and chromatic gratings of different orientation are combined the resulting plaid appears to be three dimensional (an example of the shape from shading effect). Kingdom (2003) took advantage of this phenomenon to investigate how aligned chromatic and luminance discontinuities differ from their unaligned counterparts. He took a chromatic

plaid and added a luminance grating, of the same orientation as one of the chromatic gratings. When the luminance grating was not aligned with the chromatic grating there was an impression of depth. However, when the luminance grating was aligned with the chromatic grating the shape-from-shading effect was suppressed, regardless of the colour direction of the chromatic component (Kingdom, Rangwala, & Hammamji, 2005a). This led to the proposal of a new role for colour vision in processing three-dimensional structures, where once again achromatic discontinuities represent changes in illumination and combined discontinuities represent object edges. In Bayesian terms this would mean that the visual system assigns each luminance discontinuity a probability that it arose from changes in illumination rather than reflectance and that this probability was constrained by a prior based on the spatial relationship between the achromatic and chromatic discontinuities (Kingdom, 2003).

It is generally accepted that visual performance is impaired when tested using isoluminant stimuli, however this is not necessarily due to deficits in chromatic processing (see Cavanagh 1991 for a review). If colour and luminance information are not processed independently then isoluminant stimuli would not be suitable for isolating the chromatic system (Gur & Akri, 1992). It may be that colour vision is not only for encoding chromatic information but also enhancing luminance based processing, this would mean that chromatic processing cannot be fully investigated in the absence of luminance information.

In keeping with this idea, contrast sensitivity for combined colour and luminance targets is lower than would be predicted by the contrast sensitivity of either cue alone, this facilitation necessitates that the channels are integrated, not independent (Gur & Akri, 1992). Therefore, investigation using isoluminant stimuli alone may not tell us about how that information would be processed in the presence of luminance information. The ability to discriminate a circle and an ellipse is enhanced when both colour and luminance information is present, as opposed to either channel alone. This also supports the idea that one role of colour vision is to enhance luminance-based vision and that luminance information must be present for 'normal' activation to occur in the colour system (Syrkin & Gur, 1997).

Kingdom and colleagues (2010) investigated the comparative saliency of suprathreshold colour and luminance signals. They used lattices of circles in two conditions; separated (colour and luminance modulations were temporally separated) and combined (Figure 1.4). In the combined conditions 48% more luminance contrast was required relative to when the cues were presented separately and subsequent experiments showed that this was caused by colour masking the luminance information. However, this only occurred when the colour and luminance information was present together within the circles of the lattices. If the components were segregated (each circle contained only one type of information), the reverse occurred with luminance masking colour information. It was suggested that chromatic masking occurred to facilitate segmentation by material by disregarding non-uniform luminance changes.



**Figure 1.4.** Example stimuli from Kingdom et al (2010) (a) red-cyan, (b) black-white component patterns and (c) the two combined. Figure used with permission.

Research into facilitation between colour and luminance in contrast detection highlights the complexities of the relationship between the two cues. Chromatic pedestals do not facilitate contrast detection of luminance targets although they do produce masking at higher contrasts, with very similar features to a luminance mask (K. K. De Valois & Switkes, 1983). Subthreshold luminance pedestals appear to be entirely discounted when combined with a chromatic test stimulus (Cole, Stromeyer, & Kronauer,

1990). However, a suprathreshold luminance pedestal combined with a chromatic test produces facilitation (Cole, et al., 1990; Switkes, Bradley, & De Valois, 1988). This facilitation also occurs when the pedestal is only a 'ring' rather than full field (Cole, et al., 1990), and is increased from a ~3-fold decrease in threshold to a ~7-fold decrease when low-spatial-frequency square-waves are used as opposed to spots (Gowdy, Stromeyer, & Kronauer, 1999).

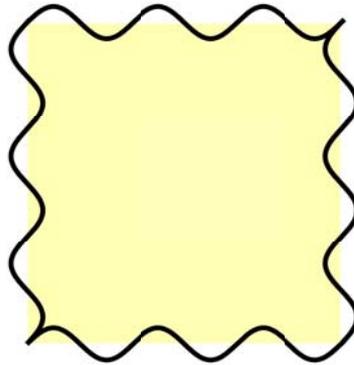
Gowdy and colleagues (1999) suggested that this increase in facilitation may be due to the sharp pedestal edges promoting segmentation; the colour is spatially demarcated by the luminance edges, then integrated between the luminance edges and finally the colour difference is compared across the luminance edges. They suggested that this represented a change in the chromatic mechanism from a 'blob' detector, tuned to broad areas of chromatic information, to an 'edge' detector, tuned to chromatic boundaries delineated by luminance information.

It has been argued that luminance plays a privileged role in edge detection when compared to colour. However it may be that luminance is not processed differently or given different weight, merely that in natural scenes luminance gains a privileged role because it has higher effective contrast than chromaticity (Rivest & Cavanagh, 1996).

### 1.5. The dominance of luminance information in edge detection tasks

It has been argued that pure isoluminant edges are rare in natural images, which would mean that in the majority of cases colour is not necessary for the detection of object edges (Zhou & Mel, 2008). However, whilst the majority of edges in natural scenes are a combination of colour and luminance, isoluminant edges are not in fact any rarer than achromatic edges, and the contrasts of the components of the combined edges have sufficient variation to be considered independent (Hansen & Gegenfurtner, 2009). This means that isoluminant edges are not inherently any less useful than luminance defined edges. Despite this there are many demonstrations of luminance information dominating chromatic information.

Illusions offer striking examples of luminance information appearing to constrain the perceived location of chromatic edges. For example, in the Boynton illusion (Figure 1.5) straight chromatic edges appear to align with nearby irregular luminance edges; the edge location is determined by luminance information (Kaiser, 1996).



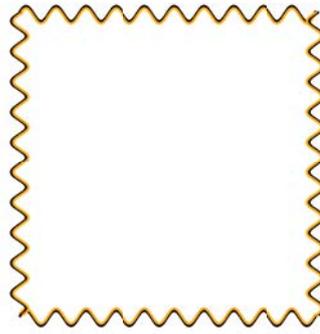
**Figure 1.5.** Boynton Illusion. When viewed at a sufficient distance the yellow square appears to fill in the luminance boundary. The yellow square never appears to cross the luminance lines but does appear to fill the full extent of the area. That is, the luminance information constrains the apparent location of the chromatic information.

Similarly in the Spanish castle illusion (Figure 1.6) chromatic afterimages appear to be constrained by luminance information. Adaptation to a negatively coloured image followed by viewing of a blank field produces a blurry, indistinct chromatic afterimage. However, if adaptation is followed by an achromatic version of the original image it appears to be sharp and normally coloured (Sadowski, 2006). The introduction of achromatic information makes the chromatic information appear sharp and no longer indistinct.



**Figure 1.6.** An adaptation of the Spanish castle illusion (Sadowski, 2006). Visual adaptation to the negatively coloured image followed by viewing the achromatic image causes the achromatic image to appear normally coloured and maintains the sharp appearance of its edges.

Chromatic filling-in also appears to be constrained by luminance information, as demonstrated by the watercolour effect (WCE, Figure 1.7), where colour appears to spread between luminance boundaries but does not cross them (Pinna, et al., 2001). The WCE could suggest that colour has a greater role in perception of surface properties as opposed to perception of edges. For example, colour constancy, specifically illuminant discounting and estimation, can be used to facilitate surface segmentation (see Foster, 2011 for a review of colour constancy). Colour has also been shown to reduce luminance noise in complex displays, such that dark achromatic targets are unmasked by chromatic variation in the background (Kingdom & Kasrai, 2006). This could be interpreted to mean that chromatic information can be used not only to segment chromatic variation but also facilitate segmentation of luminance information.



**Figure 1.7.** Watercolour Illusion. The colour from the inner orange line appears to spread causing the central area to appear to be coloured. This colour spreading effect is constrained by the outer luminance defined contour (Pinna, Brelstaff, & Spillmann, 2001).

Luminance information also appears to dominate in natural scenes: blurring only chromatic information has little effect on the percept of these scenes, whereas blurring luminance information has a profound effect on appearance (Wandell, 1995, Figure 7). This phenomenon is not well understood but has been exploited in industrial settings (Isono, Sakata, & Kusaka, 1978; Oho & Watanabe, 2001). In particular analogue television systems (PAL, NTSC, SECAM) separate the luminance and two colour components and transmit each colour at approximately one quarter the resolution of the luminance in order to save bandwidth (Hunt, 2004).

#### 1.6. The current thesis

It is clear that both chromatic and luminance information can be used in form processing. It is also clear that colour and luminance are combined asymmetrically in certain tasks. However, it is not clear whether these

asymmetries are present in cue combination of colour and luminance in edge detection.

This thesis has three aims in order to gain a better understanding of how colour and luminance are combined in edge detection. 1) To determine whether presenting colour and luminance information together improves performance. 2) To investigate how the visual system resolves conflicts between colour and luminance edge information. 3) To explore whether colour and luminance edge information is always combined in the same way.

Chapter 3 explores the masking of chromatic blur by sharp achromatic information in natural scenes. This phenomenon is well known and accepted, however, there is no previous research which quantifies this effect or investigates the mechanisms that may underlie it. It was found that blur discrimination thresholds for chromatic blur were poorer in general than those for achromatic blur. However, thresholds for chromatic blur combined with sharp luminance information were far higher than those for chromatic blur alone. Therefore, the phenomenon cannot be attributed to poorer acuity in chromatic processing. In Experiments 3.2 and 3.3 the underpinnings were further investigated and it was shown that the phenomenon could not be explained by either the lower effective contrast of chromatic information or statistical differences between the structure of colour and luminance information in natural scenes. This suggests that there is a mechanism that is prioritising luminance information regardless of the relative quality of chromatic information.

Chapter 4 investigates whether the presence of both colour and luminance information improves performance in an edge localisation task. Variability of edge localisation judgements was measured, using a staircase procedure, for isoluminant, achromatic and combined bipartite edges. The achromatic and isoluminant measurements were used to generate predictions for performance when both cues were available, according to three models; 'winner takes all', unweighted averaging and weighted averaging (MLE). The models were then compared to the behavioural data. Unfortunately, due to the small differences in model predictions it was not possible to discriminate between the performance of the three models.

In Chapter 5, method of adjustment was used, as an alternative to a staircase procedure, to see if this could allow discrimination between model predictions for cue combination of colour and luminance in aligned edges. This methodology was also not sensitive enough to discriminate between the three models. How conflicting colour and luminance edges are localised was also investigated using method of adjustment. MLE was used to predict where the participants would judge the edge to be, and behavioural data showed that participants weighted chromatic information more heavily than was predicted. This may suggest a Bayesian prior promoting the chromatic information, which may be due to chromatic information having greater utility in object edge detection in natural scenes.

A novel technique for investigating edge detection was introduced in Chapter 6; perturbation discrimination. Gratings were spatially perturbed and

participants were required to detect that perturbation in a 2IFC task. Experiment 6.1 was designed to determine whether the spatial arrangement (aligned or orthogonal) of chromatic and luminance components affected perturbation discrimination. No difference was found between perturbation thresholds for the perturbed grating alone or in combination with aligned or orthogonal cross-channel masks. This suggests that participants were able to disregard irrelevant information.

Experiment 6.2 investigated whether the type of chromatic or luminance gratings presented affected processing. Perturbation thresholds were measured for both achromatic and isoluminant square-wave and line gratings in isolation. Perturbation thresholds were then measured for these stimuli when combined with a mask grating of the other type and other channel. For example, the perturbed achromatic square-wave grating was combined with an isoluminant line-grating mask. The effect of introducing a mask was determined by measuring the difference between the target grating alone and when combined with the mask.

The introduction of a chromatic square-wave mask facilitated perturbation discrimination for luminance lines; the introduction of a chromatic square-wave mask *improved* perturbation discrimination. However, the introduction of an achromatic square-wave mask had little effect on perturbation discrimination for isoluminant lines. This asymmetric relationship demonstrates the complexity of cue combination of colour and

luminance and supports the idea that chromatic information can become 'tied' to luminance information.

In the final experimental chapter the conflicting results of Chapters 3 and 5 were investigated. Chapter 3 showed luminance information dominating chromatic information and Chapter 5 showed the reverse. This could have been due to differences in the stimuli used (natural scenes versus bipartite edges) or the task (blur discrimination versus edge localisation), therefore the method used in Chapter 3 was used to measure blur discrimination thresholds for bipartite edges. Despite the greatly simplified stimuli the results exactly replicated those found in Chapter 3. This suggests that whilst luminance information dominates in blur discrimination tasks, regardless of stimulus type, this dominance does not translate to edge localisation tasks.

## 2. General methods

The following describes the major methods and approaches used in this thesis. There are variations across the experiments and therefore specifics are given in the relevant methods sections.

### 2.1. Participants

All participants had normal or corrected-to-normal vision, were not colour anomalous and gave their informed consent to participate in the studies. All procedures were approved by the School of Psychology Ethics Committee, University of Nottingham, UK and were in accordance with the Helsinki Declaration.

### 2.2. Apparatus

Unless otherwise specified the following apparatus was used. A computer-controlled cathode-ray-tube (CRT) monitor was used to present stimuli. The monitor used was a 19-in Vision Master Pro 454 (Iiyama) with resolution of 1024 x 768, running at a refresh rate of 85 Hz. Stimuli were presented and data collected using PsychoPy (Peirce, 2007). All data collection occurred within a darkened room with a chin rest to ensure that the participant viewed the stimuli from a constant distance.

### 2.3. Bits ++ Digital Video Processor

A standard graphics card has a dynamic range of 8 bits for each of its 3 output channels (R, G and B). Where a higher luminance resolution was required (for example, to measure contrast detection thresholds) this was

increased to 14 bits by a 14-bit Digital to Analog Converter (DAC) system (Bits++, Cambridge Research Systems, Cambridge, UK). This allows the monitor to display a much larger range of contrasts, which in turn allows more precise stimulus presentation.

#### 2.4. Gamma Correction

Gamma correction was performed on all monitors used in this thesis. A photospectrometer (PR655, Photo Research, Chatsworth, CA, USA) was used to measure the luminance of 64 test patches. The gun outputs of these patches were evenly distributed from gun values of 0 to 255 and the process was carried out for each gun (red, green and blue) independently. These measurements were used to generate a set of gamma functions using Equation 2.1.

$$L(V) = a + (b + kV)^\gamma$$

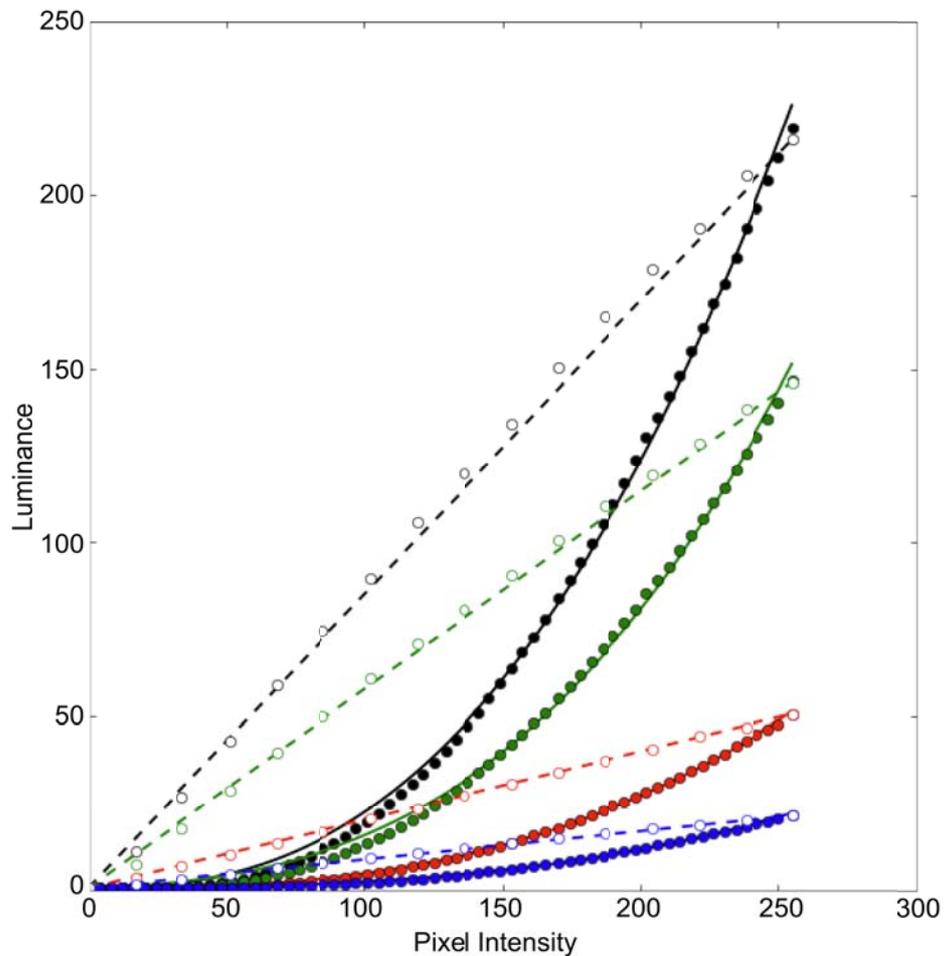
Equation 2.1

Where  $L$  is the final luminance value,  $V$  is the required intensity (from 0 to 1),  $a$  is the minimum luminance measured,  $b$  is the range of luminance measured and  $\gamma$  is the gamma value. The inverse values were then calculated using Equation 2.2 and used to build a look-up table (LUT) with linear luminance outputs.

$$LUT(V) = \frac{((1 - V)b^\gamma + V(b + k)^\gamma)^{\frac{1}{\gamma}} - b}{k}$$

Equation 2.2

The gamma correction was tested by repeating the initial measurements with the gamma corrected values (Figure 2.1). The presentation of test stimuli, gamma correction calculations and curve fitting were all performed automatically by PsychoPy.

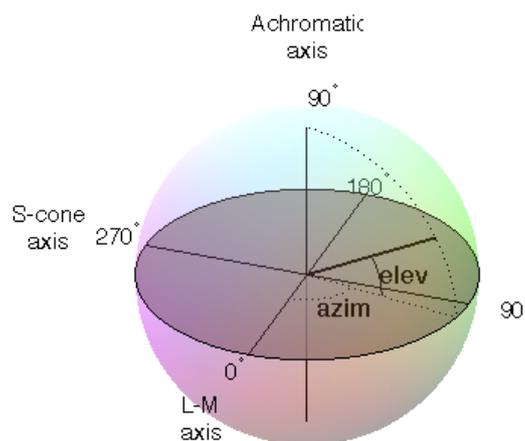


**Figure 2.1.** Example of the curves fitted before (solid dots) and after (empty dots) gamma correction for each gun individually and in combination.

## 2.5. DKL Space

LGN neurons each have a preferred colour direction and these cluster around the two cardinal isoluminant axes of MB-DKL space (L-M and S)

(Derrington, Krauskopf, & Lennie, 1984). MB-DKL space is a combination of the colour coordinates system introduced by MacLeod and Boynton (1979) and the cardinal colour directions determined by Krauskopf and colleagues (1982). Colour is represented in three-dimensional spherical space, luminance information as elevation ( $-90^\circ$  to  $+90^\circ$ ) and chromaticity across the azimuth. The axis along which only L-M information changes runs along  $0^\circ$ - $180^\circ$  with an orthogonal axis sensitive to changes in the S-(L+M) signal running along  $90^\circ$ - $270^\circ$ . In this space any light can be described in terms of azimuth, elevation and contrast (Figure 2.2).



**Figure 2.2.** A graphical representation of MB-DKL space that allows representation of lights in terms of L-M, S-(L+M) and luminance information (Figure modified from Peirce, Solomon, Forte and Lennie (2008) with permission).

The cardinal axes were determined after a series of experiments showed that detection thresholds for chromatic changes were raised

following adaptation, but that this effect was highly selective (Krauskopf, et al., 1982). There is no cross adaptation between yellow-blue, red-green or luminance defined stimuli and this selectivity is not found for intermediate directions. Therefore there are three directions that a light can be described in L-M (red-green), S-(L+M) (blue-yellow) and luminance, each representing a different visual pathway.

MB-DKL space is a useful way to describe image properties as it allows the creation of stimuli that specifically activate the different channels. This allows us to compare responses from the chromatic channels to responses from the luminance channel. Investigating the differences and similarities of these responses can be used to make inferences about how the signals may be combined. This principle underlies all the experiments described in this thesis.

## 2.6. Generating isoluminant stimuli

Isoluminant stimuli were generated in two ways in this thesis; using photometric measurements and using psychophysical measurements. When generating isoluminant stimuli using photometric measurements a PR655 spectroradiometer (Photoresearch Inc., Chatsworth, CA.) was used. The power spectrum for each gun of the monitor is measured and then converted from RGB to MB-DKL space using Smith and Pokorny (1975) cone fundamentals.

When isoluminant stimuli were created based on psychophysical measurements a minimum motion paradigm was used to measure individual

subjective isoluminant points (Anstis & Cavanagh, 1983). This technique is based on the principle that colour and luminance are not integrated temporally and so if there is no luminance component in the chromatic Gabors no clear direction of motion will be perceivable.

All Gabors were of size  $2.0^\circ$ , spatial frequency 2.0 cpd and were presented for 4 frames. The achromatic Gabors were presented at 0.1 Michelson contrast and chromatic Gabors were presented at full contrast. For half the trials an achromatic Gabor of phase 0.0 was followed by a chromatic Gabor of phase 0.25, then an achromatic Gabor of phase 0.5 and a chromatic Gabor of phase 0.75. For the remaining half the phases were reversed such that the initial achromatic grating had phase 0.75. Participants were required to indicate the direction the grating appeared to be drifting in. The elevation of the chromatic gratings was varied using a staircase procedure designed to find the 50% correct point i.e. the point where no consistent direction of motion was perceived. The elevations generated using this procedure were then applied as deviations from photometric isoluminance in MB-DKL space.

Across all participants ( $n=8$ ) deviations from photometric isoluminance ranged from  $0.185^\circ$  to  $-5.901^\circ$ , the mean deviation was  $-2.774^\circ$ . This represents deviations between 0.17% and 5.31% of the maximum possible deviation from photometric isoluminance ( $45^\circ$ ).

One participant (RJS) had their LM isoluminant point measured twice, four months apart, elevations of  $-3.478^\circ$  and  $-3.743^\circ$  were recorded. This demonstrates the reliability of the measurement and subsequently all other

participants only had their subjective isoluminant point measured on one occasion.

An individual's isoluminant point varies over time, spatial frequency and across the retina (Logothetis & Charles 1990). Therefore, there is always a risk that an isoluminant stimulus may contain luminance information. However, in this thesis the presence of a luminance artefact would only decrease the effects that we are investigating (see also Section 3.4).

## 2.7. Defining contrast

Luminance contrast can be defined in terms of Michelson contrast (Equation 2.3) or root mean square (RMS) contrast (Equation 2.4). Michelson contrast is based on the highest ( $I_{max}$ ) and lowest ( $I_{min}$ ) luminance values present in the stimulus. RMS contrast is based on the standard deviation of the luminance values present in the stimulus, where intensities  $I_{ij}$  are the  $i$ -th  $j$ -th element of a stimulus of size  $M$  by  $N$  and  $\bar{I}$  is the mean luminance value.

$$\frac{I_{max} - I_{min}}{I_{max} + I_{min}}$$

**Equation 2.3**

$$\sqrt{\frac{1}{MN} \sum_{i=0}^{N-1} \sum_{j=0}^{M-1} (I_{ij} - \bar{I})^2}$$

**Equation 2.4**

For chromatic contrast in DKL space, values are typically specified as fractions of the maximal amplitude of modulation along each of the cardinal

axes. For instance, a contrast of 1.0 along the L-M axis represents the maximum modulation along that axis, permitted by the gamut of the monitor. For the luminance axis, the maximal modulation would be identical to Michelson contrast if the monitor were able to produce a 'black' of zero luminance. In our case, the luminance range of the monitor was 0.724 – 219.20 cd/m<sup>2</sup>, giving a maximum Michelson contrast of 0.993. Therefore, our luminance contrast values are, to all intents and purposes, identical to the Michelson values and will be referred to as such (the reader can simply multiply any luminance contrast by 0.993 in order to obtain the 'true' Michelson contrast).

It should be remembered throughout this thesis that, in this colour space, the contrast values between axes are entirely arbitrary and are governed only by the gamut of the monitor.

## 2.8. Adaptive staircase procedures

Adaptive staircases allow adjustment of stimulus parameter until the feature being tested is just discriminable or the stimulus is just detectable. The intensity of the feature being tested is determined by the participant's previous responses. Initially the task is easy and then becomes more difficult until the participant gives an incorrect response (one-up) at which point the staircase reverses and the task becomes easier until the participant gives the correct response, a specified number of times (n-down), when the staircase will reverse again. This allows measurements to be taken across the psychometric function focusing on a particular percentage correct.

In this thesis, the step size is large initially gradually getting smaller to converge on the desired percentage correct and staircases are aborted after 50 trials. We use one-up, one-down staircases (designed to converge on the 50% correct point) to measure detection thresholds and one-up, three-down staircases (designed to converge on the 79.4% correct point) to measure discrimination thresholds.

### 2.9. Method of adjustment

Method of adjustment simply involves participants adjusting the stimuli until it meets some criteria. All the experiments in this thesis that use method of adjustment require the participant to use the mouse to move an edge until it is aligned with a marker. Method of adjustment produces a histogram of responses and so, unlike adaptive staircases, can detect bimodality of responses.

### 2.10. Curve fitting

Data collected using an adaptive staircase procedure can be fit to a curve using a Weibull function of the form

$$y = \textit{chance} + (1 - \textit{chance}) \left( 1 - \exp \left[ - \left( \frac{x}{\alpha} \right)^\beta \right] \right)$$

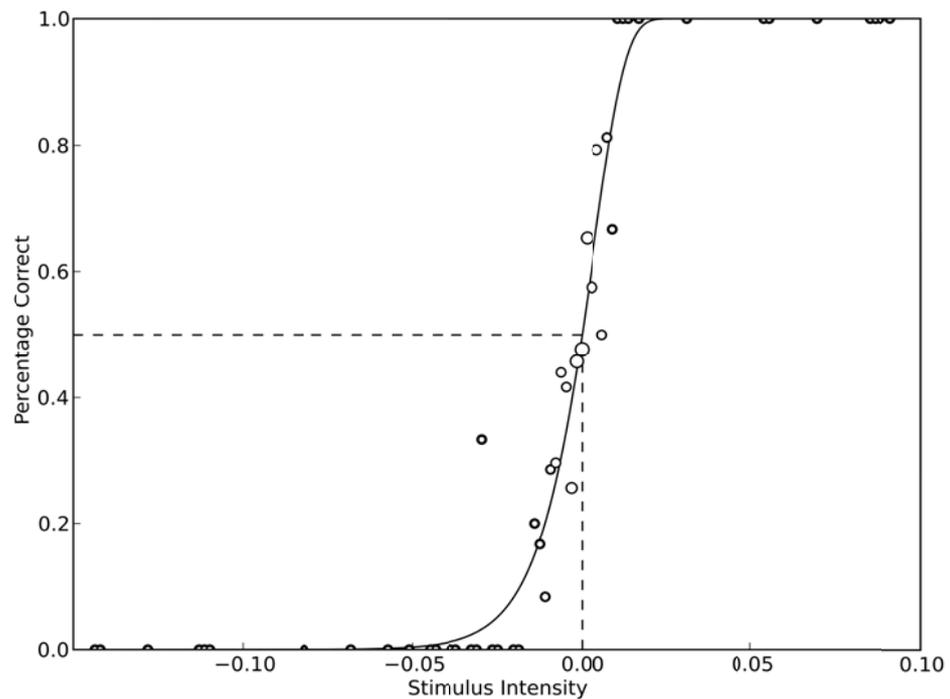
Equation 2.5.

with the inverse:

$$x = \alpha \left( -\log \left( \frac{1 - y}{1 - \textit{chance}} \right) \right)^{\frac{1}{\beta}}$$

Equation 2.6.

Where  $y$  is the probability of giving a correct response as a function of variable of interest ( $x$ ).  $\alpha$  corresponds to the value of the variable of interest at the desired percentage correct point and  $\beta$  corresponds to the slope of the psychometric function. Thresholds can be derived from this fit as the point at which the observer was at the desired percentage probability of responding correctly (Figure 2.3).



**Figure 2.3.** Example of a psychometric function created by fitting a Weibull function to staircase data. The dashed line shows the threshold at 50% correct is approximately 0.0.

### 2.11. Using the last six reversals as an estimate of threshold

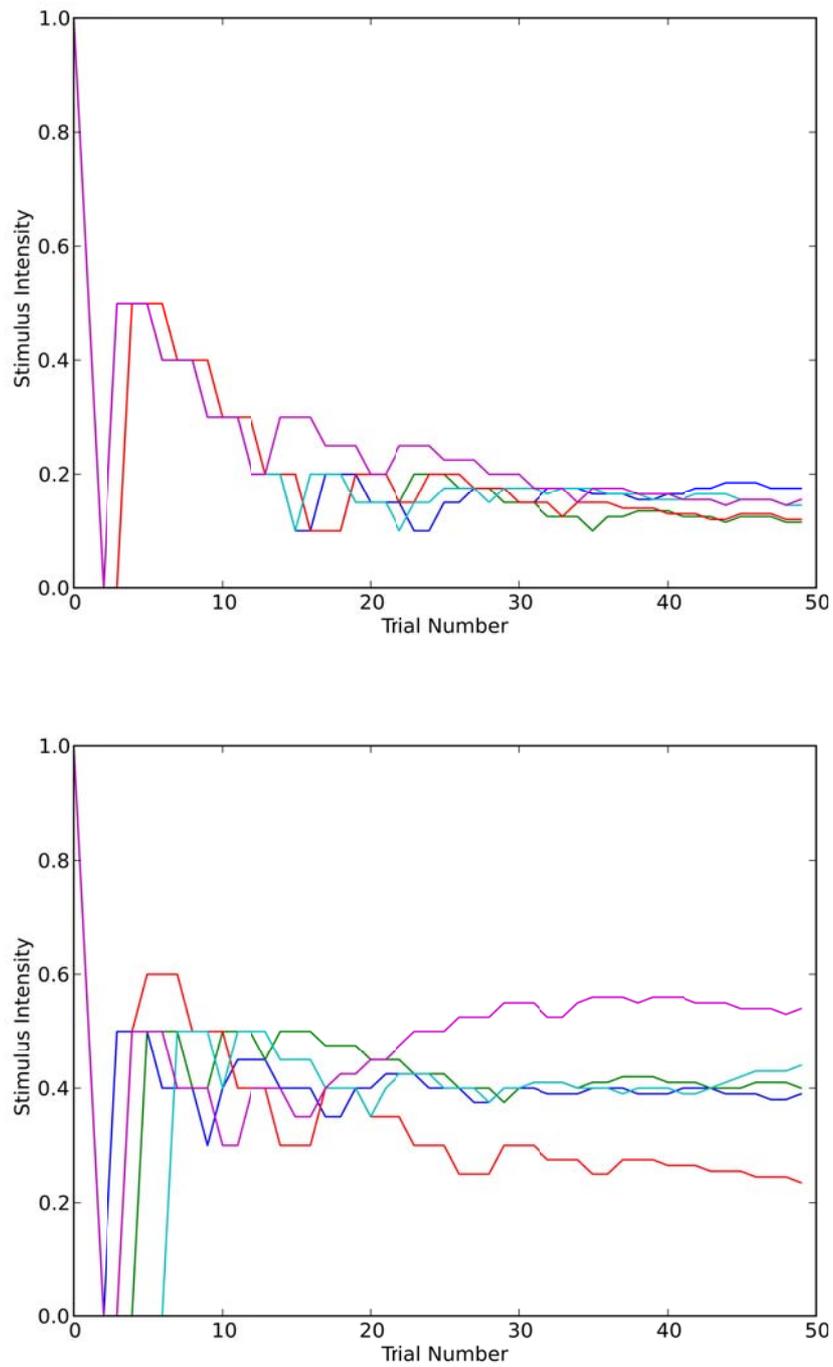
Using the fitting procedure described above is usually the best way to generate estimates of threshold as it allows inclusion of the all the data points collected. However, in some cases variability across staircases means that the quality of curve fits is poor. When data is consistent across staircases they converge at the same point, as shown in the top panel of Figure 2.4. However, when data is not consistent across staircases they no longer converge at the same point, as shown in the bottom panel of Figure 2.4.

Variability across staircases causes the data points to become scattered. When this occurs the error surface, describing the quality of the fit

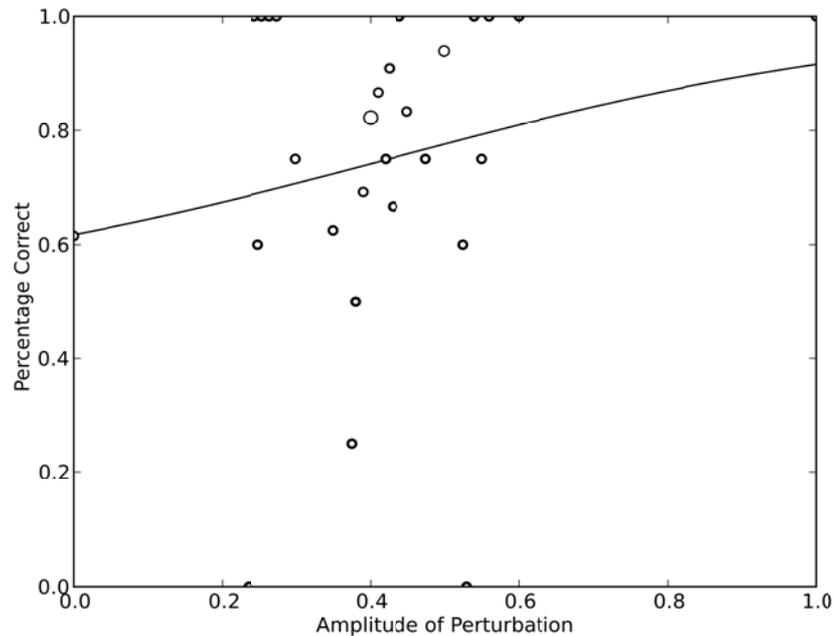
as a function of the curve fit parameters, becomes noisy and contains several local minima (Figure 2.5). Furthermore, extreme values can cause the slope of the function to become too shallow and threshold estimates to become inaccurate. In the data shown in the lower panel of Figure 2.4 and Figure 2.5 the function estimates the threshold to be 0.571, but visual inspection suggests it should be closer to 0.4.

In these cases a more reliable method to calculate threshold is to take the mean of the last six reversals (see Section 2.8 for details of when reversals occur). In the case of the data presented in the lower panel of Figure 2.4 and Figure 2.5, the mean of the last six reversals is 0.405. This appears to be a better representation of the true threshold than was generated by the fitting procedure.

This method will be used when a fitting procedure generates threshold values that are outside the theoretically possible range of values for the task in question. For example, in a contrast detection task, possible contrast values range between zero and one. Therefore, threshold values must also be between zero and one, if this was found not to be the case after a fitting procedure, the mean of the last six reversals would be used as a measure of threshold.



**Figure 2.4.** Plots of example staircase data. Each line represents a different staircase. The top graph shows data with little variability across staircases that is suitable for fitting. The bottom graph shows data with variability across staircases that is not suitable for fitting.



**Figure 2.5.** Example of the results of attempting to perform a fitting procedure on data that are variable across staircases. The way that the data points are scattered means that the function fits the data very poorly. This plot is of the data shown in the lower panel of Figure 2.4.

### 2.12. Bootstrapping

Bootstrap resampling is a technique designed to give an estimate of the variance of a measurement. Resamples from a population are generated with replacement; after each measurement is selected from the sample it is replaced and can therefore be selected more than once in any given resample. For each resample the statistic of interest e.g. PSE, mean, variance is calculated. For example, the mean can be calculated for each resample, leading to a population of possible means. The standard deviation of this is therefore the standard error of the mean and the 95% confidence interval can

be easily calculated by sorting the values in the population. This analysis can be used for any statistical feature of the sample.

In this thesis, when bootstrapping was performed, 5000 resamples were taken for each data set. The exact way bootstrapping is used for each technique is outlined in the relevant methods section.

### 3. Luminance information constrains chromatic blur discrimination in natural scene stimuli

Introducing blur into the chromatic component of a natural scene has very little effect on its percept, whereas blur in the luminance component is very noticeable (Wandell, 1995, Figure 7). In this chapter the dominance of luminance information in blur discrimination is quantified and several potential causes of the effect are examined.

In addition to the phenomenon described by Wandell (1995, Figure 7) there are several more illusions that demonstrate the dominance of luminance information and how it appears to constrain luminance information, including the Boynton illusion (Kaiser, 1996), the Spanish castle illusion (Sadowski, 2006) and the water colour effect (Pinna, et al., 2001). For a detailed discussion of these and other examples please see Section 1.5.

Despite the number of examples of luminance constraining chromatic information it is not clear why this should occur. Specifically, it is unclear whether it is due to a mechanism that gives precedence to luminance information or whether it is due to other factors. For example, chromatic blur may not be visible, simply due to poorer spatial resolution in the processing of chromatic information (Mullen, 1985). It has been demonstrated that performance is poorer for several visual tasks, when isoluminant stimuli are used, including stereopsis (Krauskopf & Forte, 2002), global shape discrimination (Mullen & Beaudot, 2002) and, importantly, blur discrimination (Wuerger, Morgan, Westland, & Owens, 2000; Wuerger, Owens, & Westland,

2001). Blur discrimination, however, is only poorer for blue-yellow modulated stimuli; red-green modulated stimuli elicit similar thresholds to achromatic stimuli. Chromatic performance is not poorer for Vernier acuity when luminance and chromatic cues are presented in equal multiples of detection threshold (Krauskopf & Forte, 2002), and so it is not clear whether poorer chromatic acuity is a sufficient explanation for the masking of chromatic blur by sharp luminance information.

In the illusions mentioned above, and in natural scenes, luminance typically has a higher effective contrast than chromatic information (Rivest & Cavanagh, 1996). This could mean that luminance is not dominant due to a neural mechanism, but rather, because it is simply more visible.

Chromatic and luminance information in natural scenes may have different statistical regularities that could affect how blur is perceived. For example, if the luminance channel contained more high spatial frequency information it would be more susceptible to the blurring process. Chromatic and luminance information have some similar features in natural scenes: there is no significant difference in the number of isoluminant and achromatic edges in natural scenes (Hansen & Gegenfurtner, 2009) and both chromatic and luminance information have  $1/f$  amplitude spectra (Parraga, Brelstaff, Troscianko, & Moorehead, 1998). However, the two types of information may differ in other ways, for example, the number of range discontinuities or the distribution of spatial frequencies.

This chapter will investigate the interaction between sharp luminance information and blurred chromatic information in natural scenes. The effect will be quantified and the potential causes outlined above will be investigated to determine whether there is evidence for a neural mechanism with a bias toward luminance information.

### 3.1. Blur discrimination

The fact that blur is more obvious when applied to the luminance channel might simply be due to poorer blur discrimination for chromatic information. To test if this was the case we examined blur discrimination for chromatic information alone and in combination with sharp luminance information.

#### 3.1.1. Methods

##### *Participants*

One male and four female volunteers (including the author), aged between 23 and 29, participated in the study. Four of the participants (one male) were naive to the purpose of the study.

##### *Stimulus Generation*

The natural images were selected from the McGill Calibrated Colour Image Database (Olmos & Kingdom, 2004). The images were from the categories; flowers, animals, foliage and fruits. The images selected were the first in each category that was entirely in focus (with no obvious depth cues), well lit (not predominantly comprised of silhouettes or large areas of

darkness), and did not contain text. The central 512x512 pixels were then cropped from each image, leading to four equally sized natural images (Figure 3.2A and Figure 3.2B).

Stimuli were then converted into MB-DKL colour space (Derrington, et al., 1984; Macleod & Boynton, 1979). There are two potential issues that could introduce luminance artefacts into the chromatic information. First, the colour space transformations were not adjusted to individual subjects' isoluminance planes. Second, cone adaptation levels can potentially vary across the extent of a natural scene, meaning that using fixed-cone sensitivities (that are implicitly assumed in the MB-DKL space) could introduce luminance artefacts into the colour channels (A. P. Johnson, Kingdom, & Baker, 2005). However, if any luminance artefacts were present they would only serve to reduce the effect as demonstrated in Experiment 3.4.

All channels were scaled down in contrast by 50% in order to ensure that the images remained within the gamut of the monitor after chromatic manipulations. At low contrasts reducing the contrast further can increase blur discrimination thresholds, but this does not occur at the high contrasts used in this study (Watson & Ahumada 2011).

Blurring was performed by filtering the relevant channel(s) with a circular Gaussian whose width was varied with a staircase procedure according to the experimental condition. Either the luminance channel alone was blurred or both the isoluminant channels (by the same degree). To present the luminance channel alone, the contrast of both chromatic

channels were set to zero and, equivalently, to present only chromatic information the luminance contrast was set to zero. After the manipulations had been made the stimuli were converted back to RGB space, for presentation on the monitor.

Stimuli were presented with a size of 10° of visual angle along each edge, with a grey background and were viewed from a constant 52 cm distance.

### *Procedure*

A two-interval forced-choice (2IFC) design was employed. Participants were presented with the two images (foil and target) for 300 ms separated by a 500 ms interstimulus interval (ISI) and asked which appeared more blurred. The presentation order of the target and foil was randomised.

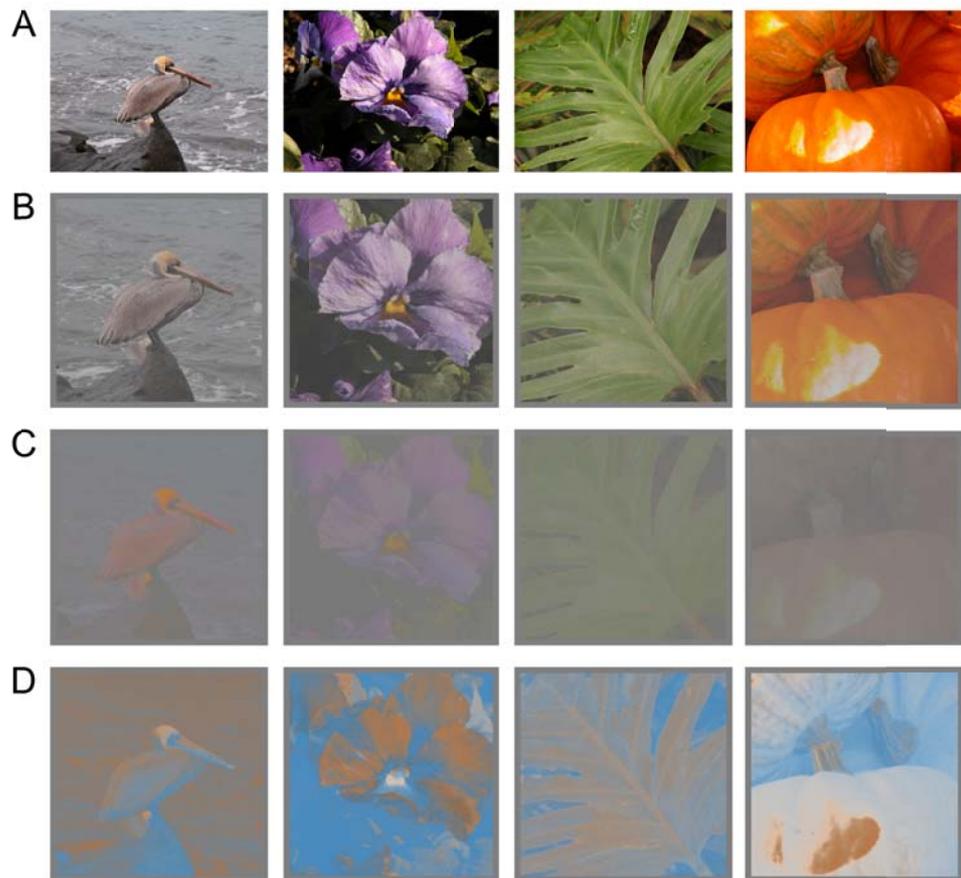
In each condition the minimal degree of blur that could be detected, the blur threshold, was measured. The blur thresholds for luminance information combined with sharp chromatic information and for chromatic information combined with sharp luminance information were measured. Furthermore, to determine whether any differences between the chromatic and luminance thresholds are simply caused by poorer blur discrimination, the thresholds for each form of information alone were measured. See Figure 3.1 for examples of the stimuli in the four conditions.

The blur threshold in each condition was determined using a one-up, three-down staircase procedure. The staircases controlled the amount of blur in the target images; a different staircase was implemented for each image

and these four staircases were randomly interleaved. Each participant collected a minimum of two staircases for each image. The author collected five staircases for each images, leading to a total of 52 staircases per condition (208 in total).



**Figure 3.1.** Creation of stimuli. The chromatic and luminance channels are first separated, then blurred to create the alone conditions (A and B). They are then recombined with their sharp counterpart to created the combined conditions (C and D). Blur discrimination thresholds were measured for these four conditions. Original image from McGill Calibrated Colour Image Database (Olmos & Kingdom, 2004).



**Figure 3.2.** Sharp stimuli. The sharp images typically used as foils in each experiment, showing the colour manipulations that were made (A) Original images before any manipulation. (B) Images used in Experiment 3.1, cropped to 512 x 512 pixels and presented at 50% contrast to allow conversions in colour space without exceeding the gamut of the monitor. (C) Examples of images used in Experiment 3.2, cropped to 512 x 512 pixels, luminance and chromatic information is presented at five times the participant's corresponding contrast detection thresholds (presented images normalised for participant AP). (D) Images used in Experiment 3.3, cropped to 512 x 512, luminance and chromatic information has been swapped in MB-DKL space, i.e., luminance information has been replaced with colour information and vice versa and presented at 50% contrast. Original images from McGill Calibrated Colour Image Database (Olmos & Kingdom, 2004).

### Data Analysis

Participants' responses were averaged for each blur intensity level presented in the staircase procedure. A Weibull function was then fit to this data to determine the threshold at the 80% correct point (See Section 2.10. for details).

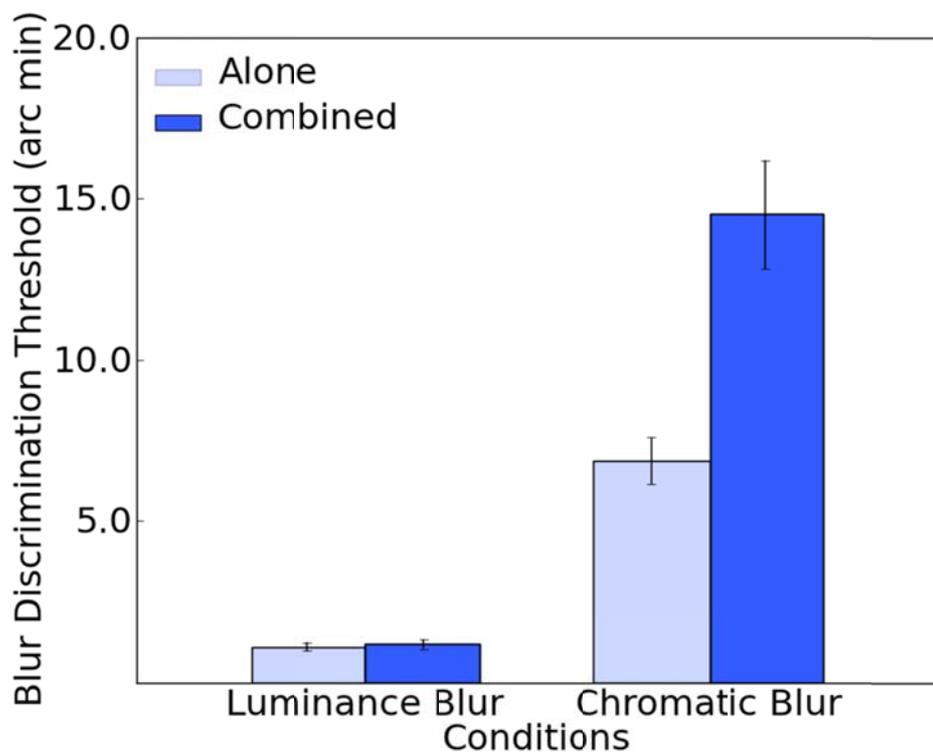
#### 3.1.2. Results

The group data are shown in Figure 3.3. A two-way ANOVA showed that observers had higher blur discrimination thresholds for chromatic than for luminance information (main effect of channel type;  $F_{(1, 76)} = 95.664$ ,  $p < 0.001$ ,  $MS_{\text{channel}} = 1300.679$ ). Critically, the elevated thresholds for chromatic blur were more pronounced in the presence of sharp luminance information (interaction between channel and combination;  $F_{(1, 76)} = 14.548$ ,  $p < 0.001$ ,  $MS_{\text{interaction}} = 197.804$ ).

Whilst the blur discrimination thresholds for the isoluminant stimuli are higher than for the luminance conditions, the thresholds when blurred chromatic information is combined with sharp luminance information were significantly higher again. For luminance defined blur, on the other hand, the presence of sharp chromatic information had no masking effect.

Lower acuity, potentially caused by the relative sparsity of S-cones (Wald, 1967), or the low-pass nature of colour vision (Mullen, 1985; Parraga, et al., 1998), may explain the generally higher thresholds for chromatic blur detection. If these factors were the source of the specific masking effect we found, there would be no difference in the blur discrimination thresholds of

the ‘chromatic blur alone’ condition and the ‘chromatic blur combined with sharp luminance’ information condition. However, blur discrimination thresholds are increased by the presence of sharp luminance information. Therefore the increase in blur discrimination thresholds caused by the introduction of sharp luminance information cannot be explained by poorer acuity of colour vision.



**Figure 3.3.** The mean blur discrimination threshold for both combined (dark blue columns) and single channel (light blue columns) across the group. Error bars represent  $\pm 1$  SEM.

### 3.2. Equating effective contrast

In natural scenes, luminance information has higher effective contrast than chromatic information (Rivest & Cavanagh, 1996) and this may cause it

to be a more effective mask. To test whether this explains the effect found in Experiment 3.1 the contrast of the channels was equated according to individual observers' discrimination thresholds.

### 3.2.1. Methods

#### *Participants*

The same participants were used as for Experiment 3.1.

#### *Stimulus Generation*

The stimuli were initially generated in the same manner as for Experiment 3.1. In addition, discrimination thresholds were measured for the luminance and the combined isoluminant channels for each participant for each image using a 2IFC task. The contrast was varied using a one-up, three-down staircase procedures and the contrast detection thresholds was extracted by fitting a Weibull function to the data from these staircases.

Rather than scaling the contrast of each channel by a uniform amount (50%) as in 3.1, the channels were each scaled independently for every image and every observer to a contrast that was five times the corresponding detection threshold for that stimulus component (Figure 3.2C).

#### *Procedure*

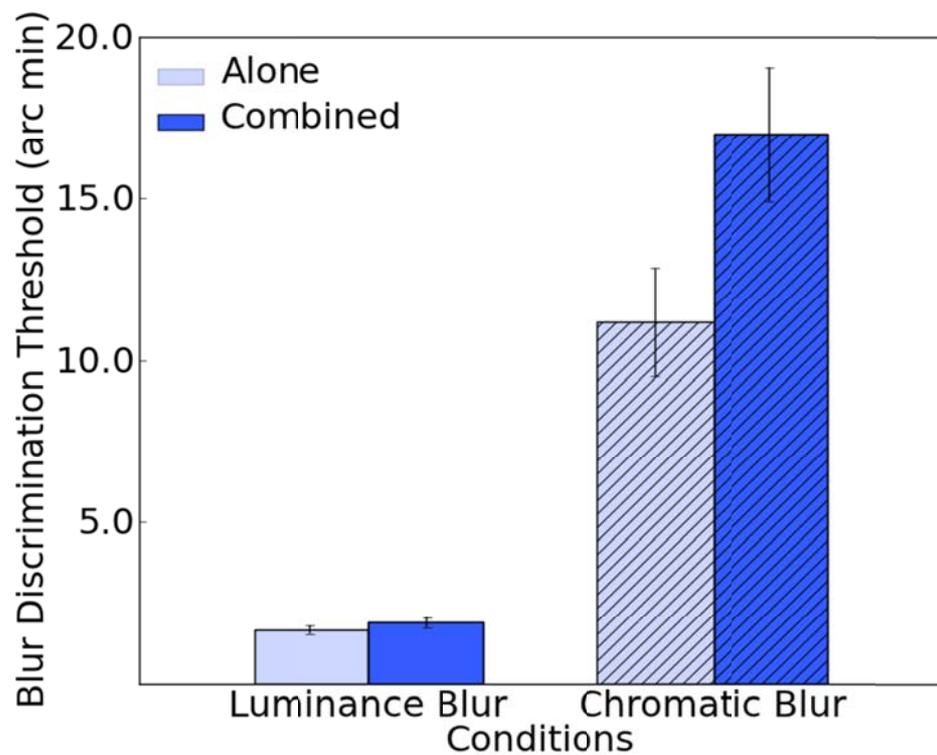
The same procedure was used as for Experiment 3.1.

### *Data Analysis*

The same data analysis was used as for Experiment 3.1. However, as a result of the lower overall contrast 13 (6.25%) staircases had to be excluded as they did not converge; three from the isoluminant condition and 10 from the blurred chromatic information combined with the sharp luminance information condition.

#### 3.2.2. Results

The main effect of channel ( $F_{(1, 67)} = 103.112$ ,  $p < 0.001$ ,  $MS_{\text{channel}} = 1503.064$ , and interaction between channel and combination ( $F_{(1, 67)} = 14.985$ ,  $p < 0.001$ ,  $MS_{\text{interaction}} = 218.418$ , were entirely undiminished (Figure 3.4); the luminance advantage is not caused by the higher effective contrast of luminance information in natural scenes.



**Figure 3.4.** The mean blur discrimination threshold, for contrast equated stimuli, for both combined (dark blue columns) and single channel (light blue columns) across the group, error bars represent  $\pm 1$  SEM. Hatched columns denote that some staircases were excluded from the analysis as they did not converge, see method for details.

### 3.3. Controlling for statistical regularities

Differences in statistical regularities between the channels may also have been the source of the original effect. For example, the appearance of high spatial frequency stimuli is more affected by blur than low spatial frequency stimuli. Therefore, if one channel contains more high spatial frequency information it may be easier to detect blur in that channel. To test if this was influencing the effect we swapped the colour and luminance channels, such that chromatic changes in the image became luminance

changes and vice versa. If the effect were caused by any difference in the statistics of the information in these natural scenes the effect should also be reversed, causing luminance blur to be masked by sharp chromatic information.

### 3.3.1. Method

#### *Participants*

Ten volunteers, aged between 18 and 29, who had not participated in previous studies, with the exception of the author, took part in this study.

#### *Apparatus*

A 22-in Vision Master Pro 513 (Iiyama) was used, running at 1280 x 1024, with an 85 Hz refresh rate.

#### *Stimulus Generation*

Stimuli were generated in the same manner as for Experiment 3.1. However, after conversion into MB-DKL space (Derrington, et al., 1984; Macleod & Boynton, 1979), the information in the LM and S channels was replaced with the luminance information and information in the luminance channel was replaced with half of the sum of the LM and S information (Figure 3.2D).

### *Procedure*

The same procedure was used as for Experiment 3.1. Each participant collected two staircases for each condition, leading to a total of 80 staircases per condition (320 staircases in total).

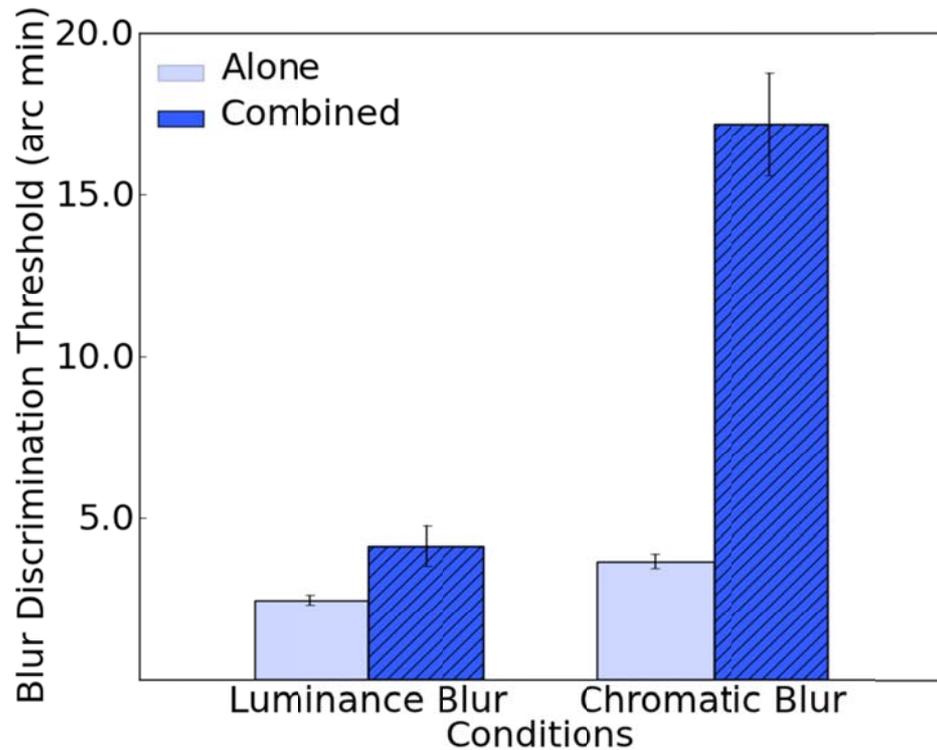
### *Data Analysis*

The method of averaging data and fitting a Weibull function could not be performed for all data in this set due to the poor performance levels in the condition combining blurred colour and sharp luminance information. For this reason, the simpler method of averaging the final six reversals from the staircase was used (see Section 2.11). Even then, 31 (9.69%) staircases had to be excluded from the analysis because the subjects' performance was so poor that the staircases did not converge. Of these, four came from the sharp chromatic information combined with blurred luminance condition and 27 came from the sharp luminance information combined with blurred chromatic information condition.

#### **3.3.2. Results**

The difference in blur thresholds between the chromatic- and luminance-only conditions was substantially reduced (Figure 3.5), to the point that it was no longer statistically significant (Fisher's least significant difference,  $p = 0.448$ ). However, chromatic blur thresholds remained poor in the presence of sharp luminance information, interaction between channel and combination ( $F_{(1, 285)} = 83.743$ ,  $p < 0.0001$ ,  $MS_{\text{interaction}} = 1932.375$ ). Clearly

the effect is not caused by differences in the information contained within the chromatic and luminance channels.



**Figure 3.5.** The mean blur discrimination threshold, for channel reversed stimuli, for both combined (dark blue columns) and single channel (light blue columns) across the group, error bars represent  $\pm 1$  SEM. Hatched columns denote that some staircases were excluded from the analysis as they did not converge, see method for details.

### 3.4. Ruling out luminance artefacts

In Experiments 3.1-3.3 isoluminance was determined photometrically; based on measurements from a photometer rather than behavioural measures. This could potentially have led to luminance artefacts in the colour channels due to individual observers having slightly different isoluminant planes (Krauskopf, Wu, & Farell, 1996). To address this issue Experiment 3.1

was repeated with isoluminance determined psychophysically and with a deliberately introduced luminance artefact. If the effect had been caused by luminance artefacts then it would be reduced when we control for the individual subject's isoluminant plane and it would be increased when we add a large artificial artefact.

### 3.4.1. Method

#### *Participants*

One male participant, aged 26, took part in the study, who was naive to the purposes of the study.

#### *Apparatus*

The same apparatus was used as for Experiment 3.3.

#### *Stimulus Generation*

The participant's psychophysical isoluminant axis was measured for the L-M and S-cone channels separately using a minimum motion procedure as described in Section 2.6.

#### *Procedure*

The same procedure was used as for Experiment 3.1. The experiment was run twice; once with the elevation values determined by the motion nulling procedure and once with a luminance artefact of 5.0° of elevation.

### *Data Analysis*

The same data analysis was used as for Experiment 3.1. Data previously collected for Experiment 3.1 were used for the photometrically determined isoluminant condition. One staircase was excluded from the sharp luminance information combined with blurred chromatic information condition in the photometric isoluminance data and one was excluded from the same condition in the luminance artefact data because they did not converge.

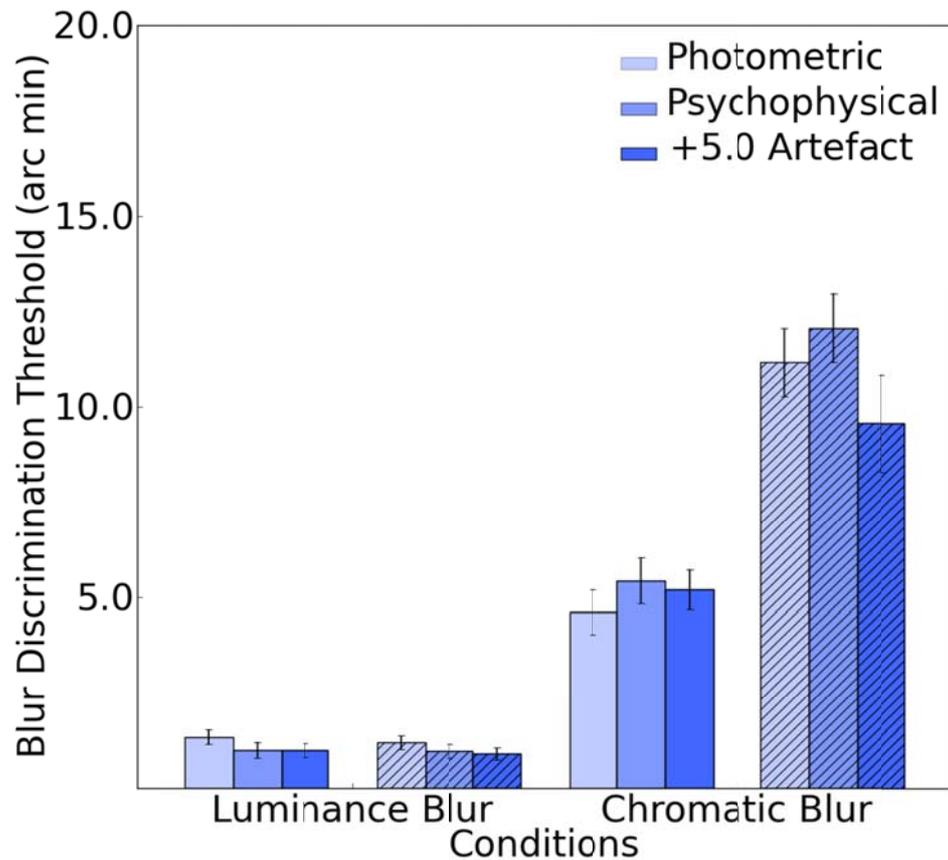
5000 within-subject bootstrap resamples were taken for each condition to produce a new set of psychometric curves. These data were then used to calculate the standard error for each condition.

#### **3.4.2. Results**

The motion nulling procedure revealed small differences between photometric and psychophysical isoluminance for the participant. Psychophysical isoluminance was at  $-0.185^\circ$  of elevation for the LM channel and  $-2.767^\circ$  of elevation for the S channel.

Stimuli generated using psychophysically determined isoluminance had slightly increased thresholds for both the isoluminant blur and chromatic blur combined with sharp luminance information conditions. Introducing a luminance artefact also slightly increased the blur threshold for isoluminant stimuli, but importantly reduced the threshold of the chromatic blur combined with sharp luminance conditions (from  $11.182^\circ$  in the photometric isoluminant condition to  $9.580^\circ$  in the luminance artefact condition, see

Figure 3.6). This demonstrates that any luminance artefacts that may have been present in the stimuli would only serve to decrease the luminance masking effect.



**Figure 3.6.** Experiment 3.4 Results. The mean blur discrimination thresholds for combined (hatched columns) and single channel for the three elevations; photometric isoluminance or 0° of elevation (light blue columns), psychophysical isoluminance or -0.185° of LM elevation and -2.767° of S elevation (mid blue columns) and 5° of elevation (dark blue columns). One staircase was excluded from each of the chromatic blur combined with sharp luminance conditions due to lack of convergence. Error bars represent  $\pm 1$  SEM.

### 3.5. Discussion

A number of previous demonstrations have suggested that luminance information is of particular importance in the detection of visual edges. Here

we quantified that dominance using a blur-discrimination task with naturalistic stimuli and tested a number of candidate explanations for it, namely whether the effect could be explained by poorer chromatic acuity, lower effective contrast, or differences in scene statistics. We found that none of these factors were able to explain the fact that subjects were unable to detect chromatic blur in the presence of sharp luminance information.

First we showed that differences in acuity are not sufficient to explain the data. Subjects *were* generally worse at detecting blur in the isoluminant stimuli which might be ascribed to poorer chromatic acuity, but they were *very much* worse at the task only when sharp luminance information was combined with the chromatic blur. Even in Experiment 3.3, for which the modifications to the images resulted in equal blur discrimination thresholds for isoluminant stimuli and achromatic stimuli, when the information was combined the chromatic blur became imperceptible.

Second, we demonstrated that the effect is not due to the higher effective contrast of luminance information in natural scenes; equating the effective contrast of the channels did not diminish the effect.

Third, the effect is not caused by differences in the statistical structure of the colour and luminance information; reversing the channels, and therefore the statistical properties of the luminance and chromatic information did not cause the effect to be reversed or even reduced.

The fact that chromatic blur alone is harder to detect than luminance blur alone is entirely consistent with previous findings. For instance, studies

have shown that blur thresholds for S-cone isolating stimuli are approximately twice as high as those for the other two channels even when cone contrast is taken into consideration (Wuerger, et al., 2000; Wuerger, et al., 2001). This may be due to reduced spatial sampling of chromatic information leading to a lower precision in chromatic processing (Peirce, et al., 2008). This reduced sampling may, in turn, be a consequence of chromatic aberration; the visual hardware may reflect the lack of spatial precision in the chromatic signals themselves (R. L. De Valois & De Valois, 1988). As a result, luminance may be used for tasks requiring high spatial precision. Conversely, colour may be used predominantly to process surface properties and to facilitate segmentation and grouping, with only a secondary role in edge detection and localisation (Mollon, 1989). If colour is mainly used to process surface properties this could explain why it appears to be discounted as a cue to edge perception when luminance information is present.

It is surprising that equating the effective contrast of the colour and luminance channels did not reduce the effect. Rivest and Cavanagh (1996) found that luminance does not play a privileged role in a contour localisation task if the luminance and chromatic channels are equated to have similar localisation thresholds when presented alone. Those authors suggested that the reason luminance appears privileged in natural scenes is due to its greater effective contrast which, at least for the perception of blur, appears not to be the case.

Colour information and luminance information in natural scenes are statistically similar in their  $1/f$  amplitude spectra (Parraga, et al., 1998) and in the numbers of achromatic and isoluminant edge that they contain (Hansen & Gegenfurtner, 2009). There might, however, be other statistical differences between the chromatic and luminance information in natural scenes, for example, in the fine structure. Even if natural scenes are not different in general, it might have been the case that the particular images used in this chapter had different image statistics in the two channels. To ensure that no such statistical artefacts could have caused the effects measured we swapped the information in the luminance and chromatic channels and repeated the experiment. The fact that this removed the advantage for the luminance channel presented alone indicates that there may have been some effect of differential statistics. However, these differences were clearly not responsible for the luminance dominance; when the reversed channels were combined subjects still gave preference to the luminance channel, even though it now contained no more information than the chromatic channel. Therefore the dominance of sharp luminance information over blurred chromatic information is not related to the statistical structure of natural scenes. At this point the evidence appears to indicate a mechanism giving active preference to luminance signals in the discrimination of blur.

It is clear from these data that the signals from chromatic and luminance information are not combined in a simple linear fashion such that it is not sufficient to consider either chromatic or luminance cues in isolation. In the current study we would not have been able to predict the masking

effect caused by combining blurred chromatic information and sharp luminance information from either the achromatic or isoluminant conditions. The masking effect could only be revealed by testing colour and luminance information in combination.

Similarly, the phase of a luminance grating overlaid on a chromatic plaid changes the appearance of the plaid (Kingdom, 2003). If the luminance grating is out of phase the plaid has a three-dimensional appearance (an example of the shape-from-shading effect). However, if the luminance grating is in phase with the chromatic information the impression of depth is suppressed.

The masking effect could indicate that chromatic blur is being bounded by the sharp luminance information, i.e. the chromatic blur does not appear to cross luminance boundaries. When reticles (thin, low-contrast, achromatic lines) are superimposed on the zero crossings of isoluminant gratings this can improve chromatic contrast sensitivity (Montag, 1997). This could be another circumstance where a chromatic gradient is bounded by luminance information. The facilitation effect caused by the reticles may be at the expense of spatial acuity of the chromatic information i.e. the chromatic information becomes tied to the luminance information (see Chapter 6 for further details). This would mean that the chromatic information would appear aligned with the luminance edges, as seen in the Boynton Illusion (Kaiser, 1996) and the results in this chapter.

There are existing accounts of edge detection such as scale space models (Georgeson, May, Freeman, & Hesse, 2007) and relative phase models (Burr, Morrone, & Spinelli, 1989). However, these do not currently attempt to incorporate the multiple channels (chromatic and luminance information) that would be necessary to model the current data.

In conclusion, the data in this chapter show that the process of combining luminance and chromatic signals is not simple linear summation. When chromatic blur is combined with sharp luminance information, chromatic blur discrimination thresholds are significantly poorer than when presented alone. The converse effect does not occur; blurred luminance information cannot be masked by sharp chromatic information. The luminance masking effect is not caused by poor acuity in the colour channels, higher contrast of luminance information or differences in the statistical properties of the information provided to each channel. This indicates an underlying mechanism that gives precedence to luminance edge information even when more precise chromatic information is available.

#### 4. Cue combination of colour and luminance in aligned synthetic edges

In this chapter synthetic edges will be used to investigate whether edge localisation is improved when both colour and luminance cues are present, compared to either in isolation. Observed data will be compared with predictions based on measurements of each cue alone. In the previous chapter, we used natural scene stimuli but, this limits the way that stimuli can be manipulated and so, in this chapter we will be using bipartite edges. This will allow us to investigate a) whether having both colour and luminance cues present improves localisation judgements and b) how the two cues might be combined.

Gur and Akri (1992) suggest that colour vision evolved not only to encode colour, but to enhance luminance processing. In support of this idea, it was found that the ability to discriminate between a circle and an ellipse is enhanced when both colour and luminance information is present, compared to either alone (Syrkin & Gur, 1997). However, contrast detection tasks show asymmetric facilitation between colour and luminance (for further details see Section 1.4). Chromatic pedestals do not facilitate detection of luminance targets (K. K. De Valois & Switkes, 1983), but luminance pedestals do facilitate detection of chromatic targets (Cole, et al., 1990; Switkes, et al., 1988). These results suggest that performance is improved by the presence of both cues for some tasks. It is not clear whether improvement will occur in an edge localisation task and, if it does, how much improvement will occur.

It may be that chromatic information is simply discounted or a 'winner takes all' strategy is being used. In this case combining the cues would not improve performance beyond the most reliable cue available. If information from both cues is combined, performance may be predicted using an unweighted averaging model, where each cue has equal influence. Alternatively a weighted averaging model, such as maximum likelihood estimation (MLE), where the cues have different amounts of influence may predict performance. See Section 1.2 for full descriptions of all three models.

One way to quantify performance is to measure variability. Here we will use a staircase procedure to generate a psychometric function and the corresponding just noticeable difference (JND). The JND is a measure of variability; the greater the spread of responses the larger the JND. Increased variability is indicative of poor performance. If a participant is performing well at a localisation task they are more likely to give consistently similar responses, if they are bad at a task there will be less consistency in their responses. JNDs will be measured for the cues in isolation and then in combination. The combined JNDs will then be compared to model predictions generated from the individual components.

In this chapter Vernier acuity (alignment) tasks will be used to investigate how colour and luminance cues are combined in an edge localisation task. The human visual system is very good at making Vernier judgements with both colour- and luminance-defined stimuli (Krauskopf & Forte, 2002) and so there is very little variability in performance in these types

of judgements. In order to test the efficacy of the three models we need to increase variation in performance and prevent ceiling effects. In particular, if there is no significant difference in performance between isoluminant and achromatic conditions there will be less difference between model predictions. For example, if both the achromatic and isoluminant conditions have a JND of 1 arc min, both ‘winner takes all’ and unweighted averaging models will predict a JND of 1 arc min (MLE would predict a JND of 0.5 arc min).

#### 4.1. The effect of Gaussian white noise

This preliminary experiment aims to test whether JNDs generated by a staircase procedure can be used to represent performance in an edge localisation task. In order to test the models fully we need to be able to control performance for the cues in isolation and generate a variety of predictions. One way that variability could be increased, and performance controlled, is by introducing Gaussian white noise.

##### 4.1.1. Methods

###### *Participants*

Two males and one female volunteer (including the author), aged between 25 and 29, participated in the study. Two participants (both male) were naive to the purposes of the study.

### *Apparatus*

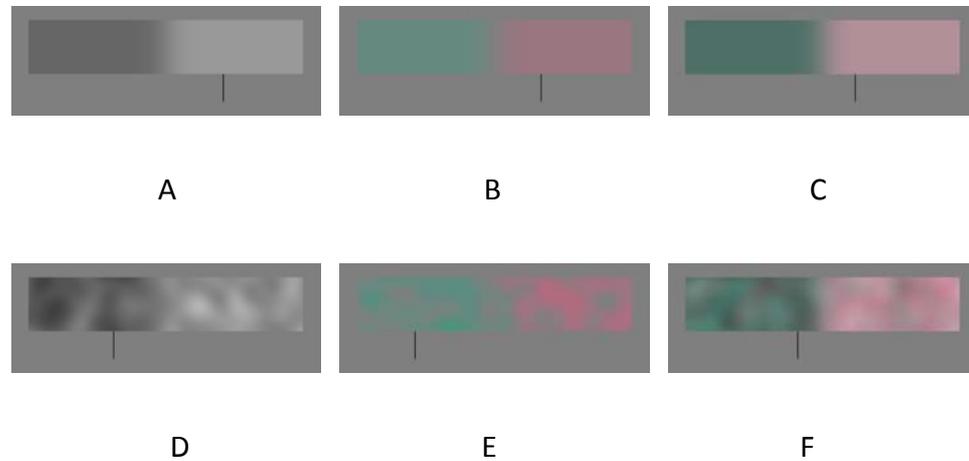
A chin rest was used to ensure participants viewed the stimuli from a constant 114cm distance, giving a viewable area that subtended  $17.95^\circ$  of visual angle.

### *Stimulus Generation*

Three bipartite edges were created in MB-DKL space (Derrington, et al., 1984; Macleod & Boynton, 1979) and presented at photometric isoluminance, one for each channel (L+M, L-M and S-(L+M)). In this instance a bipartite edge refers to a transition from one contrast polarity to the other, see Figure 4.1 for examples. The edges were Gaussian blurred ( $\sigma = 1^\circ$ ). The edges were  $10^\circ \times 2^\circ$  in size and were presented with a neutral grey background. In order to create the combined conditions (LM + Lum, S + Lum and LM + S) the relevant component edges were summed together. The single edges were presented at a Michelson contrast of 0.1, meaning that the combined edges were presented at a Michelson contrast of 0.2. A vertical marker was presented immediately below the edge. The position of the marker was randomised for each trial.

The Gaussian filter applied to the white noise had a standard deviation of  $0.1^\circ$  and the noise was presented at a Michelson contrast of 0.2. In the combined conditions white noise for both channels was added together leading to a total noise contrast of 0.4. In order to reduce loading time for the stimuli, 100 noise patterns were pre-generated. The noise pattern for each trial in the staircase was randomly selected from these patterns, such that

participants could not learn a particular pattern. Noise was not correlated across channels. See Figure 4.1 for example stimuli.



**Figure 4.1.** Example stimuli. A and D are the luminance alone conditions, B and E are the L-M alone conditions, C and F are the combined conditions. A-C are the stimuli without noise. D-F are the stimuli with noise.

### *Procedure*

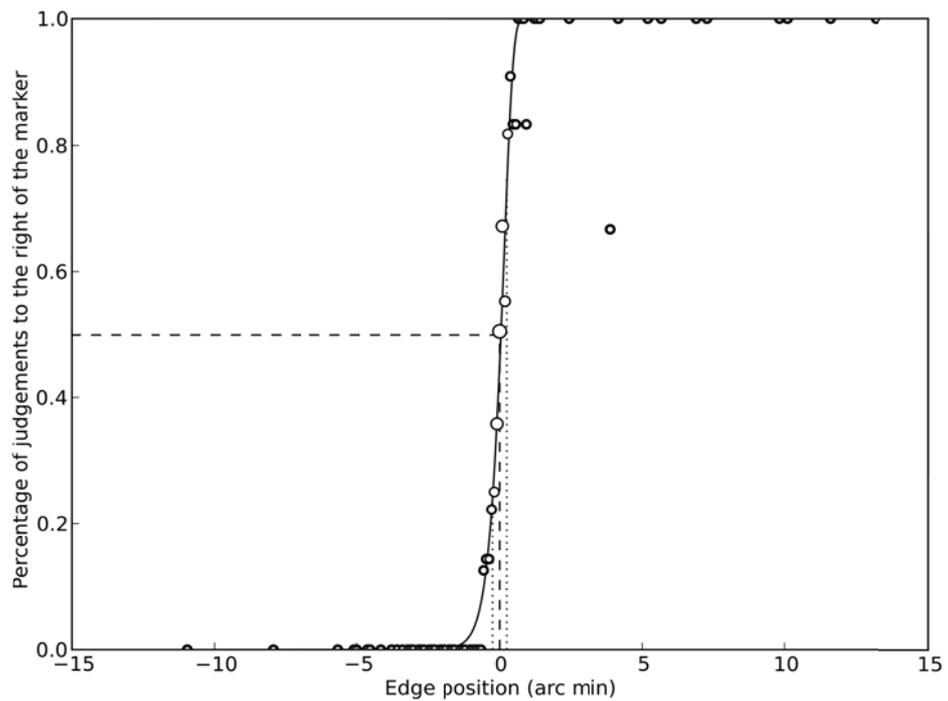
A two alternative forced choice (2AFC) design was employed. Participants were presented with the edge and marker for 300ms and asked whether the edge appeared to the left or right of the marker. Their response was followed by a 300ms inter-stimulus interval (ISI) before the start of the next trial.

The offset between the edge and marker was controlled by one-up, one-down staircase procedures, designed to converge on the point where participants were equally likely to judge the edge to be on the left or right of the marker. Each participant collected 10 staircases of 50 trials for each

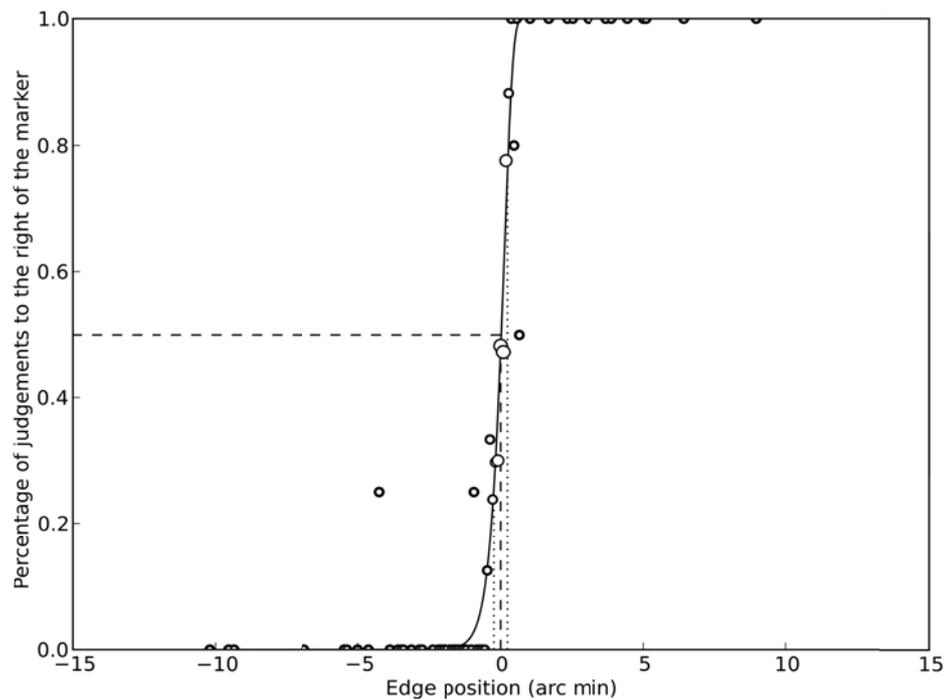
condition. The presentation order of the six conditions was randomised, but the staircases were not interleaved but run sequentially.

#### 4.1.2. Results

A logistic function was fit to the data and the just noticeable difference (JND) was calculated as the difference between the edge position where participants responded 'left' 75% of the time and the position where participants responded 'left' 25% of the time (Figure 4.2 and Figure 4.3). 5000 bootstrap resamples were taken for each condition, for each participant. Logistic functions were also fit to 5000 within-subject bootstrap resamples for each condition, for each participant. These were used to derive the standard error of the JND.

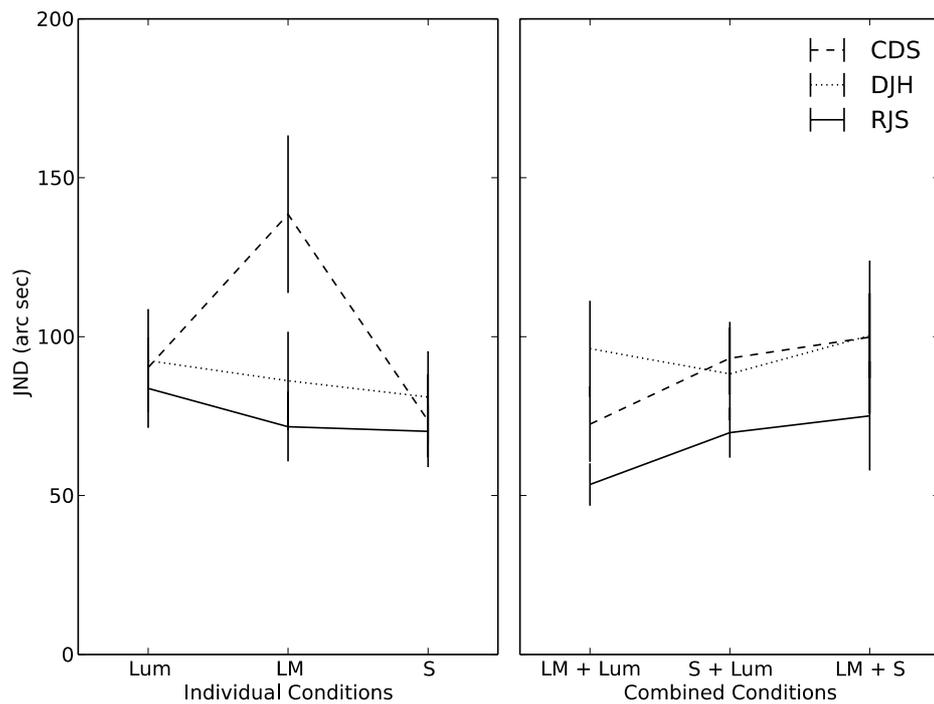


**Figure 4.2.** The psychometric function for one participant (DJH) for the achromatic edge presented at a Michelson contrast of 0.1. The solid line is the logistic fit and dashed lines represent the point at where the participant was equally likely to judge the edge to be on either side of the marker. The dotted lines represent the upper and lower bounds of the JND, the narrowness of this interval demonstrates the lack of variability in Vernier acuity judgements. Due to the fact that the data were collected according to a staircase procedure the number of trials at each level varies. The size of the data points is therefore being used to represent the number of trials at that level.



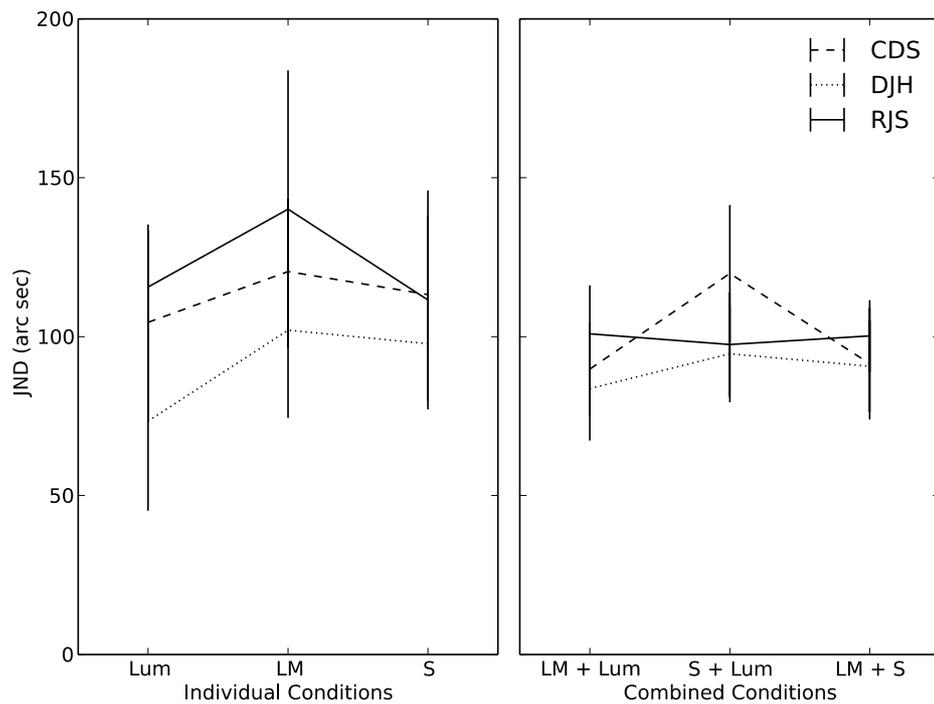
**Figure 4.3.** The psychometric function for one participant (DJH) for the isoluminant edge presented at a Michelson contrast of 0.1. The solid line is the logistic fit and dashed lines represent the point at where the participant was equally likely to judge the edge to be on either side of the marker. This demonstrates the lack of variability in Vernier acuity judgements. The dotted lines represent the upper and lower bounds of the JND, the narrowness of this interval demonstrates the lack of variability in Vernier acuity judgements.

The data do not show any systematic pattern across conditions whether noise is absent (Figure 4.4) or present (Figure 4.5). If the data were reflecting edge localisation sensitivity the L-M and S conditions would be expected to have larger JNDs than the luminance condition, however, this was not found to be the case.



**Figure 4.4.** JND values for all participants and all conditions with no noise present in the stimuli. The achromatic and isoluminant conditions alone are shown in the left panel and the combined conditions in the right panel. There is no systematic difference between any of the conditions. Error bars represent  $\pm 1$  standard error of the mean for each

The presence of noise was expected to decrease the JND values in general and prevent ceiling effects. However, as can be seen in Figure 4.5, noise had little impact on the performance of subjects in the task, and therefore cannot be used to manipulate performance.



**Figure 4.5.** JND values for all participants and all conditions with noise present in the stimuli. The achromatic and isoluminant conditions alone are shown in the left panel and the combined conditions in the right panel. There is no systematic difference between any of the conditions. Error bars represent  $\pm 1$  standard error of the mean.

#### 4.2. The effect of increasing viewing distance

In Experiment 4.1 the viewing distance was relatively short (114cm) and consequently the pixel size was quite large (68.21 arc sec). Note that participant JNDs are only slightly larger than this (Figure 4.4 and Figure 4.5). We therefore wanted to test whether pixel size was actually the limiting factor of the measurement, rather than the psychophysical performance of the subjects. In order to address this, the viewing distance was increased and the experiment repeated.

#### 4.2.1. Methods

##### *Participants*

Three male and four female volunteers (including the author), aged between 19 and 30, participated in the study. Six of the volunteers (three male) were naive to the purposes of the study.

##### *Apparatus*

A chin rest was used to ensure that participants viewed the stimuli from a constant 367cm distance, giving a viewable area that subtended  $5.62^\circ$  of visual angle. This decreased the pixel size to 21.19 arc seconds (compared to 68.21 in Experiment 4.1).

##### *Stimulus Generation*

Stimuli were generated in the same manner as for Experiment 4.1 and presented at a size of  $4.5^\circ \times 1^\circ$  with a Gaussian blur of  $0.135^\circ$ .

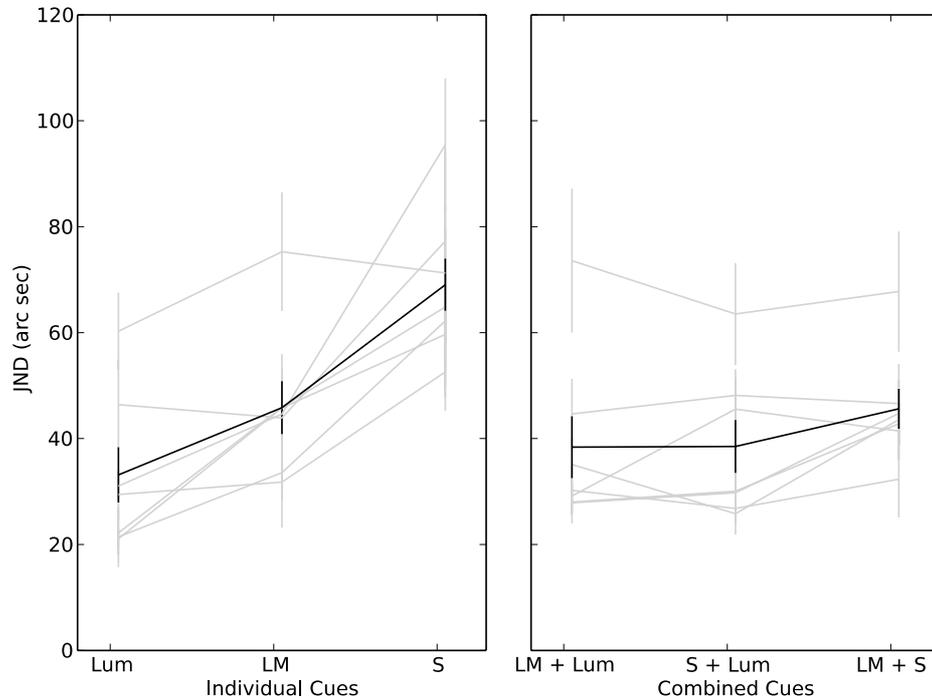
##### *Procedure*

The same procedure was used as for Experiment 4.1. In addition the luminance conditions were repeated with Michelson contrasts of 0.0225 and 0.02.

#### 4.2.2. Results

Data analysis was performed in the same manner as for Experiment 4.1. Systematic differences are clearly visible, with luminance having a smaller JND than either L-M or S conditions. This suggests that the task is now

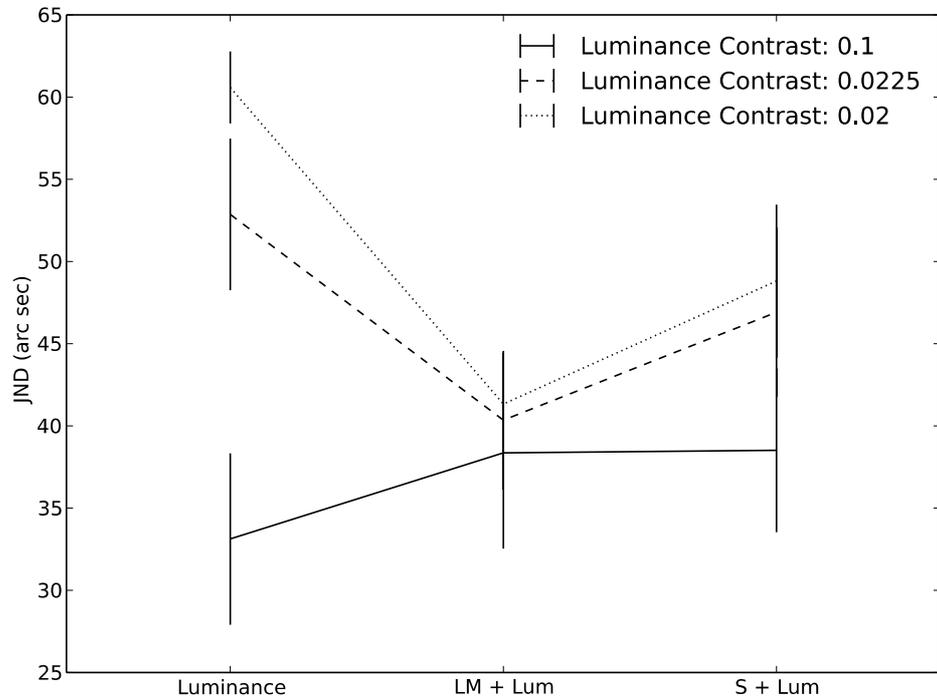
measuring edge localisation sensitivity and is no longer limited by pixel size (Figure 4.6).



**Figure 4.6.** JND values for Experiment 4.2 for all participants, shown in grey and combined across participants, shown in black. The achromatic and isoluminant conditions alone are shown in the left panel and the combined conditions in the right panel. All cues were presented at a contrast of 0.1. A clear systematic difference is shown between luminance, LM and S conditions. Error bars represent  $\pm 1$  standard error of the mean.

JNDs were measured for the three conditions that had a luminance component (luminance alone, luminance combined with L-M and luminance combined with S-(L+M)) at three Michelson contrasts (0.1, 0.0225 and 0.02). Reducing luminance contrast was found to increase the JND when the luminance edge was presented alone and for the S+Lum when luminance contrast was increased from 0.02 to 0.1 (Figure 4.7). Non-significant increases

were found in all other conditions tested. This demonstrates that contrast can be used to manipulate performance in this task.

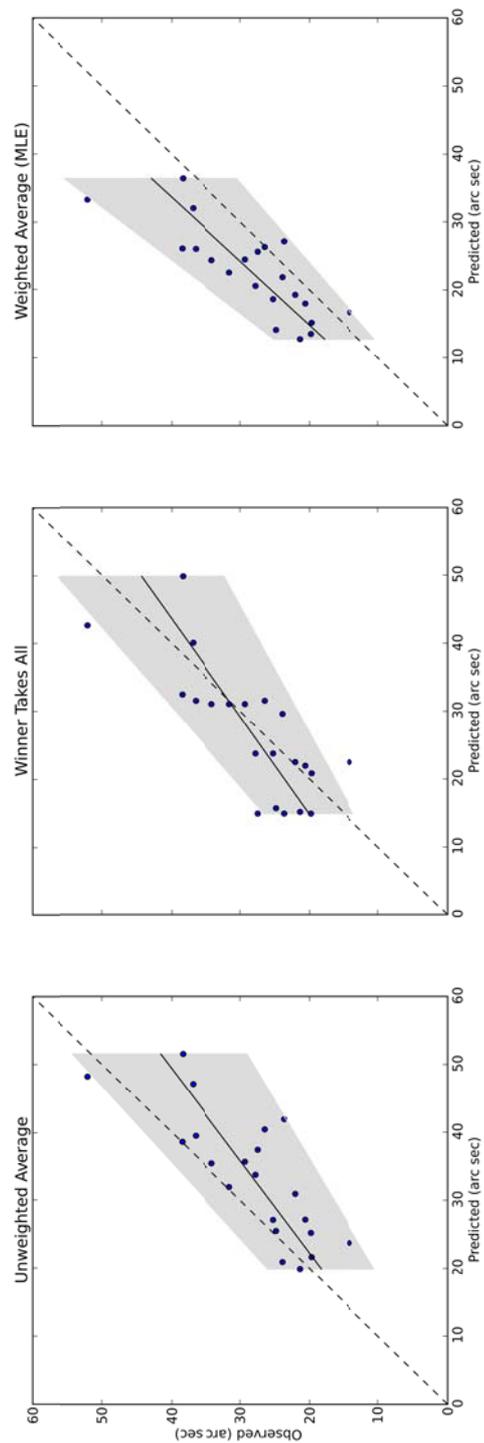


**Figure 4.7.** JND values collapsed across participants for the three luminance contrasts tested in Experiment 2. Decreasing the luminance contrast increases the JND for all conditions. Error bars represent  $\pm 1$  standard error of the mean.

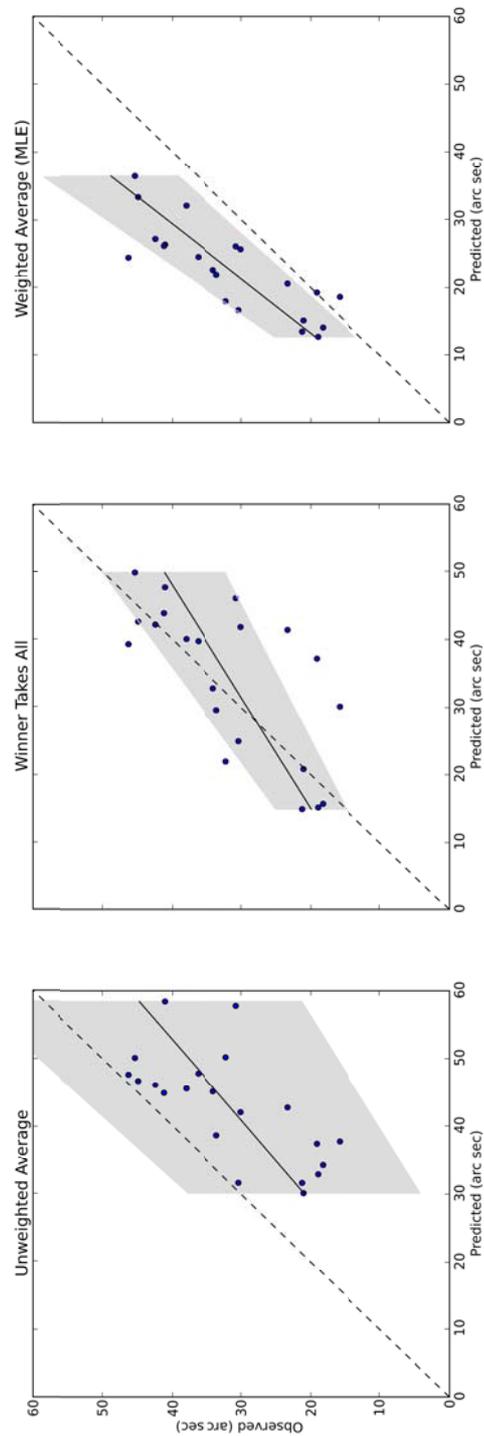
### Model Comparisons

There is no significant difference between the performance of the three models. There are large errors and, more importantly, the model predictions are not sufficiently different (Figure 4.8). Together these factors mean that no conclusions can be drawn. It appears that this methodology is not sensitive enough to differentiate between the three models.

It should be also noted that, whilst the data for S-(L+M)-combined-with-luminance is similar to those for L-M-combined-with-luminance, the errors are even larger and there is less consistency across participants. For this reason future experiments will be focused on modelling performance for L-M combined with luminance information.



**Figure 4.8.** LM + Lum model comparisons. Values are presented for the three luminance contrast conditions and all participants. Predictions are not significantly different from observed behaviour for any of the three models tested. Error bars, shown by the grey areas, represent  $\pm 1$  standard error of the mean. The dashed line represents unity and the solid line represents best fit for a linear regression.



**Figure 4.9.** S + Lum model comparisons. Values are presented for the three luminance contrast conditions and all participants. Predictions are not significantly different from observed behaviour for any of the three models tested. Error bars, shown by the grey areas, represent  $\pm 1$  standard error of the mean. The dashed line represents unity and the solid line represents best fit for a linear regression.

### 4.3. Discussion

In this chapter it has been demonstrated that: Gaussian white noise does not affect performance in the edge localisation task, a long viewing distance is necessary to prevent ceiling effects from pixel size and that contrast can be used to manipulate participants' performance. However, despite the improvements from using a longer viewing distance the staircase procedure was not sufficiently sensitive to differentiate between weighted averaging, unweighted averaging or 'winner takes all' strategies.

Experiment 4.1 shows that Gaussian white noise has little impact on participants' Vernier judgements and cannot be used to manipulate performance. This is surprising as, generally, noise limits perception causing performance to decrease (Pelli & Farell, 1999). It has been suggested that the global features of a stimulus are used in Vernier judgements; line-feature primitives are extracted before localisation takes place (Meer & Zeevi, 1986). Meer and Zeevi (1986) created a Vernier task using dots to create a vertical line, they then perturbed this line of dots into a Gaussian distribution. They found that perturbation only increased Vernier thresholds by a small amount, far less than they had predicted. They attributed this lack of effect to global information being used to overcome interference in the local information. This could explain why the introduction of Gaussian white noise had no effect on performance in Experiment 4.1. If the global form of the line was extracted prior to localisation this would have attenuated the effect of noise and

prevented deficits in performance. This also suggests that the use of a different type of noise e.g.  $1/f$  would have no effect on thresholds.

Experiment 4.2 shows that increasing contrast reduces the JNDs for edge localisation judgements and so can be used to modulate performance. It is, however, unclear whether performance will continue to improve as contrast increases or whether it will plateau. Orientation and spatial-frequency discrimination thresholds can be reduced by increasing contrast, but only for a limited range of contrasts, after which no further decrease in threshold is observed (Skottun, Bradley, Sclar, Ohzawa, & Freeman, 1987). However, Vernier acuity may be affected by a larger range of contrasts. Krauskopf and Forte (1991) demonstrated that increasing contrast reduced Vernier offset thresholds, for both chromatic- and luminance-defined stimuli, to up to ~50 multiples of detection threshold.

The predictions generated by the three models were not sufficiently different to allow them to be compared. The differences in the JNDs for isoluminant and achromatic stimuli, introduced by varying contrast, were not sufficient to allow the three predictions to be distinguished. In the future, in order to distinguish between the three models, the performance difference between the chromatic and luminance edges in isolation must be increased.

## 5. Cue combination of conflicting colour and luminance edges

In the previous chapter we attempted to study the way that information about chromatic and luminance cues to edge location are combined when they *agree*, whether or not combining them enhances sensitivity, tested with a staircase procedure. In this chapter we will replicate that experiment using method of adjustment and also investigate how the two cues are combined when they *disagree*.

Replication of Experiment 4.2 will ensure that method of adjustment is suitable for measuring edge localisation performance. This is important as, unlike staircase procedures, method of adjustment allows analysis of the distribution of judgements as well as their variability. When the two cues conflict it may be that the visual system employs a form of the ‘winner takes all’ model, where the cue used to localise the edge changes between trials; sometimes luminance determines the perceived edge location and sometimes L-M. Switching between the two cues in this manner would lead to a bimodal distribution of edge localisation judgements. If we only measured localisation judgements using a staircase procedure we would not be able to detect this bimodality, but the use of method of adjustment will allow us to check for this possibility.

We will also investigate the range of contrasts necessary to increase the difference between the predictions from the three models. As discussed in the previous chapter, it is not clear what range of contrast values can be used to modulate performance in an edge localisation task. In order to

determine this, measurements will be taken for the achromatic and isoluminant stimuli over a range of contrasts and each contrast will be used to generate a new set of model predictions. The spread of contrasts that allow for the greatest range of model predictions will then be used when measuring edge localisation performance when the cues *agree* and when they *disagree*.

When colour and luminance cues conflict there are several possible ways that edge location could be determined as both chromatic and luminance information can be used to make edge localisation judgements. When contrast is equated in multiples of detection threshold, Vernier thresholds are not significantly different for isoluminant and achromatic stimuli (Krauskopf & Forte, 2002) and luminance is not privileged in edge localisation when performance is equated (Rivest & Cavanagh, 1996). Therefore, if the cues are equated, there is no inherent reason why one should have more influence than the other. Chapter 3 showed that chromatic blur is masked by sharp luminance information. However, luminance variation can be masked by chromatic information when the cues are orthogonal (Kingdom, et al., 2010). Therefore, either luminance or chromatic information can dominate depending on the circumstance. In the case of localisation of conflicting edges it is unclear whether the cues will exert equal influence or whether one will have more influence than the other.

In summary, Experiment 4.2 will be replicated with this alternative method. Specifically, we will check that performance on luminance, L-M and S-(L+M) conditions can be differentiated (Experiment 5.1). Measurements will

be taken for isoluminant and achromatic stimuli, in order to determine the best range of contrasts to allow for model differentiation when the cues agree (Experiment 5.2). Those stimulus contrasts will then be used, in a separate experiment, measuring edge localisation performance when the colour and luminance cues *agree* (Experiment 5.3). The same range of contrasts will then be used to investigate edge localisation judgements when colour and luminance cues *disagree* (Experiment 5.4).

### 5.1. Piloting method of adjustment

This experiment tests whether the results found in Experiment 4.2, using a staircase procedure, can be replicated using method of adjustment. In particular we are aiming to determine whether method of adjustment is sensitive enough to differentiate between the achromatic and isoluminant conditions and whether contrast can be used to modulate performance. If the results of Experiment 4.2 are replicated this will allow us to use method of adjustment to investigate how colour and luminance edges are combined when they *agree* and this may improve differentiation between model performance. More importantly, it will allow us to use method of adjustment to investigate how colour and luminance edges are combined when they *disagree*; particularly, allowing us to consider the possibility of a bimodal distribution of responses.

### 5.1.1. Methods

#### *Participant*

The author, aged 30, participated in this study.

#### *Apparatus*

A chin rest was used to ensure that the participant viewed the stimuli from a constant 367cm distance, giving a viewable area that subtended 5.62° of visual angle. This viewing distance was used for the remainder of the experiments in this chapter.

#### *Stimulus Generation*

Two bipartite edges (L+M and L-M) were created in the same manner as for Experiment 4.1. These edges were Gaussian blurred ( $\sigma = 0.1^\circ$ ). Stimuli were 4.5° x 1° in size and presented with a neutral grey background. A vertical marker, with a width of one pixel, was presented immediately below the edge. The initial position of the marker was randomised for each trial.

The isoluminant edge was presented at a contrast of 0.1 and the achromatic edge was presented at Michelson contrasts of 0.1 and 0.02. The contrast of the combined conditions was the sum of the two component edges. There were three 'alone' conditions; L-M with a contrast of 0.1, luminance with a Michelson contrast of 0.1 and luminance with a contrast of 0.02 and subsequently two combined conditions.

### *Procedure*

The participant was presented with the edge and marker, and used the mouse to move the edge until they were satisfied that the two were aligned. There was no limit to the presentation time and the subject's response triggered the next trial, following a 300ms ISI. Presentation order of the conditions was randomised and 40 trials were collected per condition.

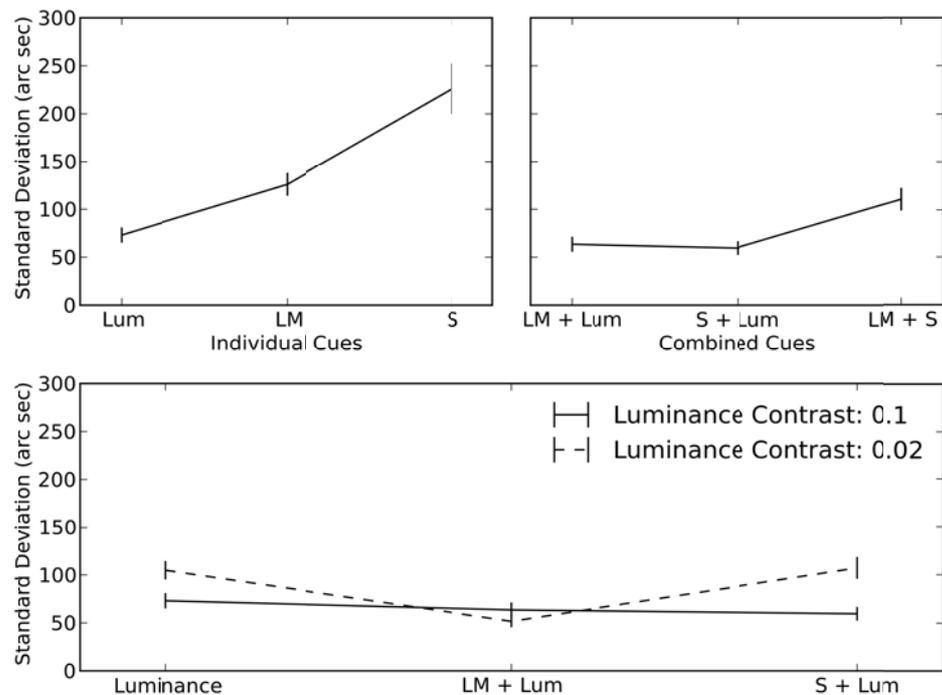
### *Data Analysis*

The absolute distance between the participant's judgement and the edge was calculated. When using method of adjustment there is a possibility of participants accidentally submitting a judgement before they are satisfied that the edge and marker are aligned. These 'mis-clicks' can skew the resulting distribution. Therefore, outliers were removed, defined as having a z-score greater than 3.0 or less than -3.0. This resulted in 7 trials (3.89%) being excluded. The standard deviation of the responses for each condition was calculated and used as a measure of performance; the more precise a participant's judgement the less variability there will be in their responses.

#### **5.1.2. Results**

The results replicate those found in Experiment 4.2, showing the systematic differences between channel and luminance contrasts (Figure 5.1). The exception to this is the luminance-combined-with-L-M condition when luminance had a Michelson contrast of 0.02. This condition had a lower variance than would be expected from Experiment 4.2; this may have been due to the small number of trials for each condition.

As the results of this experiment were largely consistent with those in Experiment 4.2 it was accepted that method of adjustment was suitable for use in future experiments.



**Figure 5.1.** Standard deviations for all conditions. The top left panel shows the alone conditions, when the contrasts were 0.1 and shows the clear differentiation between the conditions. The top right panel shows the combined conditions when the component contrasts were 0.1. The bottom panel shows the conditions with a luminance component and shows that for two of the three conditions reducing contrast increases the standard deviation. Error bars represent  $\pm$  standard error of the mean.

## 5.2. Equating performance using contrast

In order to understand how edge cues are combined it is important that the cues are equated so that each makes an equal contribution to the edge location. In addition, it is also important to maximise the difference

between model predictions. As shown previously, contrast can be used to increase or decrease the variance of edge judgements. In the following experiment, contrast was modulated in an effort to find the point where each cue was equally weighted, in accordance with MLE.

### 5.2.1. Methods

#### *Participants*

The author, aged 30, participated in this study.

#### *Stimulus Generation*

The stimuli were generated in the same manner as for Experiments 4.2, however, isoluminance was determined psychophysically (Section 2.6). ‘Flipped’ versions of the stimuli were also added in order to remove any side bias, leading to two luminance defined arrangements and two L-M defined arrangements.

#### *Procedure*

The same method-of-adjustment procedure was used as for Experiment 5.1; participants used the mouse to move the edge until they were satisfied that it was aligned with the marker. In an effort to find the point where performance was equivalent for the two cues measurements were taken for several luminance contrasts (0.02, 0.04 and 0.06 Michelson contrast) and L-M contrasts (0.1, 0.2, 0.3, 0.4, 0.6 and 0.9).

The aim of this experiment was to determine the contrast necessary to both equate the two cues, the point where they are predicted to make an

equal contribution to edge location, and maximise the difference between model predictions. Therefore, only achromatic and isoluminant conditions were tested; there were no combined conditions.

### *Data Analysis*

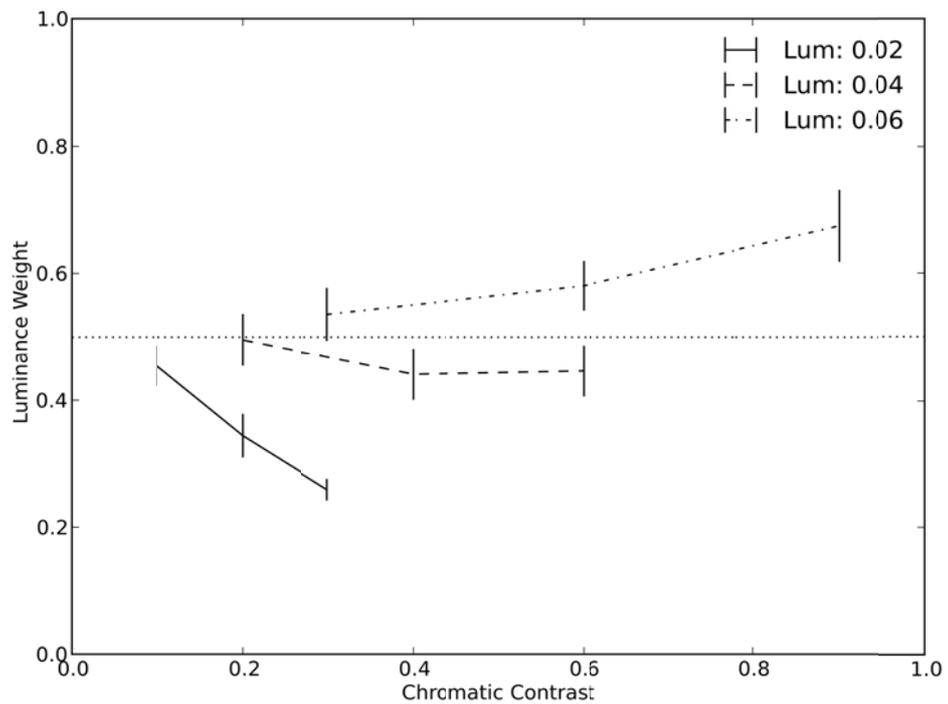
The same data analysis procedures were used as for Experiment 5.1. Furthermore, the mean difference between the flipped and non-flipped conditions was calculated and half of this was then added/subtracted from the raw values to remove any side bias.

In addition, the judgements from the achromatic and isoluminant edges were used to calculate the weights predicted by MLE. The following combinations were considered: luminance contrast of 0.02 in combination with L-M contrasts of 0.1, 0.2 and 0.3.; luminance contrast of 0.04 in combination with L-M contrasts of 0.2, 0.4 and 0.6; luminance contrast of 0.06 in combination with L-M contrasts of 0.3, 0.6 and 0.9.

#### **5.2.2. Results**

The predicted weights were closest to being equal when the luminance contrast was 0.04. However, modulations of L-M contrast had very little effect and the weights no longer varied; participant performance was approximately the same at this luminance contrast regardless of changes in the L-M contrast. The greatest range of weights occurred when luminance contrast was 0.02, although the weights were no longer equated (Figure 5.2).

The aims of this experiment were a) to determine the contrast values that would equate the cues and b) to maximise the difference between model predictions. However, in order to differentiate between weighted and unweighted averaging we must also test conditions where they are not equal; when the weights are the same, both produce the same prediction. If the weights are the same, they must both be 0.5 and this will predict that the edge will be perceived as equidistant between the two cues, this is the same as if the mean of the two edge locations was taken. Although the luminance contrast of 0.04 most closely equated the cues, it failed to generate differential model predictions. There was also very little variation in the predicted weights when the luminance contrast was 0.06. When the luminance contrast was 0.02 the weights were reasonably equated, but with greater variation in the predictions, allowing the models to be differentiated. Therefore, in the following experiments the luminance edge will be presented at a contrast of 0.02 and the L-M edges will be presented at contrast of 0.1, 0.2 and 0.3.



**Figure 5.2.** Weights for the following combinations; Luminance 0.02 and L-M 0.1, 0.2 and 0.3, Luminance 0.04 and L-M 0.2, 0.4 and 0.6 and Luminance 0.06 and L-M 0.3, 0.6 and 0.9. The dotted horizontal line represents the point where the two cues would have equal weights. Error bars represent  $\pm$  standard error of the mean.

### 5.3. Cue combination in aligned synthetic edges

Experiment 4.2 was replicated using method of adjustment and a more considered range of contrasts in an effort to differentiate between the three models: ‘winner takes all’, unweighted averaging and weighted averaging (MLE).

### 5.3.1. Methods

#### *Participants*

Three male and two female participants (including the author), aged between 20 and 31, participated in this study. Four of the participants (three male) were naive to the aims of the study.

#### *Stimulus Generation*

The component edges were the same as those used in Experiment 5.2; two luminance defined arrangements and two L-M defined arrangements ('flipped' and 'non-flipped'). In the combined conditions all possible combinations were presented to prevent bias: dark to light combined with red to green; dark to light combined with green to red; light to dark combined with red to green and light to dark combined with green to red.

#### *Procedure*

The same procedure was used as for Experiment 5.1. The participant was presented with the edge and marker and moved the edge, using the mouse, until they were satisfied that the two were aligned. There was no limit to the presentation time and there was a 300ms ISI between trials. Presentation order of the conditions was randomised and 40 trials were collected per condition.

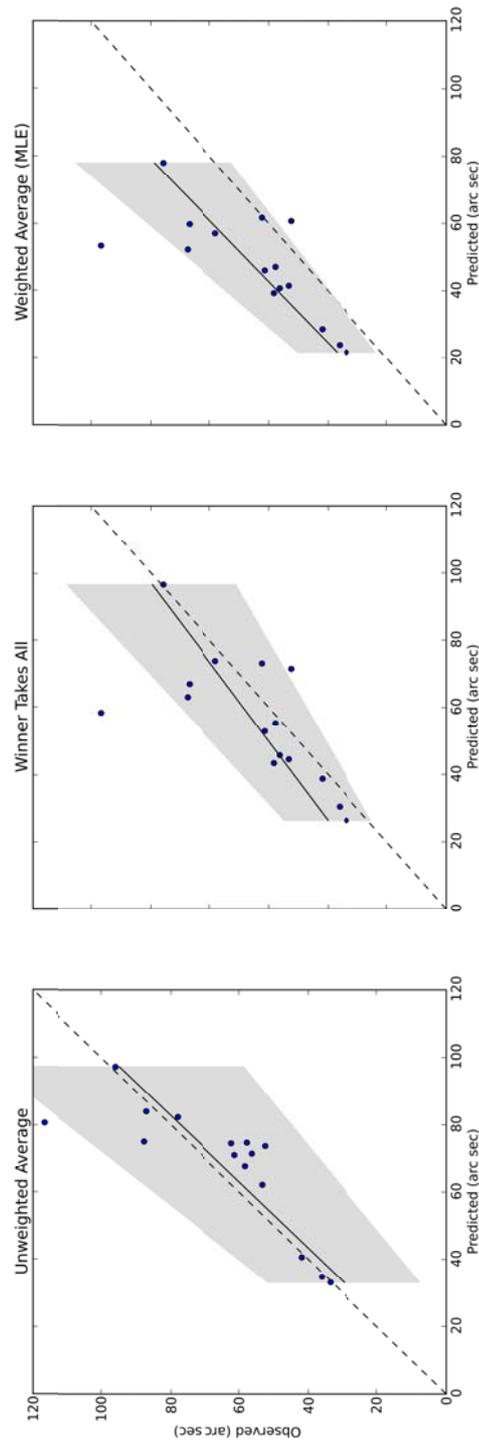
Luminance information was presented at a Michelson contrast of 0.02 and chromatic information at contrasts of 0.1, 0.2 and 0.3.

### *Data Analysis*

The data were analysed in the same manner as for Experiment 5.2. The standard deviation was measured for the component edges alone (luminance at a contrast of 0.02 and L-M at contrasts of 0.1, 0.2 and 0.3) and for the three combined edges (luminance contrast of 0.02 combined with each of the three L-M contrasts). The standard deviations from the component edges were used to calculate predictions for the three models.

#### 5.3.2. Results

The data are very similar to those recorded in Experiment 4.2 and there is no significant difference between the performance of the three models. As in that experiment there are large errors and the model predictions are not sufficiently different (Figure 5.3). There appears to be a trend away from weighted averaging (MLE), but this is not statistically significant. Once again no conclusions can be drawn as this methodology is not sensitive enough to differentiate between the three models.



**Figure 5.3.** Model comparisons. Values are collapsed across the three chromatic contrast conditions and across participants. Predictions are not significantly different from observed behaviour for any of the three models tested. Error bars represent  $\pm 1$  standard error of the mean. The dashed line represents unity and the solid line represents best fit for a linear regression.

## 5.4. Edge localisation with conflicting stimuli

It is not only important to understand how colour and luminance are combined when the edges are aligned but also when they conflict. In this experiment we will investigate participants' edge localisation judgements when the cues conflict and are not spatially aligned.

### 5.4.1. Method

#### *Participants*

Three male and two female volunteers (including the author), aged between 20 and 30, participated in the study. Four participants (three male) were naive to the aims of the study.

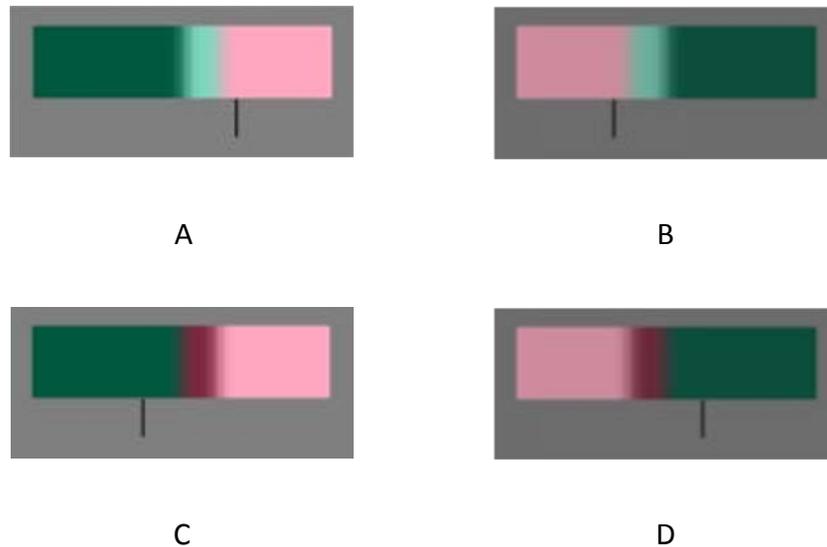
#### *Apparatus*

A chin rest was used to ensure that participants viewed the stimuli from a constant 367cm distance, giving a viewable area that subtended  $5.62^\circ$  of visual angle.

#### *Stimulus Generation*

The component edges were the same as those used in Experiment 5.3. In the combined conditions there was a gap of 3 arc min between the achromatic and L-M edges. There were four starting combined configurations; the achromatic edge central and the L-M edge a +3 arc min (Figure 5.4A) and the mirror of this (Figure 5.4B) and the L-M edge central and the achromatic

edge at  $+3$  arc min $^\circ$  (Figure 5.4C) and the mirror of this (Figure 5.4D). This gap size was selected as the largest where the edges still appeared fused and were perceived as one edge.



**Figure 5.4.** Example stimuli. A and B show the achromatic edge in the centre and the C and D show the L-M edge in the centre. The gap between the two edges is exaggerated for illustrative purposes.

### *Procedure*

As in Experiment 5.1, participants were presented with the edge and marker and moved the edge, using the mouse, until they were satisfied that the two were aligned. There was no limit to the presentation time and there was a 300ms ISI between trials. Presentation order of the conditions was randomised and 40 trials were collected per condition.

Luminance information was presented at a Michelson contrast of 0.02 and chromatic information at contrasts of 0.1, 0.2 and 0.3.

### *Data Analysis*

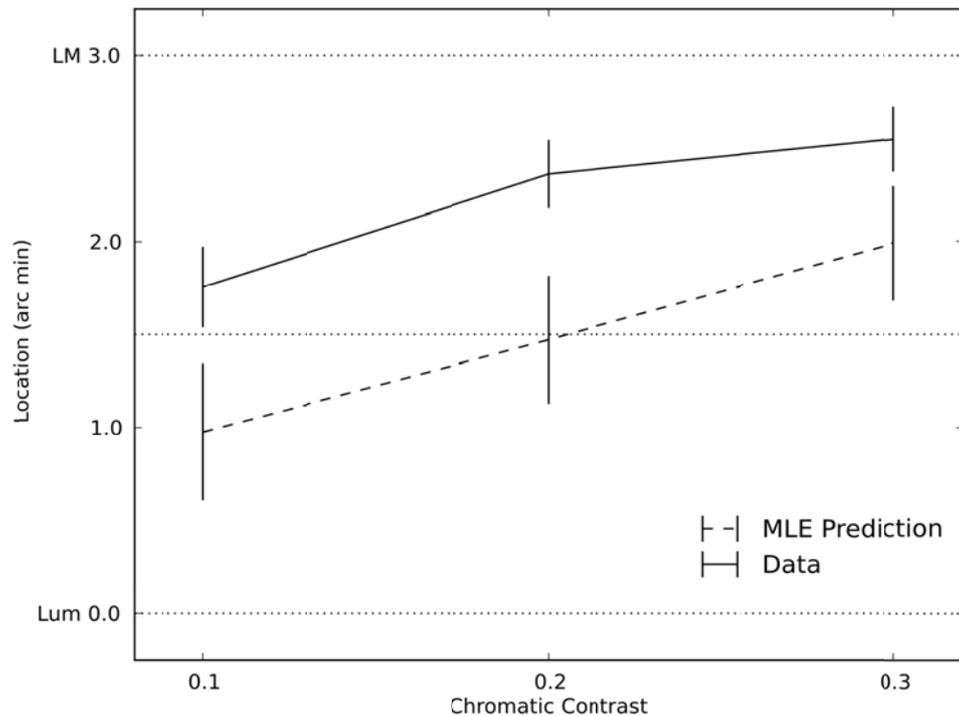
As in Experiment 5.2, the standard deviation was measured for the component edges alone (luminance at a contrast of 0.02 and L-M at contrasts of 0.1, 0.2 and 0.3) and for the three combined edges (luminance contrast of 0.02 combined with each of the three L-M contrasts).

#### 5.4.2. Results

MLE was used to generate predictions of edge localisation judgements based on participant performance with achromatic and isoluminant stimuli.

In Experiment 5.2, a luminance component with contrast of 0.02 combined with chromatic components of contrast 0.1, 0.2 and 0.3 showed variation between the L-M contrasts but did not equate the cues. However, in this experiment this combination of contrasts did equate the cues and as such judgements were predicted to be centred around the midpoint between the two edges. Experiment 5.2 only had a sample size of one, versus five in the current experiment, and so it is likely that the findings of the current experiment are more representative.

Surprisingly, participants' edge judgements were far closer to the chromatic edge than was predicted (Figure 5.5). This suggests that participants are weighting chromatic information more heavily than MLE would predict.



**Figure 5.5.** MLE predictions are shown with the dashed line, the behavioural data are shown with the solid line. The luminance edge is at position 0.0 and the chromatic edge is at 3.0 on the abscissa. The predictions show that the cues have been equated. However, the behavioural data shows that participants are judging the edge to be closer to the chromatic edge than predicted. If unweighted averaging predicted behaviour participant judgements would all be equidistant between the two edges, along the central dotted line. If 'winner takes all' predicted behaviour participant judgements could only be at either edge, the top and bottom dotted lines. Error bars represent  $\pm 1$  standard error of the mean.

### 5.5. Discussion

Despite using a range of contrasts selected to give a range of different predictions, the results from Experiment 4.2 were replicated and once again it was not possible to differentiate between the three models. It was also not possible to resolve any benefit resulting from the presence of both colour and luminance cues.

When considering conflicting edges, chromatic cues are weighted surprisingly strongly in edge localisation under the conditions tested, given their relative reliability in isolation. MLE accurately predicted the *pattern* of results across contrasts. However, under these particular conditions it consistently over-estimated the relative importance of the luminance cue. This that could be accounted for by a simple scale factor.

The results show that the weights generated from measurements of each component in isolation are not sufficient to predict edge localisation in conflicting conditions. The chromatic component requires a higher weight than would be predicted by MLE. At contrasts for which the component cues are equally reliable for localising the edge, the visual system gives more weight to the chromatic information. The weights generated using MLE represent the optimal combination of the signals based on their Bayesian likelihood functions. The fact that the predictions did not match the behaviour of participants suggests either that the system is failing to combine the signals in an optimal manner, or that it is optimal but is using additional information. In this case it may be that there is a Bayesian *prior* increasing the weight of the chromatic component, perhaps reflecting the utility of chromatic edge information in natural scenes.

Luminance has higher effective contrast in natural scenes (Rivest & Cavanagh, 1996) and is more reliable in most natural viewing conditions. Yet, it appears to be chromatic information that has the greater influence over perceived edge location when the cues are equated for reliability. This may be

because chromatic information is a more reliable cue for detecting object borders (McGraw, Whitaker, Badcock, & Skillen, 2003; Ruderman, Cronin, & Chiao, 1998). In natural scenes, achromatic information can represent variations in lighting, rather than object borders. For example, in areas of dappled shade the majority of achromatic discontinuities will represent shadows, making luminance information less useful than chromatic information for detecting object boundaries.

Hansen and Gegenfurtner (2013) compared human-labelled edges with computationally detected edges in natural scenes. The presence of chromatic (both L-M and S-cone) information in the stimuli improved performance by about 3% on average compared to luminance information alone, but reached up to 11% for some images. This type of experiment could represent a starting point not only to explain why chromatic information is weighted more heavily than is predicted by MLE, but also predict how much the weights need to be scaled.

In order to determine whether the Bayesian prior suggested above is based on image statistics, we would need to determine the relative proportions of discontinuities in each domain that represent object boundaries. If achromatic and isoluminant edges are detected using a computer algorithm, observers could then label these edges as representing edges of objects or other types of edges e.g. shadows. The percentage of edges that represent object boundaries could then be calculated for each cue. These values could then be used to scale the chromatic and luminance

weights. For example, if 100% of isoluminant edges represented object borders but only 70% of achromatic edges represented object borders, this would mean that the luminance weight should be scaled to 70% of its original value. If this procedure were carried out for a large natural image database it would allow us to calculate the mean factor that each cue should be scaled by. These mean values could then be compared to participant performance in the current task.

In conclusion, chromatic information is used more than would be predicted by weighted averaging in an edge localisation task. This may be because luminance edges in natural scenes often represent changes in illumination rather than object edges, meaning that chromatic information can have higher utility in identifying object edges.

## 6. Detection of perturbation in chromatic and luminance stimuli is modulated by context.

There is conflicting evidence as to which of luminance or chromatic edges is the more salient. One possibility is that each can be dominant under different circumstances. Findings presented earlier in this thesis and in the existing literature suggest that context i.e. the composition of the stimulus, may determine which is utilised more by the visual system.

The 'colour-shading effect' (Kingdom, 2003) occurs when a chromatic grating is added to a differently oriented luminance grating such that there is an impression of depth. This perceived depth can be suppressed by the addition of a second chromatic grating of the same orientation and spatial phase as the luminance grating. This supports the suggestion that the visual system has 'built-in' assumptions (Kingdom, 2008), in this case that aligned colour and luminance changes represent the edges of objects (hence are perceived as flat), whereas achromatic edges are perceived as changes in illumination, such as cast shadows (hence the perception of corrugation). It seems that the mere presence or absence of cues is insufficient to study interactions between colour and luminance, the spatial relationships of those cues must also be taken into account.

Chapter 3 demonstrated that chromatic blur can be masked by sharp luminance information. Similarly the Boynton illusion (Kaiser, 1996) shows how chromatic boundaries can appear shifted to align with luminance edges. These effects may be due to the luminance information constraining the

chromatic information and that can only occur when the edges are *aligned*. On the other hand, Kingdom and colleagues (2010) found that chromatic variations can suppress luminance variations, suggesting that chromatic information is dominant. This is also supported by the findings of Chapter 5, where chromatic information had more influence than predicted. However, Kingdom et al (2010) only demonstrated this effect when the chromatic and luminance edges were *orthogonal*. In Chapter 5 the edges were parallel but no longer spatially aligned. Therefore, it may be that when edges are aligned luminance is dominant, but when they are spatially separated chromatic information is dominant.

Another possible explanation for the differences in cue dominance is differences in the type of task used. Luminance dominance was found in a blur discrimination task whereas chromaticity had more influence in an edge localisation task and a saliency task. It may be that luminance dominance is restricted to blur discrimination tasks.

In order to determine whether cue dominance is determined by the spatial arrangement of the cues or the type of task, we will investigate whether there are differences between aligned and orthogonal conditions. If there are differences between these conditions that will indicate that dominance is determined by the spatial arrangement. If there is no difference between the two conditions that will suggest that dominance is task dependent.

In addition to the spatial relationship between colour and luminance information, the nature of that information may affect how it is combined; different types of grating combinations may have different effects. For example, there is increased facilitation for chromatic contrast sensitivity by a luminance pedestal for square-wave gratings versus sine-wave gratings (Gowdy, et al., 1999). This may be due to sharp edges promoting segmentation, suggesting that square-wave gratings may be combined differently to sine-wave gratings. In the same vein, luminance lines, which have sharp edges similar to those created by square-wave gratings, seem to have a facilitatory relationship with low spatial-frequency chromatic information. For example, reticles (thin, low contrast, achromatic lines) improve contrast sensitivity for chromatic gratings when they are aligned with the zero crossings (Montag, 1997). Furthermore, a thin luminance ring surrounding a uniform chromatic test facilitates contrast detection as much as a uniform luminance pedestal (Cole, et al., 1990). These effects may be based on processes similar to those that underlie the gap effect, where a gap or contour between two chromatic fields improves chromatic discrimination (Boynton, Hayhoe, & Macleod, 1977). These findings suggest that luminance-defined lines improve contrast detection of chromatic stimuli by improving segmentation. Therefore, it may be that the introduction of luminance lines could also improve performance in edge detection tasks.

A novel approach will be employed; spatial perturbation detection. Gratings will be sinusoidally perturbed in space and subjects will be asked to detect which of two stimuli is not straight. The paradigm aims to measure the

point at which the perturbation is just noticeable; the detection threshold for the perturbation. This technique can investigate whether one type of information is masked by another and whether there is interaction between the two cues. For example, if chromatic perturbation is harder to detect in the presence of luminance information i.e. it is masked by it, then this is evidence for luminance being prioritised above colour.

To test whether the spatial arrangement of gratings changes the way that chromatic and luminance information is combined, perturbation thresholds will be measured for chromatic and luminance gratings alone and in aligned and orthogonal combinations (Experiment 6.1). If cue dominance is determined by their spatial arrangement we will expect luminance information to mask detection of chromatic perturbation in the aligned conditions and chromatic information to mask detection of luminance perturbation in the orthogonal conditions. On the other hand, if luminance dominance is restricted to blur discrimination tasks we will not expect to see luminance dominance in any condition.

In addition to the question of whether the spatial relationship is critical in determining which cue is dominant, it might also be that the type of grating affects the way that chromatic and luminance information is combined. Therefore perturbation thresholds will also be measured for line and square-wave gratings alone and in combination (Experiment 6.2). If luminance-defined lines improve contrast detection of chromatic stimuli by improving segmentation we should expect the presence of a straight

chromatic square-wave to facilitate perturbation detection in luminance defined lines. However, if this is not the case we would expect the introduction of a straight grating to cause masking in all cases.

## 6.1. Aligned and orthogonal stimuli.

### 6.1.1. Method

#### *Participants*

Two male and two female volunteers (including the author), aged between 19 and 35, participated in the study. Two of the participants (one male) were naive to the aims of the study.

#### *Apparatus*

A chin rest was used to ensure that participants viewed the stimuli from a constant 367cm distance, giving a viewable area that subtended  $5.62^\circ$  of visual angle.

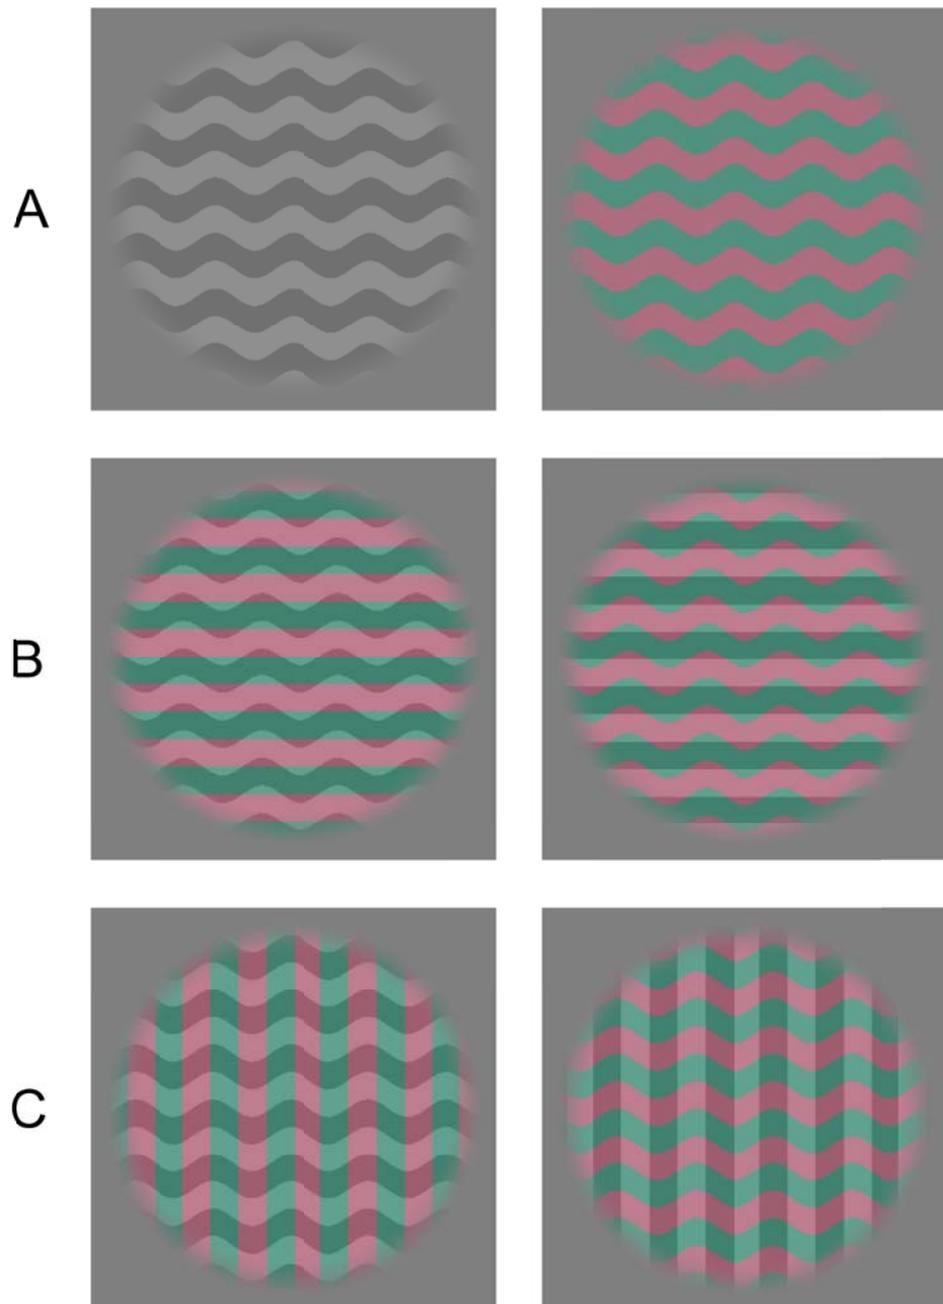
#### *Stimulus Generation*

Isoluminance was determined psychophysically for each observer. Three grating types were created; square-wave, sine-wave and line. All stimuli were presented with a raised cosine mask with a fringe width of  $0.4^\circ$  (Figure 6.1). Stimuli were presented at a size of  $2^\circ$ , spatial frequency of 0.5 cpd. Perturbed (target) gratings were presented at an orientation of  $90^\circ$  (horizontal) and straight (masking) gratings were presented at  $0^\circ$  (vertical) for the orthogonal condition and  $90^\circ$  for the aligned condition.

In order to equate the contrasts of the component gratings, contrast detection thresholds for the straight gratings alone were measured using a 2IFC, staircase procedure. Stimuli were presented at five times these detection thresholds. Two participants (BLW and RJS) had contrast detection thresholds above 0.2 for the isoluminant line stimuli and so it was not possible to present them at five times threshold (because this would be beyond the gamut of the monitor). In these cases the isoluminant component was presented at maximum contrast and the luminance component was scaled, such that the ratio was the same as if the chromatic component had been presented at five times threshold.

For the remainder of this chapter gratings that have been spatially perturbed will be referred to as the target and gratings with no perturbation will be referred to as the mask.

In all cases subjects were required to determine which stimulus contained the perturbation, and they were always aware whether the perturbation would be in the chromatic or luminance component (see Procedure for details). For each of the grating types (square-wave, line and sine-wave gratings) the signal channel was presented either with no mask (Figure 6.1A), with an aligned mask (Figure 6.1B) or with an orthogonal mask (Figure 6.1B), where this mask was always in the other channel. For example, the luminance perturbation threshold was measured with *no mask*, with an *aligned chromatic mask* and with an *orthogonal chromatic mask*.



**Figure 6.1.** Example square-wave stimuli, the same combinations were also created for line and sine-wave gratings. Luminance perturbation is shown in the left column and chromatic perturbation in the right column. (A) Shows the target gratings alone. (B) Shows the target gratings in combination with aligned masks. (C) Shows the target gratings in combination with orthogonal masks. Figure for illustrative purposes only.

*Procedure*

A 2IFC design was employed. Participants were presented with two stimuli (target and foil) for 500ms separated by a 500ms ISI and asked which appeared to have been spatially perturbed. The presentation order of the target and foil was randomised.

In each condition we measured the minimal amount of spatial perturbation that could be detected; the perturbation threshold as the amplitude of the sinusoidal modulation in arc minutes of visual angle. We measured the perturbation threshold for each grating alone as a baseline. We then measured the threshold for the colour target combined with luminance mask and luminance target combined with colour mask in both aligned and orthogonal combinations, in order to determine whether there was any masking present and, if so, whether it was dependent on the relationship between the component gratings.

The perturbation threshold for each condition was determined using a one-up, three-down staircase procedure. The staircases controlled the amount of spatial perturbation in the target stimulus. A different staircase was implemented for each condition and these staircases ran sequentially; staircases were not interleaved but the order of the conditions was randomised. As a result, participants were aware on each trial of which component would contain the perturbation and could attend accordingly. The staircases were designed to converge on the 79% perturbation discrimination threshold for each condition and aborted after 50 trials. Each participant

collected five staircases for each of the six conditions (Figure 6.1), for each of the three grating types (90 staircases in total); 360 staircases in total across participants.

### *Data Analysis*

Due to variability across the staircases it was not possible perform a fitting procedure on the data (see Section 2.11 for details). Instead, the mean of the last six reversals for each staircase was taken as a measure of perturbation threshold. The mean of these thresholds was then taken to give one value per participant, per condition i.e. 24 values overall.

#### 6.1.2. Results

Two-way ANOVAs were performed for each participant for each grating type. Pairwise comparisons were corrected for the three comparisons within each ANOVA and across the four participants (comparisons = 12). Bonferroni correction was not performed across grating type as the data were not directly compared in this way.

There was more variation between participants than might generally be expected in a psychophysical study. As a result, data analyses are presented for each individual. This allows us to investigate whether there is a pattern of differences between the participants. For example, if half the participants showed one pattern and the other half another, this could suggest that there is more than one strategy that could be used to undertake the task. The details of the variations between participants and the potential

reasons for them are explored below and data is presented graphically in Figure 6.2.

### *Square-wave Grating*

There was a main effect of chromaticity for all participants, apart from AJB; thresholds were significantly higher for L-M than luminance defined targets (BLW:  $F_{(1,24)} = 5.123$ ,  $p=0.033$ ,  $MS_{\text{colour}} = 107.751$ , JAF:  $F_{(1,24)} = 60.316$ ,  $p<0.001$ ,  $MS_{\text{colour}} = 258.596$ , RJS:  $F_{(1,24)} = 75.185$ ,  $p<0.001$ ,  $MS_{\text{colour}} = 92.278$ ).

A main effect of condition was found for two participants (JAF:  $F_{(2,24)} = 8.119$ ,  $p=0.002$ ,  $MS_{\text{condition}} = 69.621$ , RJS:  $F_{(2,24)} = 17.884$ ,  $p<0.001$ ,  $MS_{\text{condition}} = 44.376$ ). Pair-wise comparisons showed a significant difference between alone and aligned conditions for both participants (JAF:  $p=0.008$ , RJS:  $p=0.004$ ) and alone and orthogonal for RJS ( $p=0.004$ ); where alone thresholds were lower than those for combined conditions.

An interaction effect was found for participant JAF, reflecting an elevated threshold for the L-M aligned condition ( $F_{(2,24)} = 18.351$ ,  $p = 0.026$ ,  $MS_{\text{interaction}} = 36.702$ ).

### *Line Grating*

Two participants showed a main effect of chromaticity; thresholds were significantly higher for L-M than luminance defined targets (JAF:  $F_{(1,24)} = 61.285$ ,  $p<0.001$ ,  $MS_{\text{colour}} = 276.751$ , RJS:  $F_{(1,24)} = 44.542$ ,  $p<0.001$ ,  $MS_{\text{colour}} = 232.19$ ).

Participant JAF showed a main effect of condition ( $F_{(2,24)} = 3.785$ ,  $p=0.037$ ,  $MS_{\text{condition}} = 34.181$ ) but pair-wise comparisons revealed no specific differences. No other participants showed an effect of condition.

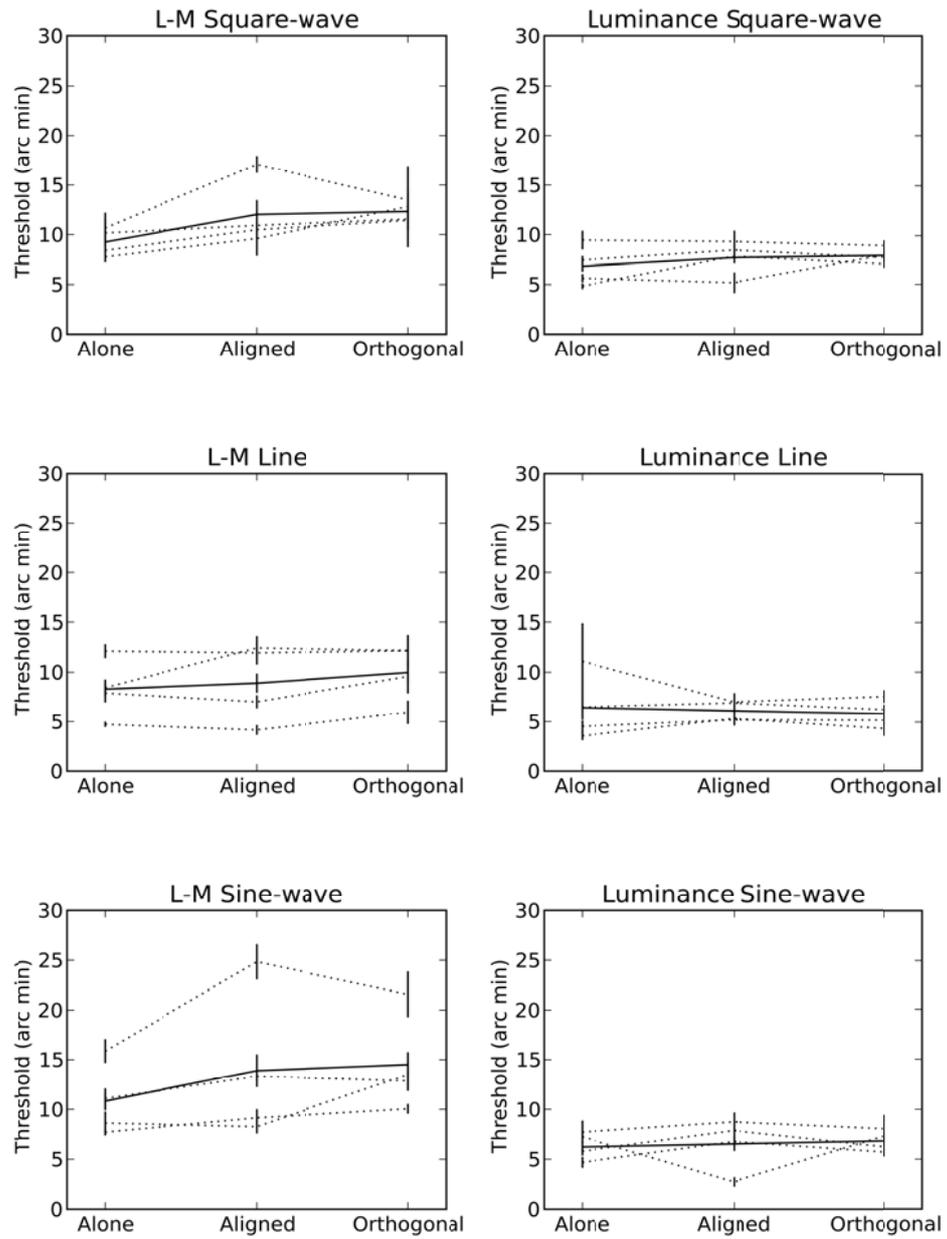
No participants showed an interaction effect.

### *Sine-wave Grating*

All participants showed a main effect of chromaticity; thresholds were significantly higher for L-M than luminance defined targets (AJB:  $F_{(1,24)} = 25.758$ ,  $p<0.001$ ,  $MS_{\text{colour}} = 147.973$ , BLW:  $F_{(1,24)} = 36.219$ ,  $p<0.001$ ,  $MS_{\text{colour}} = 145.125$ , JAF:  $F_{(1,24)} = 169.447$ ,  $p<0.001$ ,  $MS_{\text{colour}} = 1682.487$ , RJS:  $F_{(1,24)} = 15.047$ ,  $p<0.001$ ,  $MS_{\text{colour}} = 40.644$ ).

One participant (BLW) showed a main effect of condition ( $F_{(2,24)} = 5.758$ ,  $p = 0.009$ ,  $MS_{\text{condition}} = 46.144$ ) and pair-wise comparisons showed a significantly lower threshold for the alone versus orthogonal condition ( $p = 0.036$ ).

Interaction effects were present for two participants (BLW:  $F_{(2,24)} = 6.335$ ,  $p=0.006$ ,  $MS_{\text{interaction}} = 50.765$ , JAF:  $F_{(2,24)} = 11.566$ ,  $p<0.001$ ,  $MS_{\text{interaction}} = 229.678$ ). JAF once again showed an elevated threshold for the L-M aligned condition and BLW had an elevated threshold for the L-M orthogonal condition.



**Figure 6.2.** Results for Experiment 6.1. The left column shows results for L-M defined targets and the right column shows results for luminance targets. Each row represents a different grating type. Solid lines represent the data pooled across participants, dotted lines show the data for each participant. There are no significant differences between the three conditions for any of the six stimulus types. Error bars represent  $\pm 1$  standard error of the mean.

### *Summary*

There is a clear pattern of L-M defined targets having higher thresholds than luminance defined targets. This is surprising, given that the gratings were equated for contrast detection thresholds. Chromatic aberration and the low-pass nature of chromatic processing can limit chromatic acuity and could explain this difference. However, this may not be sufficient explanation, as Krauskopf and Forte (2002) showed that Vernier thresholds are not affected by chromaticity when they are equated for contrast, and this is consistent across a large range of multiples of threshold. Instead, there may be a cortical mechanism that restricts chromatic processing of perturbation. It should be noted that this effect is much more consistent for the sine- and square-wave gratings than for line gratings.

The main effect of condition and interaction effects are less consistent than the chromaticity effect. For a majority of participants there was no significant difference between the three conditions (alone, aligned and orthogonal) and no interaction. Variability between participants may have been due to the use of cognitive strategies. Debriefing revealed that as the staircases were not interleaved, two of the participants (BLW and AJB) were selecting small areas of the stimuli that gave the maximum amount of information about perturbation, effectively ignoring the straight grating. As a result the above data may not give a true reflection of perturbation thresholds in the presence of straight gratings.

It appears that the straight gratings are being disregarded, but this may be an artefact created by the use of cognitive strategies. Previous work using flankers in edge localisation judgements suggests that, in these types of judgements, irrelevant stimuli are processed involuntarily (Rivest & Cavanagh, 1996) but in the case of this experiment the conditions might have allowed that not to be the case. In order to determine whether the use of cognitive strategies is causing the lack of a masking effect it would be necessary to repeat the experiment but interleave the staircases and randomise the phase (of the perturbation) in the target gratings. However, there were no consistent significant differences between conditions for the two participants that did not report using cognitive strategies. This suggests that, whilst interleaving the staircases may reduce the variability across participants, it is unlikely to introduce significant differences between the alone, aligned and orthogonal conditions. However, variability of participant responses may underlie the main effect of chromaticity. There is more variability in the thresholds for L-M targets and if this is reduced the difference between L-M and luminance targets may also be reduced, eliminating the effect. This will be tested for the line and square-wave stimuli in Experiment 6.2. If the effect is not replicated there this will demonstrate that is likely an artefact created by the blocking of conditions.

## 6.2. Square-wave and line stimuli.

It may be that the type of gratings (e.g. square-wave, sine-wave, etc.) presented affect how colour and luminance are combined. In particular, as

discussed in the introduction the presence of a straight chromatic square-wave may not mask perturbation in luminance lines but instead produce facilitation. However, this facilitation would not be expected in any other combination of lines and square-waves. In this experiment we will be testing different combinations of line and square-wave gratings to determine whether the type of combination affects perturbation thresholds.

Experiment 6.1 found no difference between aligned and orthogonal masks when the target and mask were of the same grating type, meaning that testing both of these conditions is redundant. Using aligned combinations will allow us to test whether chromatic perturbation can be constrained by straight luminance gratings, in a similar manner to chromatic blur in Chapter 3 or chromatic information in the Boynton illusion (Kaiser, 1996). For these reasons only aligned combinations will be tested in this study.

### 6.2.1. Method

#### *Participants*

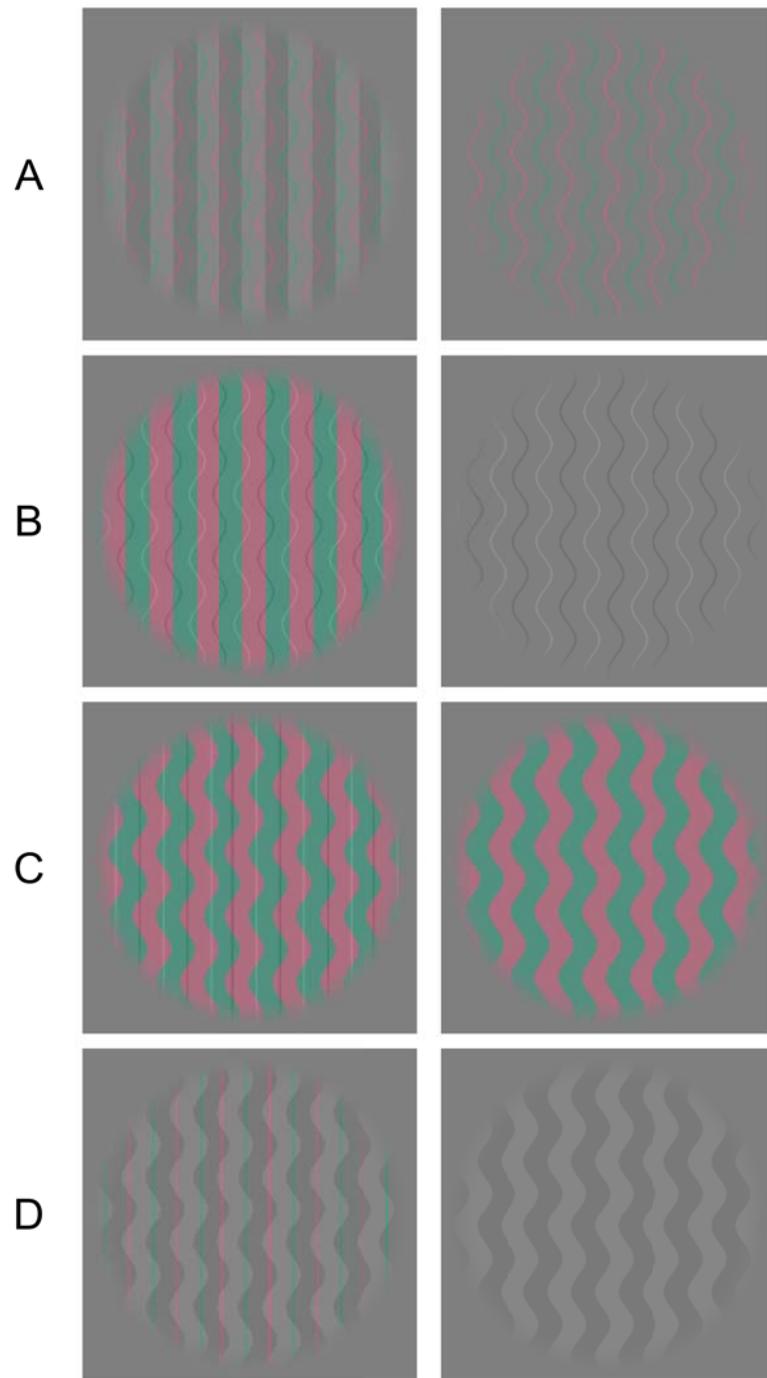
Two male and three female volunteers (including the author), aged between 22 and 35, participated in the study. Three of the participants (two male) were naive to the purposes of the study.

#### *Apparatus*

The same apparatus was used as for Experiment 6.1.

*Stimulus Generation*

The gratings were generated in the same manner as for Experiment 6.1. Gratings were presented vertically to prevent any perceptual learning carrying over from Experiment 6.1, for those participants who took part in both experiments (Fahle & Edelman, 1993). Eight conditions were created, four component gratings and four combinations: chromatic-line target alone and when combined with a luminance-square-wave mask (Figure 6.3A), luminance-line target alone and when combined with a chromatic-square-wave mask (Figure 6.3B), chromatic-square-wave target alone and when combined with a luminance-line mask (Figure 6.3C) and luminance-square-wave target alone and when combined with a chromatic-line mask (Figure 6.3D).



**Figure 6.3.** Example stimuli. The left column shows the combined conditions and the right shows the perturbed gratings alone. Figure for illustrative purposes only.

### *Procedure*

The same procedure was used as in Experiment 6.1 with the following adjustments. Staircases were interleaved, and the absolute phase of the gratings and the phase of the spatial perturbation were randomised for each trial. These adjustments were made in an effort to avoid the possibility of participants using cognitive strategies to inform their judgements. The randomisations coupled with interleaving the staircases meant that participants could not predict which cue would be perturbed or which area of the stimulus would contain the most information about the perturbation. In the combined conditions the absolute phase of the pair of gratings was kept the same, to preserve the alignment at the zero crossing point; ensuring that lines are spatially aligned with the edges in the square-wave.

### *Data Analysis*

Due to variability across the staircases it was not possible perform a fitting procedure on the data (see Section 2.11 for details). Instead, the mean of the last six reversals for each staircase was taken as a measure of perturbation threshold. The mean of these thresholds was then take to give one value per participant, per conditions i.e. 40 values overall.

Difference values were then calculated for four pairings as illustrated in Figure 6.3, each was designed to see the effect of adding a straight grating of the other type to the perturbed grating. For ease of exposition the four conditions will be referred to by the following letters the remainder of this chapter:

- A) The difference in thresholds between a chromatic-line target alone and a chromatic-line target combined with a luminance-square-wave mask.
- B) The difference in thresholds between a luminance-line grating target and a luminance-line target combined with a chromatic-square-wave mask.
- C) The difference in thresholds between a chromatic-square-wave target alone and a chromatic-square-wave target combined with a luminance-line mask.
- D) The difference in thresholds between a luminance-square-wave target alone and a luminance-square-wave target combined with a chromatic-line mask.

### 6.2.2. Results

Three staircases were excluded from all analysis as they did not converge; two from the combined chromatic-line target and luminance-square-wave mask condition and one from the combined luminance-line target and chromatic-square-wave mask condition.

#### *Planned comparisons to test for an effect of chromaticity*

Four planned comparisons were performed on the raw threshold values to determine if the main effect of colour found in Experiment 6.1 was replicated. The data were collapsed in four different ways: across grating type when the target was presented alone; across all data and for each grating

type (square-wave and line) when the target was presented alone. *P*-values were Bonferroni corrected across t-tests and participants (comparisons = 20).

For the first planned comparison, the data were collapsed across grating type when the target was presented *alone*. Independent samples t-tests were then performed for each participant to determine whether there was a significant difference between the thresholds for L-M defined targets and luminance defined targets; the thresholds for L-M targets *alone* were compared to the thresholds for luminance targets *alone*. A significant difference was found for only one participant (RJS:  $t_{18} = 3.728$ ,  $p = 0.04$ ).

For the second planned comparison, the data were collapsed across *all* conditions and grating types. Independent samples t-tests were then performed for each participant to determine whether there was a significant difference between the thresholds for L-M defined targets and luminance defined targets. The thresholds for *all* L-M targets, both alone and combined, were compared to *all* the thresholds for luminance targets, both alone and combined. Significant differences were found for two participants (RJS:  $t_{37} = 3.701$ ,  $p = 0.02$ ; BLW:  $t_{37} = 4.326$ ,  $p = 0.02$ ).

For the third and fourth planned comparisons, the data were collapsed across conditions for *each grating type* when the target was presented alone. Independent samples t-tests were performed for each participant to determine whether there was a significant difference between the thresholds for L-M defined targets and luminance defined targets. For each of the two grating types (square-wave and line) thresholds for L-M

targets *alone* were compared to thresholds for luminance targets *alone*. A significant difference was found for only one participant for the square-wave gratings (RJS:  $t_8 = 2.831$ ,  $p = 0.02$ ) but no significant differences were found for the line gratings.

In summary, for a majority of participants there was no significant difference between perturbation thresholds for L-M and luminance defined targets in any of the four comparisons performed. Therefore, it was concluded that the significant main effect of chromaticity found in Experiment 6.1 was not replicated here. It appears that the effect found in Experiment 6.1 was an artefact caused by the conditions being presented sequentially rather than in an interleaved manner.

#### *Individual participant results*

Analyses were conducted for each participant. A two-way ANOVA was performed for each participant to determine whether there were significant differences between the four difference values. Main effects were Bonferroni corrected across participants (comparisons = 5) and pairwise comparisons were Bonferroni corrected within each ANOVA and across participants (comparisons = 30).

As in Experiment 6.1 there was variability between participants (Figure 6.4). Notably participant DJH did not show a main effect of condition and demonstrated generally attenuated differences. This suggests that DJH may have a greater capacity for disregarding irrelevant information than the other participants.

The other four participants all showed a main effect of condition (ATA:  $F_{(3,16)} = 16.064$ ,  $p=0.005$ ,  $MS_{\text{condition}} = 1133.457$ , BLW:  $F_{(3,15)} = 18.516$ ,  $p=0.005$ ,  $MS_{\text{condition}} = 196.952$ , DS:  $F_{(3,15)} = 20.879$ ,  $p=0.005$ ,  $MS_{\text{condition}} = 777.233$ , RJS:  $F_{(3,15)} = 27.539$ ,  $p=0.005$ ,  $MS_{\text{condition}} = 1731.04$ ).

Pair-wise comparisons showed significant differences between conditions B and C for four participants (ATA:  $p=0.005$ , BLW:  $p=0.005$ , DS:  $p=0.005$ , RJS:  $p=0.01$ ); luminance lines have a greater masking effect than a chromatic square-wave.

Pair-wise comparisons also showed significant differences between conditions B and D for four participants (ATA:  $p=0.005$ , BLW:  $p=0.005$ , DS:  $p=0.005$ , RJS:  $p=0.005$ ); chromatic lines have a greater masking effect than a chromatic square-wave.

Significant differences were also found for three participants between conditions A and D (ATA:  $p=0.005$ , DS:  $p = 0.015$ , RJS:  $p=0.005$ ); chromatic lines have a greater masking effect than an achromatic square-wave.

There was no significant difference for any of the participants between conditions A and C. Each of the remaining pair-wise comparisons only revealed significant differences for one participant; a different participant in each case: A and B for participant BLW ( $p = 0.025$ ) and C and D for participant RJS ( $p = 0.045$ ).

*Group results*

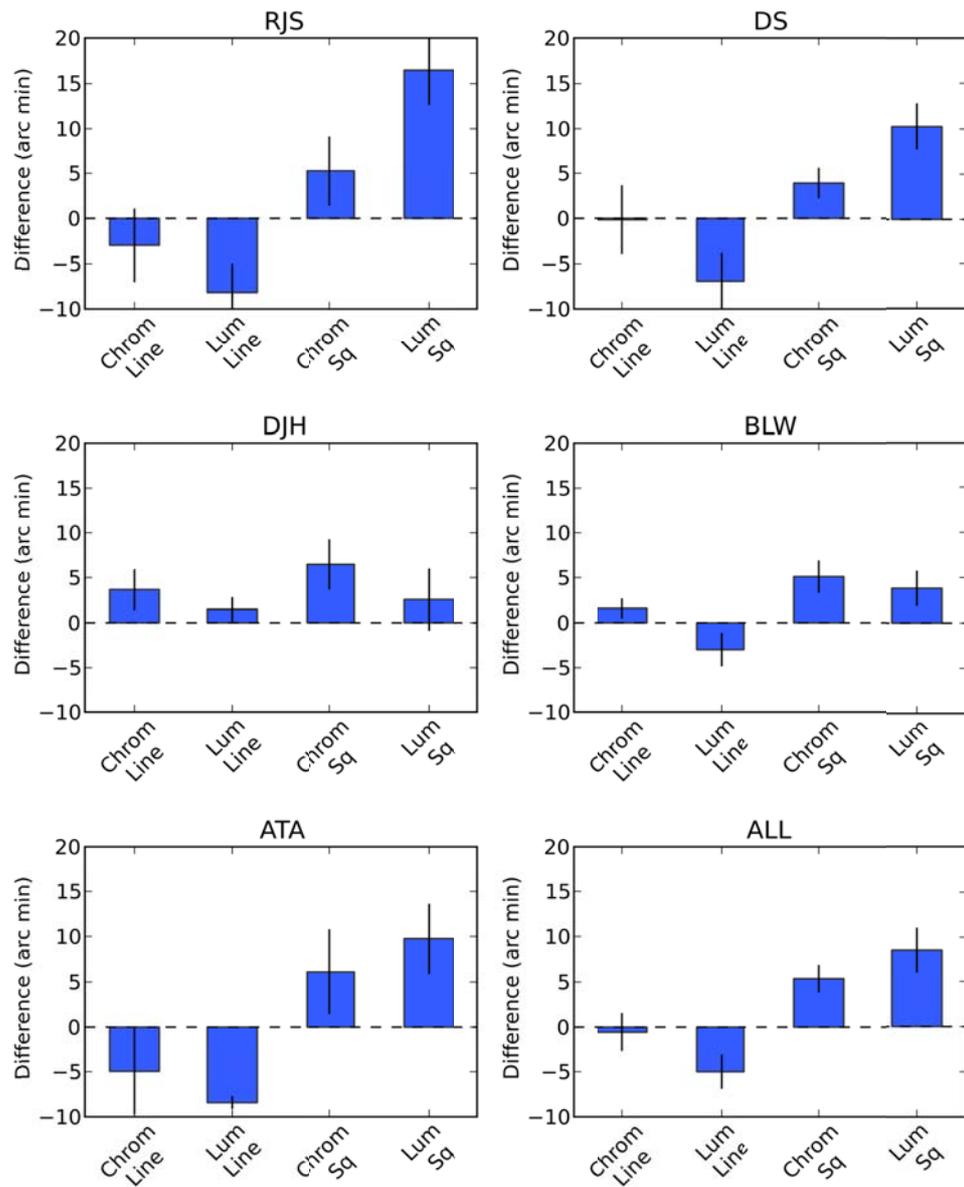
The data were pooled across participants and one-sample t-tests were performed to determine whether the difference value for each condition was significantly different from zero; a difference value of zero would mean that the mask had no effect on perturbation thresholds. Significance values were Bonferroni corrected (comparisons=4).

Condition A was not significantly different from zero. Achromatic square-wave gratings do not mask, or facilitate, perturbation detection for chromatic line targets.

Condition B was significantly different from zero ( $t_{(23)} = -5.584$ ,  $p = 0.004$ ). Chromatic square-wave gratings significantly *facilitate* perturbation detection for luminance line targets.

Condition C was significantly different from zero ( $t_{(24)} = 7.216$ ,  $p = 0.004$ ). Luminance lines significantly mask perturbation detection for chromatic square-wave targets.

Condition D was significantly different from zero ( $t_{(24)} = 6.846$ ,  $p = 0.004$ ). Chromatic lines significantly mask perturbation detection for luminance square-wave targets.



**Figure 6.4.** Results for each participant and collapsed across participants. The conditions on the abscissa denote perturbed grating type. The difference on the ordinate is the threshold for the combined condition minus the threshold for the alone condition such that a positive value represents poorer performance in the combined condition and a negative value represents improved performance in the combined condition. Four participants' performance improved in condition B, when a straight chromatic square-wave was introduced to perturbed luminance lines. Error bars represent the 95% confidence intervals.

### 6.3. Discussion

In summary, Experiment 6.1 did not show differences in perturbation threshold between the alone, aligned and orthogonal conditions. Thresholds for L-M targets were found to be significantly higher than those for luminance targets in a majority of cases. However, this was not replicated in Experiment 6.2. It is likely that the effect of chromaticity in Experiment 6.1 was an artefact created by variability between participants, particularly for L-M targets. Interleaving the staircases in Experiment 6.2 reduced variability overall, potentially explaining the lack of replication and supporting the idea that the difference found in Experiment 6.1 was an artefact.

The key findings of Experiment 6.2 are that the introduction of a L-M defined square-wave mask improves perception of perturbation in luminance lines, but the introduction of a luminance defined square-wave mask has little effect on perturbation thresholds for chromatic lines. This demonstrates that it is not merely the combination of cross-channel square-wave and line gratings that produces facilitation of perturbation thresholds, but the specific combination of luminance-line target and chromatic-square-wave mask. In the other cases tested the introduction of a line mask increased thresholds and produced a masking effect.

In Experiment 6.1 there was considerable variation between participants. This is similar to findings by Clery et al (2013). They measured the perceived depth of combinations of sine- and square-wave gratings in three conditions. In the first the gratings were orthogonal and the contrast of the

luminance grating was varied, this produced data consistent with Kingdom's previous findings (Kingdom, 2003; Kingdom, et al., 2005a). In the other two conditions the chromatic contrast was varied and the gratings were presented in aligned or orthogonal arrangements. These conditions elicited substantial individual differences and most of the participants did not produce the expected pattern of data.

One of their possible explanations for this variation is individual differences in perceived chromatic contrast. However, in Experiment 6.1 we equated contrast detection thresholds and so this could not be the case. Another possible explanation is that the heuristics suggested by Kingdom (2008), such as achromatic edges always being interpreted as shadows, may be more idiosyncratic than previously thought. Clery et al (2013) suggest that the visual system may be more flexible than previously thought and individual biases could develop based on past experience. It is also possible, as discussed above, that participants used cognitive strategies to discount the irrelevant gratings by focussing on particular narrow regions of the stimulus.

The lack of effect of condition or interaction between condition and colour in Experiment 6.1 is particularly surprising in light of the contrast detection literature (see Section 1.4 for a more in-depth discussion). Suprathreshold chromatic pedestals produce a masking effect when combined with a luminance test (K. K. De Valois & Switkes, 1983), but suprathreshold luminance pedestals produce facilitation of a chromatic test (Cole, et al., 1990; K. K. De Valois & Switkes, 1983). This would suggest that

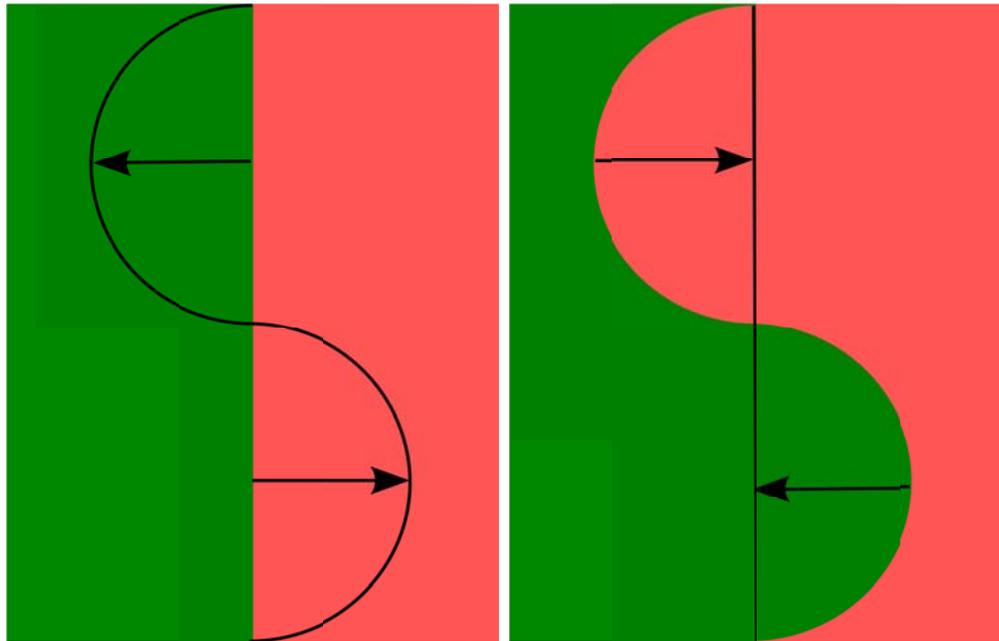
the presence of any mask should change the perturbation thresholds, whether by masking or facilitation.

It seems most likely that colour and luminance are not simply combined in the same way for perturbation detection as for contrast detection. A recurring theme throughout this thesis is that colour and luminance seem to not always be combined in the same way. In this case it may be that the visual system is processing the two components separately, and is not compelled to combine them, and the participant is, therefore, able to discount the irrelevant information at will. This would be advantageous for the visual system as it could allow the effortless separation of shadows and changes in surface material.

The results from Experiment 6.2 show that chromatic and luminance lines tend to mask perturbation in colour and luminance square-wave targets by a similar amount (conditions C and D). However, the introduction of a square-wave mask has no effect of perception of perturbation in chromatic lines (condition A) and produces facilitation for luminance lines (condition B). The greater masking effect of lines on perturbation thresholds suggests that they are a more salient stimulus for edge detection and so have a greater ability to disrupt perception of edges.

When a luminance line is presented with a chromatic edge, such as the chromatic boundaries in a square-wave grating, it appears that the chromatic information becomes 'tied' to the luminance information; the perceived location of the chromatic edge is determined by the location of the

luminance line (Figure 6.5). If the line and edge are perceived to be in the same location, this may increase the amount of information available to the visual system about the edge location. This increase in information could explain the facilitation effect found in Experiment 6.2.



**Figure 6.5.** Schematic illustration of how the perceived location of an edge in a chromatic square-wave could become tied to a luminance line. The chromatic contour is moved to align with the luminance line. In the left panel this would produce facilitation of perturbation detection and in the right panel this would produce a masking effect.

There are several pieces of work linking luminance lines to chromatic edges. The Boynton illusion demonstrates edge localisation being determined by a luminance line (Kaiser, 1996), a luminance outline facilitates contrast detection of a uniform chromatic field as much as a uniform pedestal (Gowdy, et al., 1999) and similarly luminance lines (reticles) facilitate chromatic contrast detection of gratings (Montag, 1997). Taken together, these findings

seem to demonstrate that there is a specific and important relationship between *chromatic edges* and *luminance lines*.

If chromatic edges and luminance lines are perceived as being in the same location this could have two effects. Firstly, as described above, the addition of chromatic information could improve localisation of the luminance contour. Secondly, separating two chromatic surfaces with a luminance contour improves colour discrimination (Boynton, et al., 1977). This improved colour discrimination could aid image segmentation by making similarly coloured surfaces easier to differentiate. This is in keeping with the idea that colour is primarily used to process surface properties and to facilitate segmentation and grouping (Mollon, 1989), whereas luminance is used for tasks requiring high spatial precision (Peirce, et al., 2008), such as edge localisation.

In conclusion, the introduction of aligned or orthogonal straight gratings of the same type as the perturbed grating does not affect perturbation thresholds; the irrelevant grating is disregarded. This may help to distinguish between shadows and changes in surface in natural scenes. Conversely, chromatic edges appear to become tied to luminance lines which may serve to improve chromatic discrimination and segmentation.

## 7. Luminance information constrains chromatic blur in bipartite edges

In Chapter 3 we saw that chromatic blur is constrained by sharp luminance information in natural scenes, suggesting that luminance dominates chromatic cues. However, in Chapter 5 chromatic information was shown to have more influence than expected in synthetic conflicting edges. There are two possible reasons for these contradictory results; they may be due to differences in the types of stimuli used or they may be due to differences in the task.

Chapter 3 uses natural scene stimuli, whereas Chapter 5 uses synthetic stimuli and it may be that this difference could explain the conflicting results. The complexity of natural scenes can elicit findings not predicted by simple synthetic stimuli (Felsen & Dan, 2005). However, the details of this complexity are poorly understood and difficult to control for (Rust & Movshon, 2005). For example, in Experiment 3.3 it was shown that reversing the colour and luminance channels eliminated the main effect of chromaticity; isoluminant blur discrimination was no longer poorer than achromatic blur discrimination. This suggests that there is a difference in the statistical regularities of the two types of information, however, as natural scenes contain such a wealth of information it is very difficult to determine the source of this difference. It may be that a statistical feature of natural scenes, that is absent in the bipartite edges in Chapter 5, is driving the luminance dominance found in Chapter 3.

Chapters 3 and 5 also used different tasks; blur discrimination and edge localisation respectively. These tasks may differ in the mechanisms they use; blur information and edge localisation information may be encoded differently. If this is the case then both the sensitivity of low-level mechanisms and the existence of high-level representations may differ across the domains e.g. there may be an explicit representation of chromatic edges but not chromatic blur. F.A.A. Kingdom (personal communication, 16 May 2013) investigated the contribution of chromatic information to blur *appearance*. He found that when two textures differ in both colour and luminance blur, the perception of blur is driven by the luminance component. These results were interpreted to mean that, whilst chromatic blur can be detected nearly as well as luminance blur (Wuerger, et al., 2000; Wuerger, et al., 2001), we have a very limited 'sense' of colour blur; we do not have an explicit representation of chromatic blur (F.A.A. Kingdom, personal communication, 16 May 2013).

To know whether it is plausible that blur encoding is processed differently for chromatic and luminance channels it is worth considering the current dominant model for blur detection. Scale-space analysis (Georgeson, et al., 2007) uses two stages of spatial filtering; an odd symmetric Gaussian first order derivative filter (similar to a simple cell) and a third order derivative filter (similar to a complex cell). The outputs of these are half-wave rectified before feeding forward, producing a response sensitive to one edge polarity and removing 'phantom edges' (Georgeson, et al., 2007). This process creates a scale-space response map, on which the position and scale of peaks

represent the location and blur of edges. This model has been found to accurately predict human perception for a variety of luminance profiles, but has not been tested on chromatic information. This model suggests that edge location and blur are jointly encoded for luminance information. If chromatic information is also encoded in this way we should expect that the differences in the results of Chapters 3 and 5 to be driven by differences in the stimuli, not differences in the task; edges and blur should be processed in the same way.

In order to investigate the cause of the luminance dominance found in Chapter 3, the blur discrimination task from that chapter will be replicated using bipartite edges similar to those used in Chapter 5. If the luminance dominance was caused by the visual structure of the natural scenes stimuli then luminance dominance will not be replicated with bipartite edge stimuli. However, if the differences between Chapters 3 and 5 are caused by different task demands, the luminance dominance will persist.

### 7.1. Method

#### *Participants*

One male and four female volunteers (including the author), aged between 22 and 35, participated in the study. Three participants (one male) were naive to the purposes of the study.

### *Apparatus*

A chin rest was used to ensure that participants viewed the stimuli from a constant 367cm distance, giving a viewable area that subtended 5.62° of visual angle.

### *Stimulus Generation*

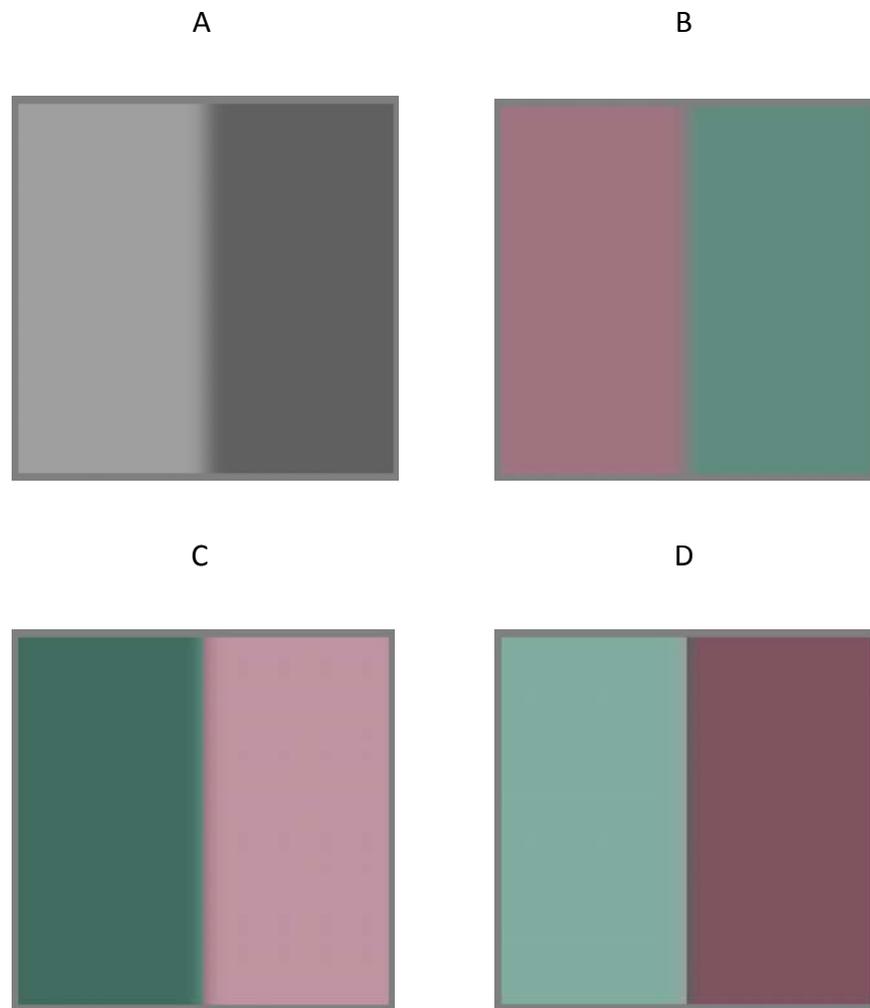
Two bipartite edges were created in MB-DKL space (Derrington, et al., 1984; Macleod & Boynton, 1979), luminance and L-M defined; the L-M defined edges were presented at psychophysical isoluminance. The stimuli were 2° x 2° in size and presented with a neutral grey background.

In Experiment 3.2 equating contrast was shown not to diminish the masking of chromatic blur by sharp luminance information, therefore contrast was not equated here. Each component edge was presented at a contrast of 0.25 meaning that the combined edges had a total contrast of 0.5.

In order to prevent bias, the polarity of the edges was randomised. In addition, randomisation of the chromatic edge polarity was independent of the randomisation of the luminance edge. Therefore the combined edges had four possible combinations; light-to-dark combined with red-to-green, light-to-dark combined with green-to-red, dark-to-light combined with red-to-green and dark-to-light combined with green to red.

There were four conditions; achromatic, isoluminant, sharp-achromatic combined with blurred-isoluminant and sharp-isoluminant combined with blurred-luminance (Figure 7.1). Blurring and stimulus

presentation was performed in the same manner as Chapter 3. Sharp edges were not blurred to any extent.



**Figure 7.1.** Example stimuli, all contain Gaussian blur ( $\sigma = 0.1^\circ$ ). The top row represents the two alone conditions; achromatic (A) and isoluminant (B). The bottom row represents the two combined conditions; luminance blur combined with sharp chromatic information (C) and chromatic blur combined with sharp luminance information (D).

### *Procedure*

A 2IFC design was employed, participants were presented with the two images (foil and target) for 300ms separated by a 500ms ISI and asked

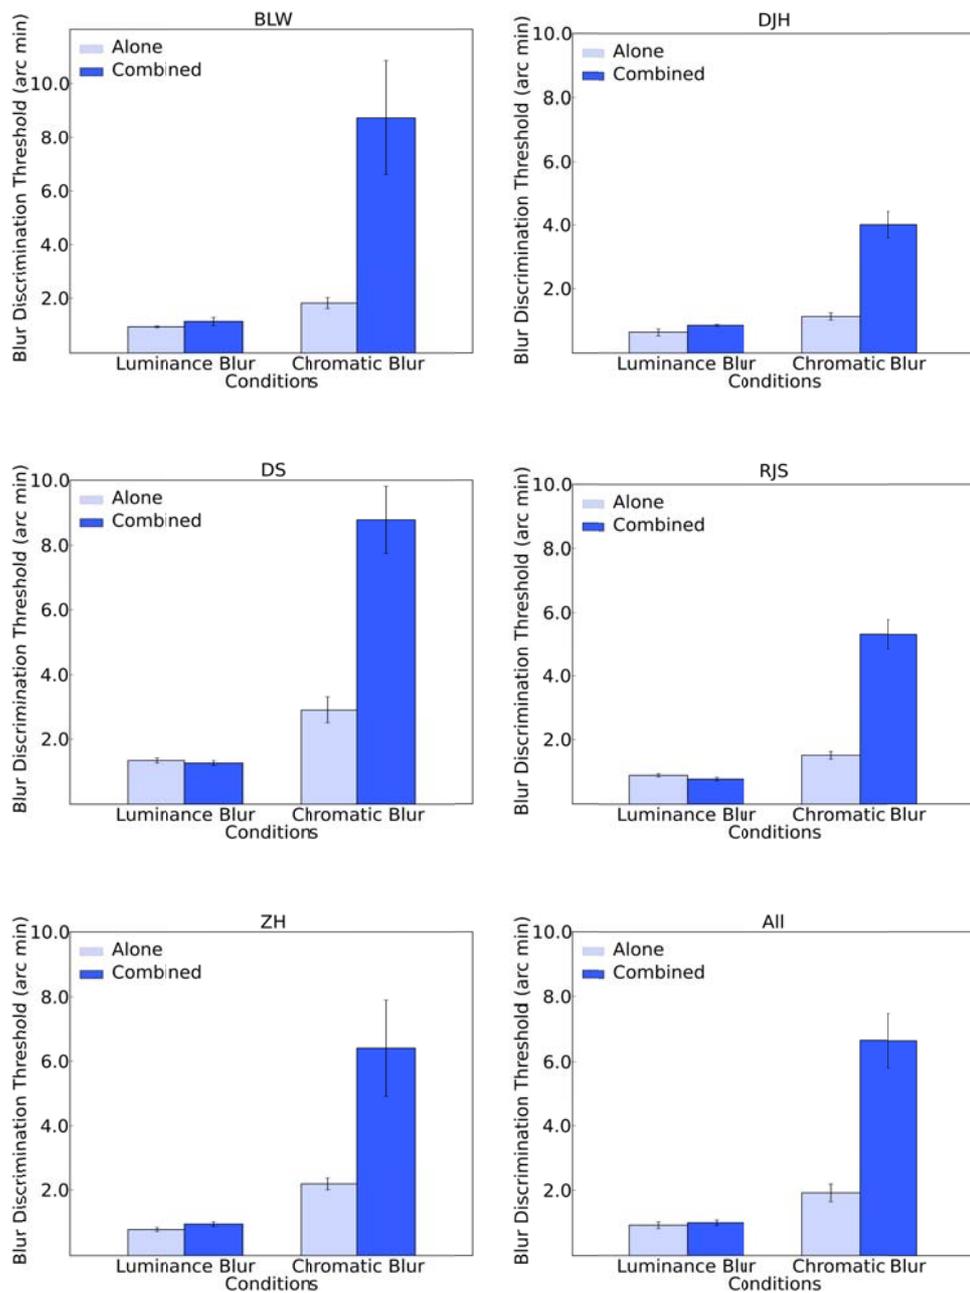
which appeared more blurred. The presentation order of the target and foil was randomised and the staircases were interleaved.

The staircase procedures used to measure blur threshold for each condition were the same as those used in Chapter 3. Each participant collected five staircases for each condition.

## 7.2. Results

Due to variability across the staircases it was not possible perform a fitting procedure on the data (see Section 2.11 for details). Instead, the mean of the last six reversals for each staircase was taken as a measure of perturbation threshold. The mean of these thresholds was then take to give one value per participant, per conditions i.e. 20 values overall.

The group data shown in Figure 7.2 exactly replicate the findings of Chapter 3. A two-way ANOVA showed that observers had higher blur discrimination thresholds for chromatic than for luminance information, main effect of channel type ( $F_{(1, 96)} = 73.118$ ,  $p < 0.001$ ,  $MS_{\text{channel}} = 0.018$ ). The elevated thresholds for chromatic blur were once again more pronounced in the presence of sharp luminance information (interaction between channel and combination;  $F_{(1, 96)} = 34.845$ ,  $p < 0.001$ ,  $MS_{\text{interaction}} = 0.009$ ).



**Figure 7.2.** The blur discrimination threshold for both combined (dark blue columns) and single channel (light blue columns) for each participant and across the group. All participants show masking of chromatic blur by sharp luminance information. Error bars represent  $\pm 1$  standard error of the mean.

### 7.3. Discussion

The results clearly replicate the findings of Chapter 3 where chromatic blur is masked by sharp luminance information. Therefore, the masking effect found in Chapter 3 is not a feature of natural scenes; it can also be generated using bipartite edges. This suggests that the difference in the findings of Chapters 3 and 5 are caused by differences between the blur discrimination and edge localisation tasks.

F.A.A. Kingdom (personal communication, 16 May 2013) suggests that we do not have a 'sense' of blur, but that blur *discrimination* can be achieved using low-level mechanisms sensitive to spatial-frequency content or edge-width. The current findings suggest that this low-level chromatic blur mechanism is modulated by luminance information; chromatic blur can be detected in isolation, but when sharp luminance information is present this dominates the percept.

The current results may be explained by distinguishing between edge localisation and blur discrimination tasks. Chapter 5 suggests that chromatic information has more influence in edge localisation tasks and the current results suggest that luminance information dominates in blur discrimination tasks. Similarly, the results of Chapter 5 suggest that we do have an explicit representation of chromatic *edges* but, Kingdom's (personal communication, 16 May 2013) results suggests that we do not have an explicit representation of chromatic *blur*.

Blur can change how size and distance are perceived and direct attention to certain areas of an image. Held and colleagues (Held, Cooper, O'Brien, & Banks, 2010) found that, whilst neither blur nor perspective cues alone are sufficient to estimate distance, they are effective depth cues when considered together. It may be therefore, that blur is generally used by the visual system as a cue for depth rather than edges. This could explain why colour is represented differently in these tasks; colour may be used differently in depth perception tasks than edge localisation tasks.

Scale-space analysis suggests that edge and blur information are jointly encoded. However, the current results suggest that this is not the case for chromatic information. The scale-space analysis concept relies on a spatially low-pass filter and a derivative (spatially-opponent) filter. Both types of filter are known to exist responding to achromatic stimuli. For colour processing, however, only the low-pass (spatially non-specific) filter is known to exist. The double-opponent (chromatic and spatially opponent) V1 cell reported by some groups is rarely tuned to the elevations close to the isoluminant plane and might not be actively involved in the percept of colour (see Shapley & Hawken, 2011 for a review). If it is the case that chromatic channels are not processed in a derivative manner, as required by Georgeson's scale-space analysis, then this might well explain why the blur percept in chromatic channels is poor and, potentially, resulting from an entirely different computation. Alternatively, scale-space analysis may be applied to chromatic signals, but the blur information generated from this may not be used for edge localisation. In order to investigate these

possibilities, predictions using the scale-space model should be tested on isoluminant profiles, in the same way that they have been for achromatic profiles.

In conclusion, chromatic blur is constrained by sharp luminance information not only in natural scenes but also bipartite edges. This demonstrates that the luminance dominance found in Chapter 3 is not caused by the structure of natural scenes. However, luminance dominance was not found in the edge *localisation* tasks used in Chapter 5, suggesting that it is specific to blur discrimination tasks.

## 8. General Discussion

This thesis has investigated cue combination of colour and luminance in edge detection using a series of psychophysical experiments. Cue combination of colour and luminance is a fundamental part of vision and due to different physiological and ecological constraints it is not clear how this combination occurs.

As discussed in Section 1.2, it is theoretically impossible for a cortical mechanism to perform both precise spatial sampling and precise chromatic sampling and this may explain poor performance in chromatic stereopsis tasks (Peirce, et al., 2008). Peirce et al (2008) suggest that binocular depth perception and binocular chromatic surface perception are controlled by different mechanisms. Similarly, it might be that, in form processing, edge detection and surface perception are performed by different mechanisms. If so then chromatic information could be less useful for edge localisation tasks and so luminance information should dominate.

However, it is not simply constraints from the visual system that must be considered when combining colour and luminance. In natural scenes, most purely achromatic edges represent shadows which do not offer information about the edges of objects, but are potentially more easily detectable by the visual system. In this case it would seem advantageous predominantly to use chromatic information in identifying and locating edges.

The experiments in this thesis have been designed to try and understand how these conflicting potential pressures interact and how the

visual system combines colour and luminance in different types of edge detection and localisation tasks.

### 8.1. Summary of findings

Chapter 3 investigated why chromatic blur is virtually imperceptible in the presence of sharp luminance information, despite being detectable in isolation. We tested various candidate sources of such effects and found that they cannot be attributed to poorer acuity of the chromatic channels, nor to the higher effective contrast of luminance information in natural scenes, nor to differences in the natural scene statistics of colour and luminance information. It appears that there is a mechanism actively promoting luminance information.

Chapter 4 considered whether the presence of both colour and luminance information improved edge localisation performance. This chapter used a staircase procedure to determine the just noticeable difference of edge localisation judgements and this was used as a measurement of performance. The small differences between edge localisation performance for isoluminant and achromatic edges meant that model predictions were not sufficiently different and the models could not be distinguished.

Chapter 5 also tried to address the question of whether the presence of both colour and luminance improves edge localisation performance, this time using a method of adjustment. Unfortunately, this method was also not sensitive enough to differentiate the models. However, this method also allowed us to measure edge localisation judgements in a new case; when the

colour and luminance edges conflicted. Chromatic information was found to be more heavily weighted by participants in the task than would be predicted by maximum likelihood estimation, suggesting a Bayesian prior actively promoting chromatic information.

It may be that the spatial relationship between colour and luminance affects how they are combined. Chapter 6 addressed this question using perturbation detection tasks. In Experiment 6.1 the introduction of a chromatic mask to a luminance target or a luminance mask to a chromatic target, in either aligned or orthogonal orientations was found to have no effect on the perception of perturbation. This suggests that in this task irrelevant information can be disregarded. In Experiment 6.2, when combinations of lines and square-waves were investigated it was found that the introduction of a chromatic-square-wave mask improved perturbation discrimination in luminance-lines. This suggests that the type of information and the spatial relationship between channels can both change how colour and luminance are combined.

The conflicting findings of Chapters 3 and 5 led to the question of whether this was driven by task differences or differences in stimuli. In Chapter 7, the blur discrimination task from Chapter 3 was performed on bipartite edges similar to those used in Chapters 4 and 5. Chromatic blur was still masked by sharp luminance information even with the far simpler stimuli. This demonstrates that this effect is not confined to natural scenes and

suggests mechanisms for blur discrimination preferentially use luminance information.

## 8.2. Implications of findings

Chromatic blur is constrained by sharp luminance information, but as suggested in Chapter 7, blur detection is not necessarily related to edge localisation. Another way to consider the constraint of chromatic blur by sharp luminance information is in terms of whether chromatic information can become ‘tied’ to luminance information such that it appears to follow the same contours as the luminance information, and fill in between them. Experiment 6.2 suggests that in the case of chromatic edges and luminance lines this is exactly what happens. This is in keeping with the fact that luminance outlines (Cole, et al., 1990) or reticles (Montag, 1997) facilitate detection of a chromatic target.

This suggests an underlying mechanism that links luminance lines and chromatic edges. It may be that this process occurs very early in the visual system leading the two pieces of information to remain tied as they progress through the different processing streams. It could potentially serve to increase the perceived contrast of chromatic information under these circumstances; improving perception of surface information. Horwitz et al (2005) found V1 neurons in macaques that responded to opposite-sign input from the S-cones and a rectified non-opponent signal from the L- and M-cones. These non-linear receptive fields responded most strongly to sharp

edges, relatively low spatial-frequency chromatic information and relatively high spatial-frequency luminance information (Horwitz, et al., 2005).

Whilst aspects of Experiment 6.2 support the idea of the perceived location of chromatic edges becoming tied to the location of luminance lines, other aspects suggest that chromatic information can have a greater influence on the perception of luminance information than vice versa. The introduction of a luminance-square-wave mask to a chromatic-line target had a smaller masking effect than introducing a chromatic-line mask to a luminance-square-wave target for the majority of participants. This might mean that chromatic information has more influence on perception of edge location, when the contrasts have been equated.

Blur discrimination tasks do not directly address edge *localisation*; they do not ask 'where is the edge?'. This could be why the findings in Experiment 5.4 show a different pattern to that described above, where chromatic information has more influence than would be predicted by its reliability. However, luminance information has a higher effective contrast in natural scenes and it is likely to drive perception in most circumstances.

It seems clear that the manner in which colour and luminance information is combined is task dependent. In blur discrimination tasks luminance dominates, even when contrast is equated. In edge localisation tasks chromatic information has more influence than would be predicted by cue reliability. In perturbation discrimination tasks there is a more

complicated relationship based on how the chromatic and luminance information is presented.

### 8.3. Relationship to previous literature

#### 8.3.1. Task Dependency

Traditionally luminance has been considered dominant for processing of lines, edges, shape and motion, with colour being more important for image segmentation (e.g. Martinovic, Mordal, & Wuerger, 2011). Evidence for this comes, in particular, from poor performance in chromatic stereopsis tasks (e.g. Krauskopf & Forte, 2002). Chromatic processing has also been found to be coarser for orientation and spatial frequency discrimination (Webster, De Valois, & Switkes, 1990).

However, chromatic performance is not poorer than luminance for all tasks. Krauskopf and Forte (2002) compared the influence of chromaticity on Vernier and stereo acuity using the same stimuli and apparatus. They found the expected deficits in chromatic stereo processing, but they found that Vernier thresholds were approximately equal for isoluminant and achromatic targets, if they were equated for spatial frequency content and contrast. This clearly shows that the utility of chromatic information is task dependent.

In this thesis chromatic blur discrimination has been shown to be poorer than luminance blur discrimination (Chapter 3) and in accordance with previous work Vernier thresholds are not poorer (Chapters 4 and 5). The thresholds for perturbation discrimination are less clear as a significant effect of chromaticity was found in Experiment 6.1 but not Experiment 6.2. It seems

most likely that the effect of chromaticity found in Experiment 6.1 was an artefact driven by the variability between participants, particularly for L-M targets. Interleaving the staircases in Experiment 6.2 reduced variability overall and the effect of chromaticity was no longer found, supporting the idea that the difference found in Experiment 6.1 was an artefact. If it is an artefact, chromatic performance in perturbation detection is not poorer than luminance.

Vernier acuity tasks, and therefore edge localisation seem to be subserved by a different chromatic mechanism than other spatial tasks. This may be related to the manner in which information coming from the randomly-arranged cone mosaic is processed by the early visual system. As described in Section 1.2 the maximum spatial sensitivity required for edge localisation can be achieved by subtracting the response of the cones to the left of a boundary from the response of the cones to the right of a boundary. The random arrangement of the cones means that there will be different numbers of L- and M-cones on either side. Therefore, subtracting one area from another will produce not only spatial information but also chromatic information. The additional availability of chromatic information may explain why Vernier acuity is not poorer for isoluminant targets.

Chromatic information is a more reliable indicator of object edges than luminance information (K. K. De Valois & Switkes, 1983; Kingdom, 2003; Kingdom, et al., 2004; McGraw, et al., 2003; Parraga, Troscianko, & Tolhurst, 2002; Ruderman, et al., 1998; Switkes, et al., 1988). The reason for this

becomes clear when considering a woodland floor where the luminance information is dominated by variations in illumination, which becomes 'noise' when performing an object edge localisation task. Colour can also unmask dark targets in complex displays (Kingdom & Kasrai, 2006) and facilitate shadow identification (Kingdom, et al., 2004). This suggests that chromatic information is not simply more useful for object edge localisation but also aids detection and discrimination of shadows. This is important as, whilst shadows act as 'noise' for object edge localisation, they are useful for other visual tasks (for a review see Dee & Santos, 2011) and in order for them to be used as cues they must be segmented from the background and labelled (Mamassian, et al., 1998).

It has been suggested that the use of colour variations to help identify lighting variations is based on the heuristic that achromatic edges are shadows and combined colour and luminance edges are object edges (Cavanagh, 1991; Kingdom, 2008). This identification then allows the chromatic information to suppress the luminance noise (Kingdom & Kasrai, 2006). It is unlikely that colour vision evolved primarily for this purpose and may be a secondary use of the chromatic system (Kingdom, et al., 2004).

### 8.3.2. Specific combinatorial rules

The relationship between colour and luminance is complex. In cross-channel contrast detection experiments there is a clear asymmetry in how the pedestal affects perception of the test; luminance pedestals, including only a luminance ring, facilitate detection of a chromatic test but chromatic

pedestals do not facilitate detection of luminance tests (Chen, Foley, & Brainard, 2000; Cole, et al., 1990, see Section 1.3 for more details). Similarly in cross-channel masking experiments, luminance gratings do not mask chromatic gratings but, chromatic gratings mask luminance gratings to a degree similar to that of luminance-luminance masking (K. K. De Valois & Switkes, 1983); also demonstrating the asymmetry of colour and luminance interactions.

The colour-shading effect offers an elegant example of the complexities of colour and luminance interactions. This illusory depth effect is triggered when non-aligned colour and luminance gratings are combined, however, it is suppressed when the cues are aligned (Kingdom, 2003); an effect of spatial arrangement. However, an attempt to replicate this illusion of depth failed for most participants, as there were large and inconsistent individual differences (Clery, et al., 2013). This challenges Kingdom's (2008) suggestion that the visual system uses a universal set of heuristics to distinguish light versus material changes and determine how colour and luminance information should be combined. In particular the idea that achromatic edges should be treated as changes in illumination and combined colour and luminance edges should be treated as changes in reflectance. Clery et al (2013) suggested that observers may use idiosyncratic heuristics formed on the basis of individual experience. This idea could add a new level of complexity to understanding cue combination of colour and luminance.

## 8.4. Unanswered questions

### 8.4.1. Can natural scene statistics be used inform models of edge localisation?

There is a growing body of research into the statistics of natural images and the same is true for cue combination. It would be nice to see more research that combines these two areas. For example, the findings of Experiment 5.4 show that colour is weighted more than would be predicted by maximum likelihood estimation, which might be explained by natural scenes statistics. If this were the case one might be able to model it by calculating the number of luminance edges that are edges of objects, versus the number of chromatic edges that are edges of objects. This ratio could then be used to scale the weights, which might form the ideal Bayesian prior for the higher weighting of chromatic edge information.

It will be necessary to differentiate between different categories to take account of different statistical features of different natural scenes. For example, images of man-made structures may contain more horizontal and vertical edges than images of flowers.

### 8.4.2. Can cue weights be changed through perceptual learning?

As previously discussed in blur discrimination tasks luminance appears to dominate and edge localisation tasks chromatic information has more of an influence than predicted by cue reliability. Are these differences 'hard-wired' into the visual system or are they learnt from exposure to the natural world? If they are learnt can they be changed through perceptual learning? The 'light

from above' prior can be modified through training (Adams, Graf & Ernst 2004), which may suggest that the relative weights of chromatic and luminance edges can also be modified.

One way to investigate this would be to manipulate the statistics of stimuli such that luminance is more reliable than chromatic information. If perceptual learning occurred and exposure to this type of stimulus causes changes in cue weightings this would suggest that the weightings found in Experiment 5.4 are learnt. A slightly different approach would be to see if participants can be trained to disregard an irrelevant cue. For example, can the masking effects found in Experiment 6.2 be reduced or extinguished with practice?

#### 8.4.3. When do aligned and orthogonal arrangements of colour and luminance cause perceptual changes?

It also remains unclear what causes the conflicting findings of Kingdom's lab (2003; 2010; 2005b), Clery and colleagues (2013) and Experiment 6.1 of this thesis, as regards the effect of aligned versus orthogonal arrangements of colour and luminance. It could be that the differences in findings are due to individual differences. However, there are significant methodological differences between the experiments that could also offer an explanation. As we have seen, colour and luminance combinations can be very sensitive to changes in task or stimulus parameters and establishing the source of the conflicting findings could offer new insight into combinatorial rules.

Another approach to the question of whether orthogonal edges are processed differently to aligned edges would be to perform the blur discrimination task on natural scenes where either the chromatic or luminance layer has been rotated 90°. In natural scenes a majority of edges are comprised of combined (aligned) colour and luminance information (Hansen & Gegenfurtner, 2009) meaning that in the experiments in Chapter 3, a majority of edges are aligned and as discussed in that chapter it is likely this is crucial to the constraint of chromatic blur by luminance information. If a layer were rotated this relationship would no longer exist and if Kingdom's (2008) heuristics are correct the masking effect would disappear or be reversed. This would help to determine whether the alignment of edges is important not just for edge localisation tasks but also blur discrimination tasks.

#### 8.4.4. What is the S-(L+M) channel used for?

In this thesis, and a considerable amount of the literature, the research focus is on the L-M channel rather than the S-(L+M) channel. The L-M system appears to be optimised for detecting reddish objects against a background of foliage and eliminates the dappled background of leaves, however the S-(L+M) system does not allow this discounting of lighting information (Parraga, et al., 2002). However, the questions remain as to what the S-(L+M) channel is optimised for and whether the limits in its performance are due to constraints such as sparse representation in the cone mosaic or chromatic aberration.

The sparseness of S-cones in the retina should mean that acuity is always lower for this type of stimulus. However, performance is not significantly different from L-M defined stimuli in supra-threshold orientation discrimination, spatial frequency discrimination or Vernier judgements (Krauskopf & Farell, 1991; Webster, et al., 1990). Blur thresholds, on the other hand, are approximately twice as high for blue-yellow gratings than red-green or luminance gratings (Wuerger, et al., 2001). Interestingly, Wuerger et al. (2001) found that red-green and luminance gratings had very similar blur discrimination thresholds. This may mean that the higher thresholds for chromatic stimuli found in Chapter 3 were driven by the inclusion of S-cone stimulating information.

### 8.5. Conclusion

The relationship between colour and luminance is complex. There is no single rule that can predict how the two cues will be combined. Cue combination of colour and luminance is task dependent and modulated by the form the cues take and their spatial arrangements.

Knowing how colour and luminance are combined is fundamental to understanding visual processing of form. It is essential to understand this relationship before plausible models of higher-level processes can be created. Colour and luminance research can also inform computer vision models, allowing improved detection of features by allowing inclusion of colour and heterogeneous illumination.

## 9. References

- Adams, W.J., Graf, E.W. & Ernst, M.O. (2004). Experience can change the 'light from above' prior. *Nature Neuroscience*, 7, 1057-1058.
- Anderson, S. J., Mullen, K. T., & Hess, R. F. (1991). Human peripheral spatial-resolution for achromatic and chromatic stimuli - limits imposed by optical and retinal factors. *Journal of Physiology-London*, 442, 47-64.
- Anstis, S., & Cavanagh, P. (1983). A minimum motion technique for judging equiluminance. In J. D. Mollon & L. T. Sharpe (Eds.), *Colour Vision: Physiology and Psychophysics* (pp. 156-166). London: Academic.
- Bartels, A., & Zeki, S. (2000). The architecture of the colour centre in the human visual brain: new results and a review. *European Journal of Neuroscience*, 12(1), 172-190.
- Bowers, J. S., & Davis, C. J. (2012). Bayesian Just-So Stories in Psychology and Neuroscience. *Psychological Bulletin*, 138(3), 389-414.
- Boynton, R. M., Hayhoe, M. M., & Macleod, D. I. A. (1977). Gap effect - chromatic and achromatic visual-discrimination as affected by field separation. *Optica Acta*, 24(2), 159-177.
- Brouwer, G. J., & Heeger, D. J. (2009). Decoding and Reconstructing Color from Responses in Human Visual Cortex. *Journal of Neuroscience*, 29(44), 13992-14003.
- Burr, D. C., Morrone, M. C., & Spinelli, D. (1989). Evidence for edge and bar detectors in human-vision. *Vision Research*, 29(4), 419-431.

- Callaway, E. M. (1998). Local circuits in primary visual cortex of the macaque monkey. *Annual Review of Neuroscience*, 21, 47-74.
- Casagrande, V. A. (1994). A 3rd parallel visual pathway to primate area - V1. *Trends in Neurosciences*, 17(7), 305-310.
- Cavanagh, P. (1991). *Vision at equiluminance* (Vol. V: Limits of vision). Boca Raton, FL: CRC Press.
- Chen, C. C., Foley, J. M., & Brainard, D. H. (2000). Detection of chromoluminance patterns on chromoluminance pedestals I: threshold measurements. *Vision Research*, 40(7), 773-788.
- Clery, S., Bloj, M., & Harris, J. M. (2013). Interactions between luminance and color signals: effects on shape. *J Vis*, 13(5).
- Cole, G. R., Stromeyer, C. F., & Kronauer, R. E. (1990). Visual interactions with luminance and chromatic stimuli. *Journal of the Optical Society of America a-Optics Image Science and Vision*, 7(1), 128-140.
- de Monasterio, F. M., & Schein, S. J. (1982). Spectral bandwidths of color-opponent cells of geniculocortical pathway of macaque monkeys. *Journal of Neurophysiology*, 47(2), 214-224.
- De Valois, K. K., & Switkes, E. (1983). Simultaneous Masking Interactions between Chromatic and Luminance Gratings. *Journal of the Optical Society of America*, 73(1), 11-18.
- De Valois, R. L., & De Valois, K. K. (1988). *Spatial vision*. New York ; Oxford: Oxford University Press.

- Dee, H. M., & Santos, P. E. (2011). The Perception and Content of Cast Shadows: An Interdisciplinary Review. *Spatial Cognition and Computation, 11*(3), 226-253.
- Derrington, A. M., Krauskopf, J., & Lennie, P. (1984). Chromatic mechanisms in lateral geniculate-nucleus of macaque. *Journal of Physiology-London, 357*(DEC), 241-265.
- Fahle, M., & Edelman, S. (1993). Long-Term Learning in Vernier Acuity - Effects of Stimulus Orientation, Range and of Feedback. *Vision Research, 33*(3), 397-412.
- Felsen, G., & Dan, Y. (2005). A natural approach to studying vision. *Nature Neuroscience, 8*(12), 1643-1646.
- Foster, D. H. (2011). Color constancy. *Vision Research, 51*(7), 674-700.
- Gegenfurtner, K. R., & Kiper, D. C. (2003). Color vision. *Annual Review of Neuroscience, 26*, 181-206.
- Gegenfurtner, K. R., Kiper, D. C., & Fenstemaker, S. B. (1996). Processing of color, form, and motion in macaque area V2. *Visual Neuroscience, 13*(1), 161-172.
- Gegenfurtner, K. R., & Rieger, J. (2000). Sensory and cognitive contributions of color to the recognition of natural scenes. *Current Biology, 10*(13), 805-808.

- Georgeson, M. A., May, K. A., Freeman, T. C. A., & Hesse, G. S. (2007). From filters to features: Scale-space analysis of edge and blur coding in human vision. *Journal of Vision, 7*(13), 21.
- Gold, J. I., & Shadlen, M. N. (2001). Neural computations that underlie decisions about sensory stimuli. *Trends in Cognitive Sciences, 5*(1), 10-16.
- Gowdy, P. D., Stromeyer, C. F., & Kronauer, R. E. (1999). Facilitation between the luminance and red-green detection mechanisms: enhancing contrast differences across edges. *Vision Research, 39*(24), 4098-4112.
- Griffiths, T. L., Chater, N., Norris, D., & Pouget, A. (2012). How the Bayesians Got Their Beliefs (and What Those Beliefs Actually Are): Comment on Bowers and Davis (2012). *Psychological Bulletin, 138*(3), 415-422.
- Gur, M., & Akri, V. (1992). Isoluminant stimuli may not expose the full contribution of color to visual functioning - spatial contrast sensitivity measurements indicate interaction between color and luminance processing. *Vision Research, 32*(7), 1253-1262.
- Hansen, T., & Gegenfurtner, K. F. (2009). Independence of color and luminance edges in natural scenes. *Visual Neuroscience, 26*(1), 35-49.
- Hansen, T., & Gegenfurtner, K. F. (2013). The contribution of colour to detect edges in natural scenes. *Journal of Vision, 13*(9), 1010.

- Held, R. T., Cooper, E. A., O'Brien, J. F., & Banks, M. S. (2010). Using Blur to Affect Perceived Distance and Size. *Acm Transactions on Graphics*, 29(2).
- Heywood, C. A., Gaffan, D., & Cowey, A. (1995). Cerebral achromatopsia in monkeys. *European Journal of Neuroscience*, 7(5), 1064-1073.
- Hillis, J. M., Watt, S. J., Landy, M. S., & Banks, M. S. (2004). Slant from texture and disparity cues: Optimal cue combination. *Journal of Vision*, 4(12), 967-992.
- 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., & Livingstone, M. S. (1987). Segregation of form, color, and stereopsis in primate area-18. *Journal of Neuroscience*, 7(11), 3378-3415.
- Hubel, D. H., & Wiesel, T. N. (1968). Receptive fields and functional architecture of monkey striate cortex. *Journal of Physiology-London*, 195(1), 215-&
- Hunt, R. W. G. (2004). *The Reproduction of Colour* (6th ed.). New York: Wiley-Blackwell.
- Ingling, C. R., & Martinezuriegas, E. (1983). The Relationship between Spectral Sensitivity and Spatial Sensitivity for the Primate R-G X-Channel. *Vision Research*, 23(12), 1495-1500.

- Isono, H., Sakata, H., & Kusaka, H. (1978). Subjective evaluation of apparent reduction of chromatic blur depending on luminance signals. *Ieee Transactions on Broadcasting*, 24(4), 107-115.
- Jacobs, G. H., & Deegan, J. F. (1999). Uniformity of colour vision in Old World monkeys. *Proceedings of the Royal Society of London Series B-Biological Sciences*, 266(1432), 2023-2028.
- Johnson, A. P., Kingdom, F. A., & Baker, C. L., Jr. (2005). Spatiochromatic statistics of natural scenes: first- and second-order information and their correlational structure. *J Opt Soc Am A Opt Image Sci Vis*, 22(10), 2050-2059.
- 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.
- Kaiser, P. K. (1996). The Joy of Visual Perception Retrieved 05/04/2011, from <http://www.yorku.ca/eye/toc.htm>
- Kersten, D., & Yuille, A. (2003). Bayesian models of object perception. *Current Opinion in Neurobiology*, 13(2), 150-158.
- Kingdom, F. A. A. (2003). Color brings relief to human vision. *Nature Neuroscience*, 6(6), 641-644.
- Kingdom, F. A. A. (2008). Perceiving light versus material. *Vision Research*, 48(20), 2090-2105.

- Kingdom, F. A. A., Beauce, C., & Hunter, L. (2004). Colour vision brings clarity to shadows. *Perception, 33*(8), 907-914.
- Kingdom, F. A. A., Bell, J., Gheorghiu, E., & Malkoc, G. (2010). Chromatic variations suppress suprathreshold brightness variations. *Journal of vision, 10*(10).
- Kingdom, F. A. A., & Kasrai, R. (2006). Colour unmask dark targets in complex displays. *Vision Research, 46*(6-7), 814-822.
- Kingdom, F. A. A., Rangwala, S., & Hammamji, K. (2005a). Chromatic properties of the colour-shading effect. *Vision Research, 45*(11), 1425-1437.
- Kingdom, F. A. A., Rangwala, S., & Hammamji, K. (2005b). Chromatic properties of the colour-shading effect. *Vision Res, 45*(11), 1425-1437.
- Kiper, D. C., Fenstemaker, S. B., & Gegenfurtner, K. R. (1997). Chromatic properties of neurons in macaque area V2. *Visual Neuroscience, 14*(6), 1061-1072.
- Knill, D. C., & Pouget, A. (2004). The Bayesian brain: the role of uncertainty in neural coding and computation. *Trends in Neurosciences, 27*(12), 712-719.
- Krauskopf, J., & Farell, B. (1991). Vernier acuity - effects of chromatic content, blur and contrast. *Vision Research, 31*(4), 735-749.
- Krauskopf, J., & Forte, J. D. (2002). Influence of chromaticity on Vernier and stereo acuity. *Journal of Vision, 2*(9), 8.

- Krauskopf, J., Williams, D. R., & Heeley, D. W. (1982). Cardinal directions of color space. *Vision Research*, 22(9), 1123-1131.
- Krauskopf, J., Wu, H. J., & Farell, B. (1996). Coherence, cardinal directions and higher-order mechanisms. *Vision Research*, 36(9), 1235-1245.
- Kulikowski, J. J., Walsh, V., McKeefry, D., Butler, S. R., & Carden, D. (1994). The electrophysiological basis of color processing in macaques with V4 lesions. *Behavioural Brain Research*, 60(1), 73-78.
- Lennie, P., Krauskopf, J., & Sclar, G. (1990). Chromatic mechanisms in striate cortex of macaque. *Journal of Neuroscience*, 10(2), 649-669.
- Leventhal, A. G., Thompson, K. G., Liu, D., Zhou, Y. F., & Ault, S. J. (1995). Concomitant Sensitivity to Orientation, Direction, and Color of Cells in Layer-2, Layer-3, and Layer-4 of Monkey Striate Cortex. *Journal of Neuroscience*, 15(3), 1808-1818.
- Lewis, A., & Zhaoping, L. (2006). Are cone sensitivities determined by natural color statistics? *Journal of Vision*, 6(3), 285-302.
- Logothetis, N.K. & Charles, E.R. (1990) The minimum motion technique applied to determine isoluminance in psychophysical experiments with monkeys. *Vision Research*, 30(6), 829-838.
- 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.

- Mamassian, P., Knill, D. C., & Kersten, D. (1998). The perception of cast shadows. *Trends in Cognitive Sciences*, 2(8), 288-295.
- Martinovic, J., Mordal, J., & Wuerger, S. M. (2011). Event-related potentials reveal an early advantage for luminance contours in the processing of objects. *Journal of Vision*, 11(7).
- McGraw, P. V., Whitaker, D., & Badcock, D. R. (2000). Localising conflicting visual attributes. *Investigative Ophthalmology & Visual Science*, 41(4), 4269B4216.
- McGraw, P. V., Whitaker, D., Badcock, D. R., & Skillen, J. (2003). Neither here nor there: localizing conflicting visual attributes. *Journal of vision*, 3(4), 265-273.
- McIlhagga, W. H., & Mullen, K. T. (1996). Contour integration with colour and luminance contrast. *Vision Research*, 36(9), 1265-1279.
- Meer, P., & Zeevi, Y. Y. (1986). Importance of Global Information in Vernier Acuity. *Journal of the Optical Society of America a-Optics Image Science and Vision*, 3(6), 880-884.
- Mollon, J. D. (1989). Tho she kneeld in that place where they grew ... The uses and origins of primate color-vision. *Journal of Experimental Biology*, 146, 21-38.
- Montag, E. D. (1997). Influence of boundary information on the perception of color. *Journal of the Optical Society of America a-Optics Image Science and Vision*, 14(5), 997-1006.

- Mullen, K. T. (1985). The contrast sensitivity of human color-vision to red green and blue yellow chromatic gratings. *Journal of Physiology-London*, 359(FEB), 381-400.
- Mullen, K. T., & Beaudot, W. H. A. (2002). Comparison of color and luminance vision on a global shape discrimination task. *Vision Research*, 42(5), 565-575.
- Mullen, K. T., Beaudot, W. H. A., & McIlhagga, W. H. (2000). Contour integration in color vision: a common process for the blue-yellow, red-green and luminance mechanisms? *Vision Research*, 40(6), 639-655.
- Newton, J. R., & Eskew, R. T. (2003). Chromatic detection and discrimination in the periphery: A postreceptoral loss of color sensitivity. *Visual Neuroscience*, 20(5), 511-521.
- Oho, E., & Watanabe, M. (2001). Natural color scanning electron microscopy based on the frequency characteristics of the human visual system. *Scanning*, 23(1), 24-31.
- Olmos, A., & Kingdom, F. A. A. (2004). A biologically inspired algorithm for the recovery of shading and reflectance images. *Perception*, 33(12), 1463-1473.
- Parraga, C. A., Brelstaff, G., Troscianko, T., & Moorehead, I. R. (1998). Color and luminance information in natural scenes. *Journal of the Optical Society of America a-Optics Image Science and Vision*, 15(3), 563-569.

- Parraga, C. A., Troscianko, T., & Tolhurst, D. J. (2002). Spatiochromatic properties of natural images and human vision. *Current Biology*, *12*(6), 483-487.
- Peirce, J. W. (2007). PsychoPy - Psychophysics software in Python. *Journal of Neuroscience Methods*, *162*, 8-13.
- Peirce, J. W., Solomon, S. G., Forte, J. D., & Lennie, P. (2008). Cortical representation of color is binocular. *Journal of Vision*, *8*(3), 10.
- Pelli, D. G., & Farell, B. (1999). Why use noise? *Journal of the Optical Society of America a-Optics Image Science and Vision*, *16*(3), 647-653.
- Pinna, B., Brelstaff, G., & Spillmann, L. (2001). Surface color from boundaries: a new 'watercolor' illusion. *Vision Research*, *41*(20), 2669-2676.
- Platt, M. L., & Glimcher, P. W. (1999). Neural correlates of decision variables in parietal cortex. *Nature*, *400*(6741), 233-238.
- Pouget, A., Dayan, P., & Zemel, R. S. (2003). Inference and computation with population codes. *Annual Review of Neuroscience*, *26*, 381-410.
- Rivest, J., & Cavanagh, P. (1996). Localizing contours defined by more than one attribute. *Vision Research*, *36*(1), 53-66.
- Ruderman, D. L., Cronin, T. W., & Chiao, C. C. (1998). Statistics of cone responses to natural images: implications for visual coding. *Journal of the Optical Society of America a-Optics Image Science and Vision*, *15*(8), 2036-2045.

- Rust, N. C., & Movshon, J. A. (2005). In praise of artifice. *Nature Neuroscience*, 8(12), 1647-1650.
- Sadowski, J. (2006). Big Spanish Castle Retrieved 04/05/2011, from [http://www.johnsadowski.com/big\\_spanish\\_castle.php](http://www.johnsadowski.com/big_spanish_castle.php)
- Sanger, T. D. (1996). Probability density estimation for the interpretation of neural population codes. *Journal of Neurophysiology*, 76(4), 2790-2793.
- Saunders, J. A., & Knill, D. C. (2004). Visual feedback control of hand movements. *Journal of Neuroscience*, 24(13), 3223-3234.
- Schiller, P. H. (1993). The effects of V4 and middle temporal (MT) area lesions on visual performance in the rhesus-monkey. *Visual Neuroscience*, 10(4), 717-746.
- Shapley, R., & Hawken, M. (2002). Neural mechanisms for color perception in the primary visual cortex. *Current Opinion in Neurobiology*, 12(4), 426-432.
- Shapley, R., & Hawken, M. J. (2011). Color in the Cortex: single- and double-opponent cells. *Vision Research*, 51(7), 701-717.
- Skottun, B. C., Bradley, A., Sclar, G., Ohzawa, I., & Freeman, R. D. (1987). The Effects of Contrast on Visual Orientation and Spatial-Frequency Discrimination - a Comparison of Single Cells and Behavior. *Journal of Neurophysiology*, 57(3), 773-786.

- Smith, V. C., & Pokorny, J. (1975). Spectral sensitivity of foveal cone photopigments between 400 and 500 nm. *Vision Research*, 15(2), 161-171.
- Sumner, P., Anderson, E. J., Sylvester, R., Haynes, J. D., & Rees, G. (2008). Combined orientation and colour information in human V1 for both L-M and S-cone chromatic axes. *Neuroimage*, 39(2), 814-824.
- Switkes, E., Bradley, A., & De Valois, K. K. (1988). Contrast Dependence and Mechanisms of Masking Interactions among Chromatic and Luminance Gratings. *Journal of the Optical Society of America a-Optics Image Science and Vision*, 5(7), 1149-1162.
- Syrkin, G., & Gur, M. (1997). Colour and luminance interact to improve pattern recognition. *Perception*, 26(2), 127-140.
- Tolhurst, D. J., Movshon, J. A., & Dean, A. F. (1983). The statistical reliability of signals in single neurons in cat and monkey visual-cortex. *Vision Research*, 23(8), 775-785.
- Wald, G. (1967). Blue-Blindness in Normal Fovea. *Journal of the Optical Society of America*, 57(11), 1289-1301.
- Walsh, V., Kulikowski, J. J., Butler, S. R., & Carden, D. (1992). The effects of lesions of area V4 on the visual abilities of macaques - color categorization. *Behavioural Brain Research*, 52(1), 81-89.
- Wandell, B. A. (1995). *Foundations of Vision*. Sunderland: Sinauer Associates, Inc.

- Watson, A.B. & Ahumada, A.J. (2011). Blur clarified: A review and synthesis of blur discrimination. *Journal of Vision*, 11(5), 10.
- Webster, M. A., De Valois, K. K., & Switkes, E. (1990). Orientation and Spatial-Frequency Discrimination for Luminance and Chromatic Gratings. *Journal of the Optical Society of America a-Optics Image Science and Vision*, 7(6), 1034-1049.
- Whiteley, L., & Sahani, M. (2008). Implicit knowledge of visual uncertainty guides decisions with asymmetric outcomes. *Journal of Vision*, 8(3), 15.
- Wozny, D. R., Beierholm, U. R., & Shams, L. (2010). Probability Matching as a Computational Strategy Used in Perception. *Plos Computational Biology*, 6(8).
- Wuerger, S. M., Morgan, M. J., Westland, S., & Owens, H. (2000). The spatio-chromatic sensitivity of the human visual system. *Physiological Measurement*, 21(4), 505-513.
- Wuerger, S. M., Owens, H., & Westland, S. (2001). Blur tolerance for luminance and chromatic stimuli. *Journal of the Optical Society of America a-Optics Image Science and Vision*, 18(6), 1231-1239.
- Zeki, S. (1983a). Color coding in the cerebral-cortex - the reaction of cells in monkey visual-cortex to wavelengths and colors. *Neuroscience*, 9(4), 741-765.

- Zeki, S. (1983b). Color coding in the cerebral-cortex - the responses of wavelength-selective and color-coded cells in monkey visual-cortex to changes in wavelength composition. *Neuroscience*, 9(4), 767-&.
- Zemel, R. S., Dayan, P., & Pouget, A. (1998). Probabilistic interpretation of population codes. *Neural Computation*, 10(2), 403-430.
- Zhou, C. H., & Mel, B. W. (2008). Cue combination and color edge detection in natural scenes. *Journal of Vision*, 8(4), 25.